MATHEMATICS OF COMPUTATION
Volume 71, Number 237, Pages 393–405
S 0025-5718(00)01297-7
Article electronically published on October 26, 2000

ARITHMETIC ON SUPERELLIPTIC CURVES

S. D. GALBRAITH, S. M. PAULUS, AND N. P. SMART

ABSTRACT. This paper is concerned with algorithms for computing in the divisor class group of a nonsingular plane curve of the form $y^n = c(x)$ which has only one point at infinity. Divisors are represented as ideals, and an ideal reduction algorithm based on lattice reduction is given. We obtain a unique representative for each divisor class and the algorithms for addition and reduction of divisors run in polynomial time. An algorithm is also given for solving the discrete logarithm problem when the curve is defined over a finite field.

1. INTRODUCTION

There is great interest in algorithms for computing in the divisor class group of an algebraic curve. Algorithms for general curves have been given by Coates [5], Huang and Ierardi [10] and Volcheck [22], but they tend not to be very suitable for practical implementation.

More efficient algorithms may be obtained if one restricts to a less general class of curves. The case of hyperelliptic curves (i.e., quadratic function fields) has been handled very successfully by Cantor [4] and Scheidler, Stein and Williams [18] (also see Paulus and Rück [16]). Some algorithms for cubic function fields have been given by Scheidler and Stein [19] and [20]. A specific case of genus three cubic extensions has been given by Barreiro, Cherdieu and Sarlabous [3] using geometric methods. An algorithm for general plane curves has also been recently given by Hess [9]. A method, using Groebner bases, for plane curves with a single rational point at infinity has been given by Arita [2]. This paper will give an algorithm for computing with the following class of curves, which is a subset of those considered by Arita.

Definition 1. Let k be a field. Let $c(x) \in k[x]$ be a monic polynomial of degree δ such that $\gcd(c(x), c'(x)) = 1$ (where $c'(x)$ is the formal derivative of $c(x)$). Let n be an integer such that $\gcd(n, \delta) = 1$ and $\gcd(n, \text{char } k) = 1$. Then the curve

$$C : y^n = c(x)$$

is called a *superelliptic curve.*

Received by the editor April 13, 1999 and, in revised form, March 17, 2000.

2000 *Mathematics Subject Classification.* Primary 14Q05, 14H40, 11G20, 11Y16.

Key words and phrases. Superelliptic curve, divisor class group, cryptography, discrete logarithm problem.

The work in this paper was carried out whilst the first author was supported by an EPSRC grant at Royal Holloway University of London, the second author was at Darmstadt University of Technology, and the third author was employed by Hewlett-Packard Laboratories.

If $c(x)$ is not monic, then (since $\gcd(n, \delta) = 1$) a change of variable can be made to obtain a monic equation. If n is odd, then all our results hold for fields of characteristic two. We tend to take $n, \delta \geq 3$ since otherwise the curve is hyperelliptic and can be handled by other means.

We show that there is a unique representative for each divisor class and we give algorithms for addition and reduction of divisors which run in polynomial time. The approach of this paper uses the representation of the divisor class group of the curve as an ideal class group. Our algorithm is the same as Cantor's algorithm [4] in the case $n = 2$, and for fixed n has complexity $O(g^2)$ just as Cantor's algorithm does. The main tool for larger values of n is a lattice reduction algorithm which provides an ideal reduction process similar to that used for computations with algebraic number fields.

Most of the results given in this paper apply to the more general case of a nonsingular plane curve with a single point at infinity. We make some comments about this at the appropriate places.

The paper is structured as follows. In Section 2 we recall some results about divisor class groups, ideal class groups, and superelliptic curves. In Section 3 we indicate the representation and computations on ideals which we borrow from the number field case. We also give a strategy for reducing ideals. In Section 4 we show how a variation of a lattice reduction algorithm due to A. Lenstra can be used as an ingredient for our ideal reduction method. In Section 5 we restrict to the case where k is a finite field and we modify the algorithm of Adleman, De Marrais and Huang [1] to obtain a heuristic method for solving the discrete logarithm problem in the divisor class group of a superelliptic curve in expected subexponential time.

2. Divisor class groups of superelliptic curves

Details of algebraic curves, their function fields, and divisor class groups can be found in [8], [21]. In particular we use the notation $\mathrm{Div}_k^0(C)$ for the set of degree zero divisors on the curve C which are defined over k, and the notation $\mathrm{Pic}_k^0(C)$ for the divisor class group $\mathrm{Div}_k^0(C)$ modulo principal divisors.

It is important to have a unique representative for each divisor class. The following result shows that such a representative exists for any curve C, over any field k, as long as there is a k-rational point on the curve.

Theorem 1. *Let C be a nonsingular curve over a field k of genus g with a given k-point, P_∞. Let $D \in \mathrm{Div}_k^0(C)$. Then there is a unique effective divisor over k of minimal degree $0 \leq m \leq g$ such that $E - mP_\infty$ is equivalent to D.*

Proof. If D is principal, then obviously $m = 0$ and $E = 0$. If D is not principal, then $l(D) = 0$. Consider the difference $l(D + (m+1)P_\infty) - l(D + mP_\infty) \geq 0$. The Riemann-Roch theorem shows that this difference is

$$l(\kappa - D - (m+1)P_\infty) + (m+1) + 1 - g - (l(\kappa - D - mP_\infty) + m + 1 - g)\,,$$

where κ is the canonical divisor. This difference is equal to $l(\kappa - D - (m+1)P_\infty) - l(\kappa - D - mP_\infty) + 1$. Now, $l(\kappa - D - (m+1)P_\infty) \leq l(\kappa - D - mP_\infty)$. It follows that the values of $l(D + mP_\infty)$ increase with m in steps of only 0 or 1.

Let m be the unique smallest positive integer such that $l(D + mP_\infty) = 1$ and let f be any nonzero function $f \in L(D + mP_\infty)$. Then for $E := (f) + D + mP_\infty$, one has that $E - mP_\infty = D + (f)$. The uniqueness of E is clear since $l(D + mP_\infty) = 1$. □

We now list some properties of superelliptic curves which we will require.

Proposition 2. *Let C be a superelliptic curve over a field k as in Definition 1. Then*

1. *C is nonsingular as an affine curve.*
2. *There is only one point, P_∞, at infinity on the normalisation of C and this point is defined over k.*
3. *The genus of C is $\frac{1}{2}(n-1)(\delta-1)$.*
4. *The integral closure of $k[x]$ in the function field $k(C)$ is*

$$\mathcal{O} := k[x,y]/(y^n - c(x)).$$

Proof. The proof of 1 is trivial, while 2 follows from repeatedly blowing up the point at infinity on the projective model. The genus can be calculated using the Hurwitz formula (see Fulton [8, pp. 8-36]). Finally, property 4 follows from Stichtenoth [21, III.5.12]. □

The next result is well known from the hyperelliptic case. It applies to any curve over k which has a single point at infinity that is k-rational.

Proposition 3. *Let C/k be a curve which has a single k-rational point on the normalization at infinity, let $\mathcal{O}$ denote the integral closure of $k[x]$ in $k(C)$, and let $\mathrm{Cl}(\mathcal{O})$ be the ideal class group of $\mathcal{O}$, then $\mathrm{Cl}(\mathcal{O}) \cong \mathrm{Pic}_k^0(C)$ as groups.*

Proof. It is standard to identify prime ideals of $\mathcal{O}$ with prime divisors on C which do not lie over the infinite place of $k[x]$. The mapping from a divisor $D = \sum_{\mathfrak{p}} n_{\mathfrak{p}}\mathfrak{p}$ to an $\mathcal{O}$-ideal $\mathfrak{a} = \prod_{\mathfrak{p} \not| \infty} \mathfrak{p}^{n_{\mathfrak{p}}}$ induces a group homomorphism from the divisor class group to the ideal class group.

Since there is only one point at infinity this map has trivial kernel. Also, since the point at infinity is k-rational, the cokernel is also trivial. □

Note that divisors of the form $E - mP_\infty$, where E is an effective divisor of degree m whose support does not contain P_∞, correspond to integral ideals of degree m. Therefore Theorem 1 implies that in every ideal class of $\mathcal{O}$ there is a unique integral ideal of minimal degree.

Definition 2. The unique integral ideal of minimal degree in an $\mathcal{O}$-ideal class is called the *reduced ideal* in the class.

The following sections are concerned with computations involving these ideals.

3. Arithmetic on ideals

3.1. **Representation of ideals.** It is necessary to have a good representation for the ideals under consideration. Following the number field case (see Cohen [6, Section 4.7]) we will represent integral $\mathcal{O}$-ideals as $k[x]$-modules in Hermite Normal Form (HNF). The details follow immediately from [6] so we merely state the final result.

Proposition 4. *Every integral $\mathcal{O}$-ideal $\mathfrak{a}$ can be represented as a $k[x]$-module with basis of the form*

$$\{a_{1,1}(x), a_{2,2}(x)y + a_{2,1}(x), \ldots, a_{n,n}(x)y^{n-1} + a_{n,n-1}(x)y^{n-2} + \cdots + a_{n,1}(x)\},$$

where the $a_{i,j}(x) \in k[x]$ and where $\deg a_{j,i}(x) < \deg a_{i,i}(x)$ for all $1 \le i < j \le n$. This basis is unique and it has the further property that $a_{i+1,i+1}(x) | a_{i,i}(x)$ for all

$1 \leq i < n$. *The norm of the ideal* $\mathfrak{a}$ *is* $N_{k(C)/k(x)}(\mathfrak{a}) = \prod_{i=1}^{n} a_{i,i}(x) \in k[x]$. *The degree of the ideal* $\mathfrak{a}$ *is* $\deg(\mathfrak{a}) = \deg_x(N_{k(C)/k(x)}(\mathfrak{a}))$.

Note that the divisor consisting of a k-rational point (x_0, y_0) corresponds to the ideal whose basis is $\{x - x_0, y - y_0, y^2 - y_0^2, \ldots, y^{n-1} - y_0^{n-1}\}$. Similarly, to calculate the points in the support of a divisor given the HNF ideal representation, one would first find the irreducible factors of the polynomials $a_{i,i}(x)$. The polynomial $a_{1,1}(x)$ determines the x-coordinates of the points in the divisor while the other terms give information about the corresponding y-values and multiplicities.

3.2. **Multiplication of ideals.** The next step in providing an explicit arithmetic on ideals is to explain how to multiply two ideals. This is once again the same as the number field case (see Cohen [6, Section 4.7.1]). Given two ideals with bases $\{b_1, \ldots, b_n\}$ and $\{b'_1, \ldots, b'_n\}$, we compute all products $b_i b'_j$ and write them in terms of the standard basis $\{1, y, \ldots, y^{n-1}\}$ for $k(C)/k(x)$. We then reduce the resulting $n \times n^2$ matrix over $k[x]$ into HNF using an analogue of Algorithm 2.4.5 of Cohen [6].

We now estimate the complexity (in terms of operations in k) of this ideal multiplication process. Suppose that all entries $b_{i,j}(x)$ in the bases $\{b_i\}$ and $\{b'_i\}$ are polynomials over $k[x]$ of degree bounded by B. The first step of computing all $b_i b'_j$ takes $O(n^2 \cdot n^2 \cdot (B^2 + 2B\delta))$ operations in k. The resulting matrix has size $n \times n^2$ and its entries are polynomials of degree bounded by $2B + \delta$.

The HNF algorithm on an $N_1 \times N_2$ matrix involves at most $N_1 N_2$ iterations of computing an extended greatest common divisor and taking a linear combination of two columns. One problem with HNF computations is that the intermediate steps can cause the size of the matrix entries to grow very large. This can be avoided by working modulo the determinant of the matrix (see Cohen [6] Section 2.4.2). In our application, since we begin with ideals represented in HNF, the determinant of the matrix can be easily calculated. Hence we will assume that all operations are performed modulo the determinant and we will write M for the degree of the determinant. An extended greatest common divisor of two polynomials of degree bounded by M can be performed in $O(M^2)$ operations. Therefore the total complexity of the HNF algorithm in this case is $O(N_1 N_2 (M^2 + N_1 M^2) = O(N_1^2 N_2 M^2)$. When reducing a product of ideals we will have $N_1 = n$ and $N_2 = n^2$.

3.3. **Ideal reduction strategy.** We now consider the process which computes the unique integral ideal of smallest degree in a given class. First we recall the strategy used in the number field case.

In the number field case an ideal $\mathfrak{a}$ is reduced by finding a small element $\alpha \in \mathfrak{a}$ and defining the reduction to be $\mathfrak{b} := \mathfrak{a}/(\alpha)$, which is a fractional ideal such that $1 \in \mathfrak{b}$. One drawback with this approach is that the reduction process is not necessarily unique. This problem is not so serious as one can determine that two "reduced" ideals $\mathfrak{b}_1$ and $\mathfrak{b}_2$ are equivalent by reducing $\mathfrak{b}_1 \mathfrak{b}_2^{-1}$ to a principal ideal. A more serious drawback of this reduction method is that it does not seem to be possible to give a useful bound on the degree of the resulting "reduced" ideal.

For these reasons we propose a different strategy for the reduction of ideals which is motivated by the following elementary result.

Lemma 5. *Let* $\mathfrak{a}$ *be an integral ideal of* $\mathcal{O}$. *Let* α *be an element of* $\mathfrak{a}$ *of minimal degree and define* $\mathfrak{b} := (\alpha)/\mathfrak{a}$. *Then* $\mathfrak{b}$ *is the reduced ideal in the class of* $\mathfrak{a}^{-1}$.

Proof. First note that $\mathfrak{b}$ is an integral ideal which is equivalent to $\mathfrak{a}^{-1}$. Minimality of the degree of $\mathfrak{b}$ is then clearly equivalent to the minimality of the degree of α. □

Ideal reduction algorithm.
Input: A $k[x]$-module basis for an integral $\mathcal{O}$-ideal $\mathfrak{a}$.

1. Compute an integral ideal $\mathfrak{a}'$ which is equivalent to the ideal $\mathfrak{a}^{-1}$.
2. Find an element $\alpha \in \mathfrak{a}'$ of minimal positive degree.
3. Output a basis for the ideal $\mathfrak{b} := (\alpha)/\mathfrak{a}'$ in Hermite Normal Form (HNF).

We now give further details on how to implement these three steps.

3.4. **Computing the inverse of an ideal.** One method to compute $\mathfrak{a}'$ is to take $\prod_{\sigma \neq 1} \mathfrak{a}^{\sigma}$ where σ runs over elements of the Galois group of $k(C)/k(x)$ (which is a cyclic group of order n). We choose a generator σ by fixing an nth root of unity ζ_n and imposing the action $\sigma(y) = \zeta_n y$. To calculate a basis for $\mathfrak{a}^{\sigma}$ from a basis for $\mathfrak{a}$, simply multiply each $a_{i,j}(x)$ by ζ_n^{j-1}. Note that the intermediate calculations involve extending the ground field from k to $k(\zeta_n)$, though the final result is defined over k. It is clear that $\mathfrak{a}'$ is an integral ideal in the ideal class $\mathfrak{a}^{-1}$.

This method involves $n-1$ iterations of the ideal multiplication process. It is also necessary to perform an HNF reduction at each step to prevent the matrices from becoming too large. Note that the norm of the ideal and the determinant of the matrix representation are easy to keep track of, but that they grow quite large: if the initial ideal has determinant of degree M, then by the final stage the determinant has degree $(n-1)M$.

An alternative method to compute $\mathfrak{a}'$ which does not require extending the ground field is to use the following strategy (Cohen [6, Section 4.8.4]). Precompute an $n \times n$ matrix $T = (\mathrm{Tr}_{k(C)/k(x)}(y^{i+j}))_{i,j=0}^{n-1}$. Also precompute the different which is known in this case to be the principal ideal (y^{n-1}) (see Neukirch [13, Satz III.2.4]).

Given an ideal $\mathfrak{a}$, represented as a matrix A in HNF, compute the matrix $(A^t T)^{-1}$ as a matrix in $k(x)$ (i.e., the entries will be ratios of polynomials in $k[x]$). The columns can be taken as a basis for a $k[x]$-module, say $\mathfrak{b}$. Multiplying $\mathfrak{b}$ by the different (y^{n-1}) gives an ideal $\mathfrak{a}'$ which is equivalent to $\mathfrak{a}^{-1}$ (see Cohen [6, Proposition 4.8.19]).

We now compare these two ideal inversion strategies. A major problem with algorithms such as these is the growth of the entries of the matrices. As we have seen in subsection 3.2, this can be controlled in the case of multiplication of ideals by working modulo the discriminant. The authors are not aware of any similar technique for matrix inversion. Hence our discussion will focus on the first strategy for ideal inversion. In practice, particularly in those cases where $k(\zeta_n)$ is a large degree extension of k, the second method may be useful.

Theorem 6. *Let $\mathfrak{a}$ be an integral $\mathcal{O}$-ideal of degree bounded by M represented in HNF. Let $d = [k(\zeta_n) : k]$. Then the first strategy above for computing $\mathfrak{a}'$ requires $O((n-1)^3 n^4 M^2 d^2)$ operations in k.*

Proof. The algorithm performs the ideal multiplication and HNF algorithms of subsection 3.2 a total of $n-1$ times. Note, however, that the degree of the determinant grows from M to $M' = (n-1)M$ in the course of the algorithm. Also note that the basic arithmetic operations (i.e., multiplication, addition, etc.) now take place in $k(\zeta_n)$.

For the ideal multiplication algorithm we have $B, \delta \leq M'$, and so the complexity is $O(n^4(M')^2d^2)$ operations in k. For the HNF reduction we have $N_1 = n, N_2 = n^2$, so the complexity is also $O(n^4(M')^2d^2)$ operations in k. The result follows. □

3.5. **Remaining steps.** To perform the second step of the algorithm, we proceed as follows. The ideal $\mathfrak{a}'$ is represented as a $k[x]$-module of rank n. For small values for n the smallest element can be found by classical algorithms (reduction of quadratic forms when $n = 2$ and Voronoi's algorithm when $n = 3$). For larger values of n we propose using lattice reduction techniques. The details and complexity statement are given in Section 4.

The third step looks as if it requires another ideal division. However, notice that $(\alpha)/\mathfrak{a}' = (\alpha)\mathfrak{a}/N_{k(C)/k(x)}(\mathfrak{a})$. We already have $N_{k(C)/k(x)}(\mathfrak{a})$ so this step just involves multiplying the ideal $\mathfrak{a}$ by the function $f(x, y) := \alpha/N_{k(C)/k(x)}(\mathfrak{a})$. Any denominators in the function $f(x, y)$ must occur in all entries of the HNF representation of $\mathfrak{a}$ so the division is straightforward. The final task is to ensure that the output ideal is represented in HNF. The complexity of this step is therefore dominated by the HNF reduction.

For some applications the goal of divisor reduction is to find an element of the Riemann-Roch space, i.e., to find a function f such that $D_1 = D_2 + (f)$ (where D_2 is the original divisor and D_1 is the reduced divisor). Observe that the function $f(x, y)$ defined in the previous paragraph is such a function.

3.6. **Complexity of addition in the divisor class group.** We now return to the context of the divisor class group of a superelliptic curve. We are given two divisors of the form $E_1 - m_1P_\infty$, $E_2 - m_2P_\infty$ represented as ideals $\mathfrak{a}_1$, $\mathfrak{a}_2$, and we know that $\deg(\mathfrak{a}_i) = m_i \leq g$. To perform the addition in the divisor class group, we first multiply the two ideals to obtain an integral ideal $\mathfrak{a}$. We then perform the ideal reduction process to obtain an ideal $\mathfrak{a}_3$ corresponding to the unique reduced divisor $E_3 - m_3P_\infty$. The ideal $\mathfrak{a}$ is a product of two ideals of degree bounded by g, so the determinant of the associated matrix has degree less than $M = 2g$.

We now consider the first step of the ideal reduction process. Let $d = [k(\zeta_n) : k]$. Putting the estimates above into Theorem 6 gives a total of $O(n^7g^2d^2)$ operations in k to construct the integral ideal $\mathfrak{a}'$.

We now consider the second step of the ideal reduction process. The lattice we reduce using the methods of Section 4 has entries of degree bounded by $2(n-1)g$, and so Proposition 12 shows that the complexity of the lattice reduction is $O(n^7g^2)$. Therefore the total complexity is dominated by the process of step one and we have proved the following result.

Theorem 7. *Let $d = [k(\zeta_n) : k]$. The complexity of the addition algorithm for reduced divisors in the divisor class group of a superelliptic curve C over a field k is $O(n^7d^2g^2)$ operations in k. This is polynomial time in terms of the input size.*

Algorithms for addition in the divisor class group of a curve have often been described for curves in families with a fixed value for n (e.g., hyperelliptic curves or cubic function fields). We observe that if n is fixed, then our complexity is $O(g^2)$ which is the same as the hyperelliptic case.

Due to considerations of space, our analysis has been very crude. It is likely that the algorithm would perform much better than this in practice. We note that the methods of this section are completely general. The reduction strategy we have

proposed could be applied to ideal class groups of any order in the function field of any algebraic curve.

4. Finding an element of minimal degree

In this section we give a method for finding an element of minimal positive degree in an integral $\mathcal{O}$-ideal. Once again our approach is motivated by the number field situation (Cohen [6, Section 6.5]). In that case the strategy is to consider the embeddings of the number field into $\mathbb{C}$ and then perform lattice reduction using the usual absolute value as a notion of size. The function field analogue of Minkowski's geometry of numbers was developed by Mahler [12]. In our case we use the fact that $k(C)$ may be embedded in a certain field of Puiseux series. We also use the fact that Puiseux series are equipped with a norm which arises from the natural extension of the valuation at infinity to the completion of $k(C)$.

Definition 3. Let K be a field and n a positive integer. Then

$$K\langle x^{1/n}\rangle := \left\{ \sum_{i=-\infty}^{m} a_i x^{i/n} \mid a_i \in K, a_m \neq 0 \right\}$$

is called the *field of Puiseux series*. The *norm* of an element of $K\langle x^{1/n}\rangle$ is defined to be $\left|\sum_{i=-\infty}^{m} a_i x^{i/n}\right| := m/n$.

We will now recall how to embed $k(C)$ into such fields.

Proposition 8. *Let C be a superelliptic curve over k with equation $y^n = c(x)$. Let $k_n := k(\zeta_n)$ be the field extension of k containing the nth roots of unity. Then there exist n distinct elements $\rho_1, \ldots, \rho_n \in \mathbb{K} := k_n\langle x^{1/n}\rangle$ such that $\rho_i^n = c(x)$. The n distinct choices $y \mapsto \rho_i$ induce n distinct homomorphisms Ψ_i from $k(C)$ to $\mathbb{K}$.*

For the proof see Walker [23] or [17, Theorem 9]. The elements ρ_i differ from each other only by powers of ζ_n.

The embeddings of elements of $k(C)$ into $\mathbb{K}$ give rise to an embedding of ideals in the following way. Let $\mathfrak{a}$ be an $\mathcal{O}$-ideal and consider the map $\Psi : \mathfrak{a} \to \mathbb{K}^n$ given by

$$\Psi : \alpha = \sum_{j=0}^{n-1} a_j(x) y^j \longmapsto (\Psi_1(\alpha), \Psi_2(\alpha), \ldots, \Psi_n(\alpha)) = \left(\sum_{j=0}^{n-1} a_j(x) \rho_i^j \right)_{i=1,\ldots,n}.$$

The following result is then immediate.

Proposition 9. *The image of $\mathfrak{a}$ under Ψ is a $k[x]$-module in $\mathbb{K}^n$ of rank n. In other words, $\Psi(\mathfrak{a})$ is a lattice in $\mathbb{K}^n$ over $k[x]$.*

We may define a norm on $\mathbb{K}^n$ by

$$\|(\alpha_1, \ldots, \alpha_n)\| := \max_{i=1,\ldots,n} \{|\alpha_i|\}.$$

Recall that our goal is to find an element $\alpha \in \mathfrak{a}$ of smallest degree. Therefore, the following easy result is important.

Proposition 10. *Let α be an element of an $\mathcal{O}$-ideal $\mathfrak{a}$. Then*

$$\deg(\alpha) = n\|\Psi(\alpha)\|.$$

The strategy for solving our problem is now clear: finding $\alpha \in \mathfrak{a}$ of minimal degree is equivalent to finding the shortest vector in the lattice $\Psi(\mathfrak{a}) \subset \mathbb{K}^n$ with respect to the norm on $\mathbb{K}^n$. The shortest vector can be efficiently found using a modified version of the lattice reduction algorithm due to A. Lenstra [11] (also see [15]). We note that in this case it is known that the lattice reduction algorithm always yields a minimum of the lattice (unlike in the number field case where it can only be proven that a rather small vector is found).

The above strategy can be adapted to the problem of reducing ideals in any function field. However, a drawback with this method is that computations with Puiseux series are required. In the case of superelliptic curves we can avoid computations with Puiseux series due to the following observation.

Theorem 11. *Let C/k be a superelliptic curve and let $\mathcal{O}$ be as above. Let $\alpha = \sum_{j=0}^{n-1} a_j(x)y^j \in \mathcal{O}$. Then*

$$\|\Psi(\alpha)\| = \max_j \{\deg_x(a_j(x)) + \delta j/n\}.$$

Proof. It is easy to see that all $|\rho_i| = \delta/n$. Therefore, each term $\|\Psi(a_j(x)y^j)\| = \deg_x(a_j(x)) + \delta j/n$. Now, since $(\delta, n) = 1$, each of these values is in a different class in $\frac{1}{n}\mathbb{Z}/\mathbb{Z}$. It follows that there can be no cancellation of terms, and the result follows. □

This is an important result as it shows that we do not actually need to compute with Puiseux series or to extend the ground field to $k(\zeta_n)$. Instead we can perform the lattice basis reduction process directly on the $k[x]$-module representation of the ideal $\mathfrak{a}$ and merely "pretend" that we are working in a lattice of Puiseux series.

We now give some details of the lattice reduction algorithm in this case. The lattice we are trying to reduce is the $\mathcal{O}$-ideal $\mathfrak{a}$, which is given as a $k[x]$-module via a basis $\{b_1, \ldots, b_n\}$. The lattice is endowed with a norm (which we will also denote by $\|\cdot\|$) via the Puiseux embedding above. Each b_i can be written as $\sum_{j=0}^{n-1} b_{i,j}(x)y^j$ and Theorem 11 shows that the norm of b_i is $\|b_i\| = \max_j\{\deg_x(b_{i,j}(x)) + \delta j/n\}$.

To compute the determinant of the lattice, we consider the modified vectors $b_i' = (b_{i,0}(x), b_{i,1}(x)x^{\delta/n}, \ldots, b_{i,n-1}(x)x^{(n-1)\delta/n})$.

Definition 4. The *determinant* of the lattice $\mathfrak{a}$ (denoted $d(\mathfrak{a})$) is the determinant of the $n \times n$ matrix B' over $k[x]$ which has the modified vectors b_i' as rows. The *orthogonality defect* of a basis $\{b_1, \ldots, b_n\}$ for $\mathfrak{a}$ (denoted $OD(b_1, \ldots, b_n)$) is defined to be

$$\sum_{i=1}^{n} \|b_i\| - \deg_x d(\mathfrak{a}) \in \frac{1}{n}\mathbb{Z}.$$

It is easy to see that $OD(b_1, \ldots, b_n) \geq 0$. The usual notion of a reduced basis [11] also applies with these modified notions and so we make the following definition.

Definition 5. The basis $\{b_1, \ldots, b_n\}$ for $\mathfrak{a}$ is *reduced* if $OD(b_1, \ldots, b_n) = 0$.

The lattice reduction algorithm of [11] and its complexity analysis can therefore be easily adapted to the current setting. We do not give the details here.

Proposition 12. *Let $b_1, \ldots, b_n$ be vectors in $k[x]^n$ corresponding to a basis for an ideal $\mathfrak{a}$ as above. Let B be a bound on the degrees of all entries $b_{i,j}(x)$ and let S be a*

bound on the number of values for $\sum_{i=1}^{n} \|b_i\|$. *Then the lattice reduction algorithm computes a reduced basis in time*

$$O(n^3 \cdot B \cdot S).$$

In the case of reduction of divisors on a superelliptic curve we apply the lattice reduction algorithm to the ideal $\mathfrak{a}'$ which has degree less than $(n-1)g$. To deduce the complexity of lattice reduction in this case we use the very crude bounds $B = O(ng)$ and $S = n^2B = O(n^3g)$ (since $\|b_i\|$ must be multiplied by n to become integer valued) and obtain a complexity of $O(n^7g^2)$.

5. The discrete logarithm problem

We now restrict to the case where k is a finite field $\mathbb{F}_q$. In this case $\mathrm{Pic}^0_{\mathbb{F}_q}(C)$ is a finite group. Suppose that D_1 is a divisor class in $\mathrm{Pic}^0_{\mathbb{F}_q}(C)$ and that D_2 lies in the subgroup generated by D_1. We will assume that the order of the divisor D_1 in $\mathrm{Pic}^0_{\mathbb{F}_q}(C)$ is known and, for simplicity, that it is a prime L. Then the discrete logarithm problem is to find an integer λ such that $D_2 = \lambda D_1$. The discrete logarithm problem arises in the context of cryptography using algebraic curves.

In this section we describe how the algorithm due to Adleman, De Marrais and Huang [1], which was developed for hyperelliptic curves, can be modified to apply to superelliptic curves. The algorithm is an index calculus method and so the main process involves generating relations amongst elements of a "factor base". The central idea behind [1] is to find these relations by considering the decomposition of principal divisors coming from functions of the form

$$\phi = a(x) + b(x)y.$$

The algorithm of this section uses exactly the same strategy. Before describing the details, it is necessary to study the support of the principal divisor (ϕ).

Proposition 13. *Let* $C/\mathbb{F}_q$ *be a superelliptic curve* $y^n = c(x)$ *and let* $\mathcal{D}(x)$ *be the discriminant of the extension* $\mathbb{F}_q(C)/\mathbb{F}_q(x)$. *Let* ϕ *be the function* $a(x)+b(x)y$ *where* $a(x)$ *and* $b(x)$ *are coprime elements of* $\mathbb{F}_q[x]$. *The principal divisor* (ϕ) *is of the form* $E - (\deg E)P_\infty$ *where* E *is an effective divisor. A prime* $p(x)$ *of* $\mathbb{F}_q[x]$ *lies in the support of the divisor* E *if and only if* $p(x)$ *divides*

$$N_\phi = N_{\mathbb{F}_q(C)/\mathbb{F}_q(x)}(\phi) = a(x)^n + (-1)^n c(x) b(x)^n.$$

Suppose $p(x)$ *is a prime of* $\mathbb{F}_q[x]$ *such that* $p(x) | N_\phi$ *and* $\gcd(p(x), \mathcal{D}(x)) = 1$. *Then there is at most one prime divisor* $\mathfrak{p}$ *lying above* $p(x)$ *in the support of* (ϕ) *and this prime divisor has* $e_\mathfrak{p} = f_\mathfrak{p} = 1$ *(where* $e_\mathfrak{p}$ *and* $f_\mathfrak{p}$ *are the ramification and residue class degrees, respectively).*

Proof. First observe that the poles of the function ϕ occur at the point at infinity on C. It remains to deal with the zeroes of the function ϕ.

That the primes in the support divide N_ϕ is clear. Those with $\gcd(p(x), \mathcal{D}(x)) = 1$ necessarily have $e_\mathfrak{p} = 1$. It remains to prove that $f_\mathfrak{p} = 1$ and that at most one prime divisor above any such $p(x)$ appears.

Let $\mathfrak{p}$ be any prime divisor over $p(x)$ such that ϕ is zero on $\mathfrak{p}$. Over an algebraic closure of $\mathbb{F}_q$ we can think of the divisor $\mathfrak{p}$ as a sum of points (x_0, y_0). Since $\phi = a(x) + b(x)y$ is zero at each (x_0, y_0), it follows that $y_0 = -a(x_0)/b(x_0)$ and

thus it is not possible to have two different points (x_0, y_0) for the same x_0. Furthermore, it follows that $y_0 \in \mathbb{F}_q(x_0)$, which means that $f_{\mathfrak{p}} = (\mathcal{O}/\mathfrak{p} : \mathbb{F}_q[x]/(p(x))) = (\mathbb{F}_q(x_0, y_0) : \mathbb{F}_q(x_0)) = 1$.

Finally, suppose that the point (x_0, y_0) has multiplicity greater than one above the point x_0 of the x-line. In this case x_0 would be a root of $\mathcal{D}(x)$ and, since all Galois conjugates of x_0 would have the same property, $p(x)|\mathcal{D}(x)$. □

This result shows that decomposing (ϕ) is easy: simply factor N_ϕ into irreducibles, $\prod p_i(x)^{t_i}$. For those $p_i(x)$ which do not divide $\mathcal{D}(x)$ we set $r(x, y) = y - \big(a(x)b(x)^{-1} (\mathrm{mod}\ p_i(x))\big)$ and observe that the prime divisor $(p_i(x), r(x, y))$ lies in the support of ϕ with multiplicity t_i. Divisors for which $p(x)$ divides $\mathcal{D}(x)$ are seen to have $f_{\mathfrak{p}} = 1$ and can only appear with multiplicity one.

Proposition 13 implies that only relations involving ramified prime divisors and those with residue class degree one can be found using this method. Therefore, the factor base is taken to consist only of prime divisors with these properties. It is important that the factor base generate the full divisor class group so we need the following modified version of Theorem 2 of Müller, Stein and Thiel [14].

Theorem 14. *Let C be any curve over $\mathbb{F}_q$. Define next_prime(x) to be the smallest prime $p \geq x$. The divisor class group of C is generated by the set of prime divisors of residue class degree one whose degree is less than d, where*

$$d := \mathit{next\_prime}\left(\max\left\{n + 1, 2\log_q(4g - 2)\right\}\right).$$

Proof. The proof of this result is a slight adaption of the proof in [14] (though note that we use the notation $f_{\mathfrak{p}}$ for the residue class degree rather than the total degree). We assume, for a contradiction, that the statement is false and we let χ be a nontrivial character on $\mathrm{Pic}^0_{\mathbb{F}_q}(C)$ which is trivial on the group generated by the prime divisors in question.

Let $\prod^\dagger$ denote the product over all prime divisors other than those of residue degree one and degree less than or equal to q^d. We obtain, as in [14],

$$\prod_{i=1}^{2g-2}(1 - \omega_i(\chi)u) = \frac{\prod_{i=1}^{2g}(1 - \omega_i u)}{(1-u)(1-qu)} \prod{}^\dagger \frac{1 - u^{\deg \mathfrak{p}}}{1 - \chi(\mathfrak{p})u^{\deg \mathfrak{p}}}.$$

Taking the logarithmic derivative of this equation and equating coefficients of u^{d-1} gives

$$-\sum_{i=1}^{2g-2}\omega_i(\chi)^d = -\sum_{i=1}^{2g}\omega_i^d + (1 + q^d) + A \text{ where } A = \sum{}^+(\deg \mathfrak{p})(\chi(\mathfrak{p})^{d/\deg \mathfrak{p}} - 1)$$

and the sum $\sum^+$ is over all prime divisors of degree dividing d and of residue class degree greater than one. The argument in [14] gives the required contradiction as long as $A = 0$.

Since d is prime, the only divisors in the last sum must be of degree d or 1. But they cannot be of degree one since $\deg \mathfrak{p} \geq f_{\mathfrak{p}} > 1$. So the sum is over all divisors of degree d and of residue class degree greater than one. Now, again since d is prime, the residue class degree is equal to d, but this is impossible since the residue class degree must be $\leq n$. Hence the sum is empty and $A = 0$. □

We now give some further details of the algorithm, though our presentation is sketchy since the basic idea is very similar to [1] whilst a good implementation

would include numerous techniques which are well known from research on factoring algorithms such as the number field sieve. The input consists of a superelliptic curve $C/\mathbb{F}_q$ of genus g and two divisors D_1 and D_2 of prime order L.

The first step is to construct the factor base $\mathcal{F}$. We define a number $S_1 = c_1(\delta \log_q(\delta))^{1/2}$ where c_1 is a constant whose value is chosen to optimise the running time of the algorithm. The set $\mathcal{F}$ consists of

- all prime divisors of C above the primes $p(x)$ of $\mathbb{F}_q[x]$ which divide the discriminant $\mathcal{D}(x)$ (equivalently, divide $c(x)$).
- all unramified prime divisors of residue class degree equal to one which lie above a prime $p(x)$ of $\mathbb{F}_q[x]$ of degree less than or equal to S_1.

Note that there are some "automatic" relations between some of these factor base elements which arise from the splitting and ramification behaviour. These relations should be added to the relation matrix.

The second step is to "smooth" the initial divisors D_1 and D_2 so that they are expressed in terms of elements of $\mathcal{F}$. This is more complicated than in [1] since the supports of the D_i may contain primes which do not have residue class degree one and so we cannot possibly smooth them using functions of the form ϕ. To deal with this we take random combinations $D = m_1 D_1 + m_2 D_2$ in the divisor class group until the support of D consists entirely of residue degree one places, clearly this step requires our earlier algorithm for adding divisors. We could repeat until the divisor D actually decomposes entirely over $\mathcal{F}$ (as in the Hafner-McCurley style algorithm). However, the strategy we propose is to simply add the prime divisors in the support of D to $\mathcal{F}$ and find relations by sieving (by using an analogue of the lattice sieve method of the NFS factoring algorithm).

Once a few relations have been found which link D_1 and D_2 to the factor base, it remains to find a full set of relations amongst elements of $\mathcal{F}$. This is done by attempting to decompose random functions ϕ, as in [1]. This step can be speeded up by using a sieving operation like the one described in [7]. Once enough relations have been found then sparse linear algebra modulo L is performed to obtain the solution to the discrete logarithm problem.

It is clear that, if it terminates, the algorithm will give a solution to the discrete logarithm problem. Justifying why the algorithm terminates is less easy since it might not be possible to generate a full set of relations by considering functions of the form $\phi = a(x) + b(x)y$. For this reason a practical implementation would take several different functions θ such that $k(x, \theta) = k(C)$ and sieve over functions of the form $a(x) + b(x)\theta$ (one then needs a slightly more general form of Proposition 13). Indeed, this provides a natural way to parallelise the algorithm.

The analysis of the complexity of the algorithm relies on heuristic statements regarding the smoothness of the polynomials N_ϕ. Rather than give the details we merely state the complexity as a conjecture. A more precise statement is the subject of further research by the authors.

Conjecture 1. *Let n and q be fixed. Then there exists a constant $c > 0$ such that, for all sufficiently large δ, the algorithm proposed above will solve the discrete logarithm problem on any superelliptic curve over $\mathbb{F}_q$ (of degrees n and δ as in Definition 1) in expected time*

$$O\big(\exp\big(c\sqrt{\log(q^g)\log\log(q^g)}\big)\big).$$

In terms of the size of the original discrete logarithm problem (which is in a subgroup of size L) we observe that this algorithm is most useful when L is of size close to q^g.

We note that the results in this section can be applied to the case of the divisor class group of any plane curve with a single point at infinity. Hence it is reasonable to conjecture that a suitable modification of the above method would give a subexponential time algorithm for solving the discrete logarithm problem on any nonsingular plane curve over $\mathbb{F}_q$ with a single point at infinity.

References

[1] L. Adleman, J. De Marrais, and M.-D. Huang. A subexponential algorithm for discrete logarithms over the rational subgroup of the Jacobians of large genus hyperelliptic curves over finite fields. In *ANTS-1 : Algorithmic Number Theory*, L. Adleman and M.-D. Huang, editors. Springer-Verlag, LNCS 877, 1994. MR **96b:**11078

[2] S. Arita. Algorithms for computations in Jacobian group of $C_{a,b}$ curve and their application to discrete-log-based public key cryptosystems. in *The mathematics of public key cryptography, Fields Institute* A. Odlyzko et al (eds.), 1999.

[3] E. R. Barreiro, J.-P. Cherdieu and J. E. Sarlabous. Efficient reduction on the Jacobian variety of Picard curves. in *Coding theory, cryptography and related areas*, J. Buchmann et al (eds.), Springer, 2000.

[4] D.G. Cantor. Computing in the Jacobian of a hyperelliptic curve. *Math. Comp.*, **48**, 95–101, 1987. MR **88f:**11118

[5] J. Coates. Construction of rational functions on a curve. *Proc. Cam. Phil. Soc.*, **68**, 105–123, 1970. MR **41:**3477

[6] H. Cohen. *A Course in Computational Algebraic Number Theory.* Springer-Verlag, GTM 138, 1993. MR **94i:**11105

[7] R. Flassenberg and S. Paulus. Sieving in function fields. *Experimental Mathematics*, **8**, No. 4, 339–349, 1999. CMP 2000:07

[8] W. Fulton. *Algebraic Curves.* Benjamin, 1969.

[9] F. Hess. Zur divisorenklassengruppenberechnung in globalen funktionenkörpern. Thesis, T-U Berlin 1999.

[10] M.-D. Huang and D. Ierardi. Efficient algorithms for the Riemann-Roch problem and for addition in the jacobian of a curve. *J. Symbolic Comp.*, **18**, 519–539, 1994. MR **96h:**14077

[11] A. Lenstra. Factoring multivariate polynomials over finite fields. *J. Computer and Systems Science*, **30**, 235-248, 1985. MR **87a:**11124

[12] K. Mahler. An analogue to Minkowski's geometry of numbers in a field of series. *Ann. Math.*, **42**, 488–522, 1941. MR **2:**350c

[13] J. Neukirch. *Algebraische Zahlentheorie.* Springer, 1990. MR **92a:**01057

[14] V. Müller, A. Stein and C. Thiel. Computing discrete logarithms in real quadratic function fields of large genus. *Math. Comp.*, **68**, 807–822, 1999. MR **99i:**11119

[15] S. Paulus. Lattice basis reduction in function fields. In *ANTS-3 : Algorithmic Number Theory*, J. Buhler, editor. Springer-Verlag, LNCS 1423, 567–575, 1998. CMP 2000:05

[16] S. Paulus and H.-G. Rück. Real and imaginary quadratic representations of hyperelliptic function fields. *Math. Comp.*, **68**, 1233–1241, 1999. MR **99i:**11107

[17] M. Pohst and M. Schörnig. *On integral basis reduction in global function fields.* In *ANTS-2 : Algorithmic Number Theory*, H. Cohen, editor. Springer-Verlag, LNCS 1122, 273–283, 1996. MR **98c:**11125

[18] R. Scheidler, A. Stein and H. C. Williams, Key-exchange in real quadratic congruence function fields. *Designs, Codes and Cryptography.*, **7**, 153–174, 1996. MR **97d:**94009

[19] R. Scheidler, A. Stein. Voronoi's algorithm in purely cubic congruence function fields of unit rank 1. *Math. Comp.* **69** (2000) 1245–1266. CMP 2000:11

[20] R. Scheidler. Ideal arithmetic and infrastructure in purely cubic function fields. Preprint 1999.

[21] H. Stichtenoth. *Algebraic function fields and codes.* Springer Universitext, Springer, 1993. MR **94k:**14016

[22] E. J. Volcheck. Computing in the Jacobian of a plane algebraic curve. In *ANTS-1 : Algorithmic Number Theory*, L. Adleman and M.-D. Huang, editors. Springer-Verlag, LNCS 877, 221–233, 1994. MR **96a:**14033

[23] R. J. Walker. *Algebraic Curves.* New York: Springer 1978. MR **80c:**14001

Institute for Experimental Mathematics, Ellernstr. 29, 45326 Essen, Germany
E-mail address: galbra@exp-math.uni-essen.de

Kopernikusstrasse 15, 69469 Weinheim, Germany
E-mail address: sachar.paulus@t-online.de

Department of Computer Science, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, United Kingdom
E-mail address: nigel@cs.bris.ac.uk