Index Calculus in Class Groups of Plane Curves of Small Degree

Claus Diem

April 18, 2005

Abstract

We present a novel index calculus algorithm for the discrete logarithm problem (DLP) in degree 0 class groups of curves over finite fields. A heuristic analysis of our algorithm indicates that asymptotically for varying q, "essentially all" instances of the DLP in degree 0 class groups of curves represented by plane models of a fixed degree d over \mathbb{F}_q can be solved in an expected time of $\tilde{O}(q^{2-2/(d-2)})$.

A particular application is that heuristically, "essentially all" instances of the DLP in degree 0 class groups of non-hyperelliptic curves of genus 3 (represented by plane curves of degree 4) can be solved in an expected time of $\tilde{O}(q)$.

We also provide a method to represent "sufficiently general" (nonhyperelliptic) curves of genus $g \ge 3$ by plane models of degree g+1. We conclude that on heuristic grounds the DLP in degree 0 class groups of "sufficiently general" curves of genus $g \ge 3$ (represented initially by plane models of bounded degree) can be solved in an expected time of $\tilde{O}(q^{2-2/(g-1)})$.

1 Introduction

This work is concerned with the discrete logarithm problem (DLP) in degree 0 class groups of arbitrary (non-singular, projective, geometrically irreducible) curves over finite fields. The motivation for the work is derived from cryptanalytic applications: First, the DLP in degree 0 class groups of non-hyperelliptic curves of genus 3 has been suggested as a cryptographic primitive for public key cryptosystems (see e.g. [4], [3] as well as [11], [12]). Second, the method of "covering attacks" (aka Weil descent attacks) (cf. [7, Appendix], [8], [21], [13, Section 4.4]) sometimes allows to transfer the DLP in elliptic curves (or in degree 0 class groups of hyperelliptic curves) over extension fields into the DLP in degree 0 class groups of curves over smaller fields. Often, the resulting curves are not hyperelliptic anymore. It is well-known among many cryptographers that one can *in principle* adapt the known index calculus attacks from hyperelliptic curves to more general curves but so far the efficiency of these attacks has been questioned.

We present a novel index calculus attack on the DLP in degree 0 class groups of curves. Our algorithm is particularly efficient if the (non-singular) curve is represented by a (possibly singular) plane model of *small degree*. A heuristic analysis of our algorithm gives rise to (see Section 4):

Heuristic Result 1 Let us consider the DLP in degree 0 class groups of curves represented by plane models of a fixed degree $d \ge 4$ over finite fields \mathbb{F}_q . Then "essentially all" instances of the DLP in such groups can be solved in an expected time of $\tilde{O}(q^{2-\frac{2}{d-2}})$.

Here, the \tilde{O} -notation means that we suppress logarithmic factors.

Additionally to the index calculus algorithm, we present a method to find plane models of degree g+1 of "sufficiently general" (non-hyperelliptic) curves of genus $g \ge 3$ (see Section 5).

By applying our algorithm to such a plane model, we obtain that on heuristic grounds the DLP in degree 0 class groups of "sufficiently general" non-hyperelliptic curves of a fixed genus $g \ge 3$ (initially represented by plane models of bounded degree) can be solved in an expected time of

$$\tilde{O}(q^{2-\frac{2}{g-1}}).$$

This result should be compared with the following heuristic result which can be obtained with an adaption of the double large prime variation algorithm by Gaudry, Thériault and Thomé ([15]) as well as Nagao ([28]) from hyperelliptic to arbitrary curves (see Section 2).

Let us consider the DLP in degree 0 class groups of curves of a fixed genus g over finite fields \mathbb{F}_q represented by plane models of bounded degree. Then "essentially all" instances of the DLP in such groups can be solved in an expected time of $\tilde{O}(q^{2-\frac{2}{g}})$.

An important special case for our algorithm is constituted by the DLP in degree 0 class groups of *non-hyperelliptic* curves of genus 3 over finite fields \mathbb{F}_q : Every such curve can (via the canonical embedding) be represented as a plane curve of degree 4. By applying our algorithm to such a model, we obtain a heuristic running time of $\tilde{O}(q)$ whereas the adaption of the algorithm in [15] and [28] leads to a heuristic running time of $\tilde{O}(q^{4/3})$. Our heuristic result for non-hyperelliptic genus 3 curves has been confirmed by an experimental study.

Curves, divisors, etc.

In this work, a *curve* is always non-singular (i.e. smooth), projective and geometrically irreducible.

In the presentation above we implicitly used the following conventions concerning the representation of curves, divisors and divisor classes:

Let q be a prime power. We fix a homogeneous coordinate system X, Y, Zof $\mathbb{P}^2/\mathbb{F}_q$. We think of every curve in question as being the normalization of a possibly singular curve in \mathbb{P}^2 . We distinguish the two by calling the latter one a *plane model* of the curve. We use a defining homogeneous polynomial to represent the plane model (and thus the curve itself).

We represent every point of $\mathcal{C}_{pm}(\mathbb{F}_{q^r})$ (with $r \geq 1$) by its coordinates in $\mathbb{P}^2(\mathbb{F}_{q^r})$. Points of $\mathcal{C}(\mathbb{F}_{q^r})$ not lying over singular points are represented by the corresponding points of $\mathcal{C}_{pm}(\mathbb{F}_{q^r})$. For every singular point of $\mathcal{C}_{pm}(\overline{\mathbb{F}}_q)$, we fix an enumeration on the corresponding points of $\mathcal{C}(\overline{\mathbb{F}}_q)$. We represent each closed point of \mathcal{C} with residue degree r by one of points the corresponding Galois orbit in $\mathcal{C}(\mathbb{F}_{q^r})$.

If not stated otherwise, we think of divisors on C as being represented as a formal sum of closed points in C. Following [20], we call this representation the *free representation*.

Let \mathcal{C}/\mathbb{F}_q be a curve of genus g with a fixed plane model \mathcal{C}_{pm} , and let $P_0 \in \mathcal{C}(\mathbb{F}_q)$ be a fixed point. For some divisor D of \mathcal{C}/\mathbb{F}_q , let us denote the corresponding divisor class by [D]. Then by the Riemann-Roch theorem, every element of $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ is of the form $[D] - g[P_0]$ for some effective divisor D of degree g. For our algorithmic purposes, we think of every element of $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ as being represented by such a divisor D. (This applies in particular to the input values to algorithms.) We note that every element of $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ is uniquely represented by $D - sP_0$ for s being minimal and some divisor D of degree s, but we do not need this unique representation.

Notation

Throughout the work, we will use the following notation:

Additionally to the homogeneous coordinate system X, Y, Z on $\mathbb{P}^2/\mathbb{F}_q$, we fix a coordinate system X_a, Y_a on $\mathbb{A}^2/\mathbb{F}_q$. We think of \mathbb{A}^2 as being included included in \mathbb{P}^2 via $(x_a, y_a) \mapsto (x : y : z) := (x_a : y_a : 1)$. Note that via this inclusion we have $X_a = \frac{X}{Z}, Y_a = \frac{Y}{Z}$.

We fix a defining homogeneous polynomial F(X, Y, Z) of \mathcal{C}_{pm} and let $f(X_a, Y_a) := F(X_a, Y_a, 1)$. We denote the "affine plane model" defined by f (i.e. the intersection of \mathcal{C}_{pm} with \mathbb{A}^2) by \mathcal{C}_a . (We only consider plane models of curves of genus ≥ 1 such that \mathcal{C}_{pm} cannot be contained in the line Z = 0.) We use the same notation for a function of \mathbb{A}^2 , its restriction to \mathcal{C}_a and the induced element in $\mathbb{F}_q(\mathcal{C})$.

Overview over this work

Our algorithm can be viewed as a variant of the recent double large prime variation algorithm by Gaudry, Thériault and Thomé ([15]) as well as Na-

gao ([28]).

To facilitate the description, we start off with the a generalization of the original index calculus algorithm by Gaudry with an improvement by Harley (see [14], in particular the "Conclusion", as well as [33]).

We then present a preliminary algorithm which can be viewed as a variant of the algorithm by Gaudry and Harley but which can also be viewed as a (simplified) adaption of the index calculus algorithm by Adleman, De-Marrais and Huang ([1]) to the small degree (and small genus) situation. We provide a heuristic analysis of the algorithm which is based on similar assumptions as the analysis of the algorithm by Adleman, DeMarrais and Huang in [1].

In Section 4 we include a double large prime variation into our algorithm and give the corresponding analysis. We note that the final result is heuristic but arguably the heuristic assumptions are milder than the assumptions for the heuristic analysis of the preliminary algorithm.

In Section 5 we provide a method to find plane models of degree g + 1 of "sufficiently general" (non-hyperelliptic) curves of genus g.

We finish with a discussion on the "full cost" of our algorithm applied to genus 3 curves and state some interesting lines for future research.

We have implemented our algorithms in the computer algebra system Magma. At the end of Sections 3 to 5 (in which we present our new algorithms), we report on experiments conducted with the respective algorithms. Moreover, at the end of Sections 2 to 4, we give some information on practical aspects of the algorithms and implicit constants.

Calculating the group order

In our algorithms, we always assume that the order of the cyclic subgroup in question (or the order of the full class group) is known. From a theoretical point of view this is however not an obstacle because it can be shown that the *L*-polynomials of curves C/\mathbb{F}_q represented by plane models of bounded degree can be calculated in (deterministic) polynomial time in $\log(q)$. (This result follows from [30, Theorem H] which in turn relies on Pila's extension of the point counting algorithm by Schoof ([32]) to abelian varieties ([29]).) Moreover, in cryptographic situations, the order of the cyclic subgroup in question is always known.

The heuristic nature of our results

The analyses of all index calculus algorithms presented in this work are heuristic. It is conceivable that there is a sequence of instances which violates the stated running times. This is why we talk about "essentially all instances".

A rigorous interpretation of our claims can be given as follows:

Let us fix the degree d. Now for a prime power q, let S(q) be the set of all instances of the DLP in curves over \mathbb{F}_q represented by plane models of degree d. (With the representations described above.)

The (conjectural) claim is now that there exist subsets $S_1(q)$ of S(q) with $\#S_1(q)/\#S(q) \longrightarrow 1 \ (q \longrightarrow \infty)$ such that the instances in $S_1(q)$ can be solved in the stated time.

In our variant of the Gaudry-Harley algorithm, we additionally fix the genus. Here a rigorous interpretation of our heuristic result can be given by considering all instances with a fixed genus.

We note that for a combination of Heuristic Result 1 and the results of Section 5, one needs that for a fixed degree d the statement in Heuristic Result 1 holds for essentially all instances where the curves have genus d-1 (resp. d-2).

Above, we also used the term "sufficiently general". This term will be defined in Section 5.

2 The algorithm by Gaudry and Harley

We assume that the reader is familiar with index calculus algorithms. A good overview to these algorithms in a general setting is given in [10]. (The algorithms presented in Sections 3 and 4 are however not specifications of the general description in [10].) As we are interested in index calculus in class groups of curves of small genus, all running time estimates in this section are given with respect to a *fixed genus g*. In order to bound the running time for the arithmetic in the class group in terms of field operations, we further assume that the curves are represented by plane models of *bounded degree*.

Index calculus on class groups of hyperelliptic curves (in imaginary quadratic representation) of small genus over finite fields was pioneered by Gaudry ([14]). Using all rational points as factor base, he obtained that heuristically, the DLP in class groups of hyperelliptic curves of genus g can be solved in an expected time of $\tilde{O}(q^2)$. The running time is thereby asymptotically dominated by the linear algebra part. In the Conclusions of [14] an idea of Harley's is mentioned: One reduces the factor base and balances the running times of the relation search and the linear algebra part. With this approach one obtains heuristically an expected running time of $\tilde{O}(q^{2-2/(g+1)})$ (see also [33]).

In [33], Thériault introduced and analyzed the algorithm with a large prime variation, and in [15] as well as in [28], Gaudry, Thériault, Thomé as well as Nagao introduced and analyzed the algorithm with a double large prime variation. With the double large prime variation one obtains heuristically an expected running time of $\tilde{O}(q^{2-2/g})$.

Our approach leads to analogs of all these four algorithms. Before we

come to our analogs, we present below an adaption of the reduced factor base index calculus algorithm by Gaudry and Harley to more general than hyperelliptic curves. In the next section we discuss what changes one has to make in order to obtain the corresponding new algorithm, and we present the algorithm. It is then not difficult to combine this new algorithm with a double large prime variation, and we study the resulting algorithm in Section 4.

Arithmetic in class groups of curves

Let \mathcal{C}/\mathbb{F}_q be a curve of genus $g \geq 1$, represented by a plane model \mathcal{C}_{pm} , and let $P_0 \in \mathcal{C}(\mathbb{F}_q)$ be a fixed point (used for the representation of the elements of $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$).

We need an efficient method to add two elements in $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$, that is, we need an efficient (randomized) algorithm which given two divisors D_1, D_2 of degree g on \mathcal{C}/\mathbb{F}_q calculates a representative $D_3 - gP_0$ of $[D_1] - [gP_0] + [D_2] - [gP_0]$. For this, one can use any algorithm to compute Riemann-Roch spaces. Indeed, if h is any non-trivial element of the Riemann-Roch space $L(D_1 + D_2 - gP_0)$ (which is always non-trivial), then $D_3 := \operatorname{div}(h) + D_1 + D_2 - gP_0$ is an effective divisor with $[D_1] - [gP_0] + [D_2] - [gP_0] = [D_3] - [gP_0]$.

The following proposition seems to be very classical. It follows for example from the work by Volcheck ([34]) as well as Huang and Ierardi ([22]) who use a method by Brill and Noether.

Proposition 2 Let us consider all curves C/\mathbb{F}_q represented by plane models C_{pm} of bounded degree. (Thus in particular the genus of C is bounded.)

Let \mathcal{C}/\mathbb{F}_q be one of these curves. Then after some precomputation for \mathcal{C} which takes a randomized polynomial time, one can calculate bases of Riemann-Roch spaces of divisors of bounded height on \mathcal{C}/\mathbb{F}_q in a bounded number of field operations (independent of \mathcal{C}).

In particular, if one considers curves \mathcal{C}/\mathbb{F}_q , represented by plane models of bounded degree, with a fixed point $P_0 \in \mathcal{C}(\mathbb{F}_q)$ (used for the representation of the elements of $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$), an addition of two elements $a, b \in \mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ can be performed in randomized polynomial time. One can thereby guarantee that the effective divisor used to represent a + b is uniformly distributed among all possible divisors.

Here, by the *height* of a divisor we mean the maximum of the degrees of the zero- and pole-divisor.

In the following, whenever in the algorithms we say that a divisor representing an element $a \in \operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ should be calculated, we mean that the divisor is selected uniformly at random among all possible effective divisors D with $[D] - [gP_0] = a$. **Remark 3** For the calculations in $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ (as always in free representation), it seems to be necessary to factorize polynomials. For this, no deterministic polynomial-time algorithm is available.

Remark 4 In relation to [34], the running time (after some precomputation and for suitable models) of the computation of Riemann-Roch spaces (and thus of the addition in the class group) was later considerably improved by Heß using an ideal- and function field theoretic approach ([20]) and by Khuri-Makdisi using a different "geometric" approach ([23]) (see "Practical aspects" at the end of this section).

Given Proposition 2, it is possible to generalize the known index calculus algorithms from hyperelliptic curves to the more general situation.

The reduced factor base index calculus algorithm for curves of small genus by Gaudry and Harley adapted to arbitrary curves

Input. A curve \mathcal{C}/\mathbb{F}_q of genus ≥ 1 , represented by a plane model, with a fixed point $P_0 \in \mathcal{C}(\mathbb{F}_q)$ (used for the representation of the elements of $\mathrm{Cl}^0(C/\mathbb{F}_q)$), $a, b \in \mathrm{Cl}^0(C/\mathbb{F}_q)$ such that $b \in \langle a \rangle$, the number $N := \#\langle a \rangle$, and a positive rational number r < 1. *Output.* An x such that $x \cdot a = b$.

- 1. Fix a subset (the "factor base") $\mathcal{F} \subseteq \mathcal{C}(\mathbb{F}_q)$ and an enumeration $\mathcal{F} = \{F_1, F_2, \ldots\}$ with $\#\mathcal{F} = \lceil q^r \rceil$, selected uniformly at random from all possible subsets. (If no such set exists, output "failure" and terminate.)
- 2. Construct a sparse matrix R over $\mathbb{Z}/N\mathbb{Z}$ as follows: Repeat { Choose randomly and uniformly $\alpha, \beta \in \mathbb{Z}/N\mathbb{Z}$. Calculate a divisor D such that $[D] - [gP_0] = \alpha a + \beta b$. If D splits completely, i.e. $D = \sum_{k=1}^{g} P_k$ for some $P_k \in \mathcal{C}(\mathbb{F}_q)$, { determine whether all P_k lie in the factor base, and if this is the case, { determine the relation $D = \sum_j r_j F_j$. (Now we have $(\sum_j r_j [F_j]) - g[P_0] = \alpha a + \beta b$.) Store the sparse vector $(r_j)_j$ as a new row of the sparse matrix $R = (r_{i,j})_{i,j}$ over $\mathbb{Z}/N\mathbb{Z}$. }

 $\}$ until the matrix R has more rows than columns.

3. Try to find a non-trivial element in γ ∈ ker(R^t) with a randomized algorithm for sparse linear algebra (with variants of Lanczos' or Wiedeman's algorithms). Thereby do computations modulo the prime divisors of N and use the Chinese Remainder Theorem to find γ. (If N is not square-free, use the lifting procedure described in [10, Section 4].) If the procedure fails, return to Step 2 and include some more relations into R.

(We now have

$$\left(\sum_{i} \gamma_i \alpha_i\right) a + \left(\sum_{i} \gamma_i \beta_i\right) b = \sum_{i,j} \gamma_i r_{i,j} (F_j - P_0) = 0.$$

4. If $\sum_i \gamma_i \beta_i \in (\mathbb{Z}/N\mathbb{Z})^*$, output

$$x = -\frac{\sum_{i} \gamma_{i} \alpha_{i}}{\sum_{i} \gamma_{i} \beta_{i}} \in \mathbb{Z}/N\mathbb{Z}$$

otherwise return to Step 2.

Remark 5 The "algorithm" is in fact an *algorithm scheme* because we have not specified how to do some essential computations (i.e. how to compute a divisor D with $[D] - g[P_0] = \alpha a + \beta b$ or how to perform the linear algebra. Moreover, the "algorithm" is randomized and not deterministic. Nonetheless, we refer to it as an "algorithm". The same remark applies to the "algorithms" in the following sections.

Complexity

As stated, we give the complexity estimates for fixed genus and bounded degree.

The factor base can be determined in the following way: First one enumerates $\mathcal{C}(\mathbb{F}_q)$. For this, one calculates for all possible *x*-coordinates the corresponding points in $\mathcal{C}(\mathbb{F}_q)$. This can be done in an expected time of $\tilde{\mathcal{O}}(q)$. After that the factor base itself can be determined in an expected time of $\tilde{\mathcal{O}}(q^r)$.

By Proposition 2, each iteration of Step 2 can be performed in a randomized polynomial time in $\log(q)$ (and one can thereby guarantee that Dis selected uniformly at random among all possible divisors). Heuristically, it seems reasonable to assume that one needs $\approx g! \cdot q^r \cdot q^{(1-r) \cdot g}$ iterations in Step 2 (leading to an asymptotic running time of $\tilde{O}(q^r \cdot q^{(1-r) \cdot g}))$ for "most" curves and "most" choices of the factor base. As the linear algebra can be performed in $\tilde{O}(q^{2r})$ operations, we obtain a heuristic running time of

$$\tilde{O}(q^{r+(1-r)\cdot g} + q^{2r}) \tag{1}$$

for "most" runs of the algorithm.

For r := 1 - 1/(g + 1), heuristically, the running time of the relation search then balances asymptotically with the running time of the linear algebra part (up to logarithmic factors) for "most" runs of the algorithm.

It is however conceivable that for some choices of the factor base, the number of iterations needed is much worse even though other choices of the factor base would lead to the stated number of iterations. We thus propose the following *modification* of the algorithm:

A modification

Whenever in Step 2 after $2 \cdot g! \cdot q^r \cdot q^{(1-r) \cdot g}$ iterations R does not have more rows than columns, the algorithm outputs "failure" and terminates.

A crucial *heuristic assumption* is now that there exists some P < 1 such that for r = 1 - 1/(g + 1), for essentially all inputs, this happens with probability < P.

Under this assumption (and some further assumptions for the linear algebra part), we obtain:

Heuristic Result 6 With the algorithm by Gaudry and Harley for arbitrary curves with the modification stated above, one can asymptotically for varying q solve essentially all instances of the DLP in degree 0 class groups of curves of a fixed genus g over \mathbb{F}_q , represented by plane models of bounded degree, in an expected time of $\tilde{O}(q^{2-\frac{2}{g+1}})$.

Remark 7 If one modifies the relation generation step (Step 3) along the lines of the relation generation step of the general algorithm in [10, Section 3], one can prove that the algorithm has the stated expected running time when applied to hyperelliptic curves \mathcal{C}/\mathbb{F}_q in imaginary representation with $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q) = \langle a \rangle$.

Remark 8 One can combine the algorithm with a double large prime variation as presented in [15] and [28]. Although the heuristic analysis in [15] is only stated for hyperelliptic curves, it also applies in the more general setting we consider here. One obtains that heuristically, one can solve essentially all instances of the DLP in degree 0 class groups of curves of a fixed genus g over \mathbb{F}_q represented by plane models of bounded degree in an expected time of $\tilde{O}(q^{2-\frac{2}{g}})$. This was already mentioned in the introduction.

Implicit constants and practical aspects

For practical purposes one should consider the following specifications and modifications of the algorithm and the way the size of the factor base is determined. Asymptotically, up to logarithmic factors, these modifications do not alter the running time.

- 1. One can construct a factor base in an expected time of $O(q^r)$ by considering points with prescribed x-coordinate (instead of first enumerating $\mathcal{C}(\mathbb{F}_q)$).
- 2. The running time of the relation generation step (Step 2 of the algorithm) depends critically on the efficiency of the arithmetic in the class groups. Even though every algorithm to compute Riemann-Roch spaces which fulfills the conclusions of Proposition 2 can be used to obtain Heuristic Result 6, for practical purposes some algorithms are more suited than others.
- 3. Heß ([20]) gave a fast general purpose algorithm to calculate Riemann-Roch spaces. His algorithm can be viewed as a generalization of the algorithm by Cantor ([6]) for computations in class groups of hyperelliptic curves in imaginary quadratic representation. Let us give some information on his approach:

The algorithm relies in an *ideal theoretic representation* of divisors in a *function field theoretic* setting:

One chooses a non-constant element $\tilde{X} \in \mathbb{F}_q(\mathcal{C})$ such that the extension $\mathbb{F}_q(\mathcal{C})|\mathbb{F}_q(\tilde{X})$ is separable. After that one needs to calculate bases of the so-called *finite order*, the integral closure of $\mathbb{F}_q[\tilde{X}]$, and the *infinite order*, the integral closure of \mathcal{O}_{∞} , the local ring of the place "infinity". These bases can be calculated in polynomial time $\log(q)$ and in the total degree of f, and this has to be done only once.

For the arithmetic itself one represents divisors by *two ideals* with respect to the two orders. (The ideals themselves are represented by a $\mathbb{F}_q[\tilde{X}]$ - and a \mathcal{O}_{∞} -basis respectively). The running time thereby depends crucially on the degree of the extension $\mathbb{F}_q(\mathcal{C})|\mathbb{F}_q(\tilde{X})$.

According to [18] for every curve \mathcal{C}/\mathbb{F}_q , there exists a suitable extension $\mathbb{F}_q(\mathcal{C})|\mathbb{F}_q(\tilde{X})$ such that the arithmetic in $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ in ideal representation takes $O(g^4)$ field operations for varying g (after some precomputation). For comparison, the arithmetic (in ideal representation) in class groups of hyperelliptic curves in imaginary representation takes $O(g^2)$ field operations.

4. Another approach for fast arithmetic in class groups of curves was provided by Khuri-Makdisi ([23]). His algorithms of geometrical nature rely on an embedding of a symmetric power of the curve into a Grassmannian. The computations in the class group then reduce to linear algebra computations. This approach can also be used to calculate bases of Riemann-Roch spaces (which can then in turn be used to calculate "free representations" of divisors). Similarly to [20], the arithmetic (in the specific representation used in the algorithm) can be performed in $O(g^4)$ field operations.

5. Before one calculates a free representation of a divisor, one should first check whether the projection to some coordinate axis splits completely. Like this, one avoids unnecessary factorizations of polynomials. (Note that if the divisor splits completely, so does its projection, and the converse holds essentially always in practice.)

If one uses the ideal arithmetic, one should include 1 into the basis of the finite order. Then one can perform the arithmetic in such a way that the "finite" ideal contains a univariate polynomial in \tilde{X} which gives the \tilde{X} -coordinates. (Note that this is a direct generalization of the approach in [14] and [33] which uses the so-called "Mumford representation".)

- 6. For special classes of curves, one can use more efficient algorithms to perform the arithmetic. For example, for non-hyperelliptic genus 3 curves, one can use the algorithms in [4] and [12].
- 7. If one chooses r = 1 1/(g+1), the relation generation part is usually much more costly than the linear algebra part. One should rebalance the two by including more elements into the factor base. (Under certain conditions on g and q, it might even pay off to include all \mathbb{F}_q -rational points into the factor base, as was originally suggested by Gaudry in [14].)

3 Our variant of the algorithm by Gaudry and Harley

We now come to our modifications of this algorithm. As above, let \mathcal{C}/\mathbb{F}_q be a curve of genus ≥ 1 represented by a plane model \mathcal{C}_{pm} of degree d. Let $\varphi: \mathcal{C} \longrightarrow \mathcal{C}_{pm}$ be the canonical projection, and let us denote the associated morphism $\mathcal{C} \longrightarrow \mathbb{P}^2$ also by φ . As stated in the introduction, the affine part of \mathcal{C}_{pm} is denoted by \mathcal{C}_a . Let $\mathcal{C}_{a,ns}$ be the non-singular part of \mathcal{C}_a .

In the sequel, we identify zero-dimensional subschemes on \mathcal{C} with effective divisors on \mathcal{C} .

As stated in the introduction, we use the same notation for a function on \mathbb{A}^2 , its restriction to \mathcal{C}_a and the induced element in $\mathbb{F}_q(\mathcal{C})$. Similarly, we use the same notation for an element of $\Gamma(\mathbb{P}^2, \mathcal{O}(1))$ (the space of "homogeneous coordinates" on \mathbb{P}^2) and its pull-backs to $\Gamma(\mathcal{C}, \varphi^*(\mathcal{O}(1)))$. (For example, we write Z instead of $\varphi^*(Z)$.) To distinguish the divisor of zeros of elements of $\Gamma(\mathcal{C}, \varphi^*(\mathcal{O}(1)))$ and their pull-back to $\Gamma(\mathcal{C}, \varphi^*(\mathcal{O}(1)))$, we write div_C for the latter. (See [17, II. §7] for information on the divisor of zeros.)

Definition 9 Let $D_{\infty} := \operatorname{div}_{\mathcal{C}}(Z) \in \operatorname{Div}(\mathcal{C})$ be the divisor of zeros of Z on \mathcal{C} .

Remark 10 If C is non-singular "at infinity", D_{∞} is the intersection of C_{pm} with the hyperplane defined by Z = 0 (regarded as divisor on C). In general, it is equal to $\varphi^{-1}(\operatorname{div}(Z))$, the scheme-theoretic preimage (or pull-back) in C of the hyperplane Z = 0, with other words, the scheme-theoretic preimage of the scheme-theoretic intersection of C_{pm} with the hyperplane Z = 0 to C.

The same remark holds for divisors defined by pull-back of other "homogeneous coordinates" to \mathcal{C} .

As C_{pm} has degree d, so has the divisor D_{∞} . We remark that even though we use D_{∞} in the following description of our modifications, it does not occur explicitly in the calculations.

We fix a point $P_0 \in \mathcal{C}_{a,ns}(\mathbb{F}_q)$ which will serve as "base point" as above. Instead of searching for relations of the form

$$\sum_{j} r_{i,j}([F_j] - [P_0]) = \alpha_i a + \beta_i b$$

we just search for *two* relations involving a non-trivial right-hand side:

$$\sum_{P \in \mathcal{C}_{a,ns}(\mathbb{F}_q)} k_P([P] - [P_0]) = \alpha d$$

as well as

$$\sum_{\in \mathcal{C}_{a,ns}(\mathbb{F}_q)} l_P([P] - [P_0]) = \beta b$$

P

(with $\alpha, \beta \neq (\mathbb{Z}/N\mathbb{Z})^*$ and $k_P, l_P \in \mathbb{N}_0$). Only after we have found these relations, we fix the factor base which is a subset of $C_{a,ns}(\mathbb{F}_q)$. We thereby guarantee that the points involved in these relations are in the factor base. The first two rows of the "relation matrix" R are then made up by the relations for αa and βb respectively.

All other relations are just between the elements of the factor base:

$$\left(\sum_{j} r_{i,j}[F_j]\right) - [D_{\infty}] = 0 \qquad (\text{with } \sum_{j} r_{i,j} = d).$$

We find such relations by intersecting C_{pm} with lines which are defined by passing through two points of the factor base. (For more information see subsection "Correctness" below.)

An important and subtle aspect of our algorithm is that the factor base has to chosen appropriately in order that we are able to generate enough relations to solve the DLP. We discuss this below.

In contrast to the previous algorithm, in this algorithm, we restrict ourselves to the the case that $N = \#\langle a \rangle$ is square-free. (This is because we cannot guarantee that the "lifting procedure" in [10, Section 4] can be carried out successfully.) This is however not an obstacle because one can with repeated calls to the algorithm also tread instances of the DLP with arbitrary N.

In the algorithm, as stated in the introduction, we denote the coordinate ring of \mathbb{A}^2 by $\mathbb{F}_q[X_a, Y_a]$, and we denote an "affine line" by L_a .

Our variant of the Gaudry-Harley algorithm

Input. A curve \mathcal{C}/\mathbb{F}_q of genus ≥ 1 , represented by a plane model, with a fixed point $P_0 \in \mathcal{C}_a(\mathbb{F}_q)$ (used for the representation of the elements of $\mathrm{Cl}^0(C/\mathbb{F}_q)$), $a, b \in \mathrm{Cl}^0(C/\mathbb{F}_q)$ such that $b \in \langle a \rangle \ a, b \in \mathrm{Cl}^0(C/\mathbb{F}_q)$ such that $b \in \langle a \rangle$, the number $N := \#\langle a \rangle$, which is assumed to be square-free, and a positive rational number r < 1.

Output. An x such that $x \cdot a = b$ or "failure".

Let F(X, Y, Z) be the defining homogeneous polynomial of the plane model.

- 1. Calculate the non-singular part $C_{a,ns}$ of C_a . If $P_0 \notin C_{a,ns}(\mathbb{F}_q)$, substitute it by a point in $C_{a,ns}(\mathbb{F}_q)$ (also called P_0) and change the representations of a and b.
- Choose randomly and uniformly α ∈ (Z/NZ)\{0} and calculate a divisor D with [D] - [gP₀] = αa, until α ∤ N and D splits completely into points of C_{a.ns}(F_q).

Choose randomly and uniformly $\beta \in (\mathbb{Z}/N\mathbb{Z}) \setminus \{0\}$ and calculate a divisor D with $[D] - [gP_0] = \beta b$, until $\beta \nmid N$ and D splits completely into points of $\mathcal{C}_{a,ns}(\mathbb{F}_q)$.

Fix a "factor base" *F* ⊂ *C*_{n,s}(𝔽_q) with #*F* = ⌈q^r⌉, such that *F* contains the points for the relations for *αa* and *βb*, selected uniformly at random from all possible subsets. (If no such set exists, output "failure" and terminate.)

Store the "left-hand sides" of the relations in Step 2 as the first two rows of a sparse matrix R.

4. Let f(X_a, Y_a) ← F(X, Y, 1) ∈ F_q[X_a, Y_a]. Construct a sparse matrix R over Z/NZ as follows: For all pairs (i, j) with i < j ≤ #F do
{
Calculate the defining polynomial l of the line L_a : Y_a = l(X_a)
passing through F_i and F_j (in A²). (For simplicity we assume
that X_a(F_i) ≠ X_a(F_j). If the two are equal, a simple
modification can be applied.)
If the polynomial f(X_a, l(X_a)) splits completely and has degree d,

```
{
factorize it and
calculate the Y_a-coordinates corresponding to its roots.
If all points obtained in this way lie in the factor base,
{
let (r_j)_j be the sparse vector whose non-zero entries
correspond to these points with values being the
multiplicities from the polynomial factorization.
Store (r_j)_j as a new row in the sparce matrix R.
}}}
```

- 5. Calculate a random vector $\gamma \in \ker(R^t)$ with algorithms from sparse linear algebra (cf. previous algorithm).
- 6. If $\gamma_1 \in (\mathbb{Z}/N\mathbb{Z})^*$, output $x = -\frac{\gamma_1 \alpha}{\gamma_2 \beta} \in \mathbb{Z}/N\mathbb{Z}$. Otherwise output "failure".

Correctness

It follows from the following lemma that each row of the matrix R corresponds to a relation in $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$.

Lemma 11 Let $L : \lambda X + \mu Y + \nu Z = 0$ be a line in $\mathbb{P}^2/\mathbb{F}_q$, and assume that L intersects \mathcal{C}_{pm} in non-singular points. Let D be the scheme-theoretic intersection of L with \mathcal{C}_{pm} , considered as divisor on \mathcal{C} . Then the function $\lambda Y_a + \mu X_a + \nu \in \mathbb{F}_q(\mathcal{C})$ has the principal divisor $D - D_{\infty}$.

Proof. The function $\lambda X_a + \mu Y_a + \nu = \lambda \frac{X}{Z} + \mu \frac{Y}{Z} + \lambda \frac{Z}{Z} \in \mathbb{F}_q(\mathcal{C})$ has the divisor $\operatorname{div}_{\mathcal{C}}(\lambda Y + \mu X + \nu Z) - \operatorname{div}_{\mathcal{C}}(Z)$ which by definition is $D - D_{\infty}$. \Box

Say that the affine part of such a line L is given by $L_a : Y_a = l(X_a)$ as in the algorithm. (As stated in the algorithm, if $\lambda = 0$, an easy modification can be applied.) Let D_a be the "affine part" of D, i.e. we disregard points "at infinity" in the support of D. Then the roots of $f(X_a, l(X_a))$ in $\overline{\mathbb{F}}_q$ give the X_a -coordinates of the support of D_a over $\overline{\mathbb{F}}_q$, and the multiplicities correspond to each other.

Now, in Step 4 of the algorithm, we consider only such lines which intersect C only in $C_{a,ns}$. Indeed, we explicitly rule out singular intersection points. We rule out intersection points "at infinity" by demanding that $\deg(f(X_a), l(X_a)) = d$ which implies that $\deg(D_a) = \deg(\mathcal{C}) = \deg(D_{\infty})$, i.e. $D = D_a$.

It follows that a vector $(r_j)_j$ calculated in Step 4 corresponds to a relation $\sum_i (r_j[F_j]) - [D_\infty] = 0$ in $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$.

One easily sees that any $\gamma \in \ker(\mathbb{R}^t)$ leads to an equation

$$\gamma_1 \alpha a + \gamma_2 \beta b = 0.$$

This proves the correctness of the algorithm.

Complexity

We now discuss the *complexity* of the various steps of the algorithm. We do the estimations for a *fixed degree* d and varying q.

The singular locus of C_a can be determined by considering the common roots of the two partial derivatives of the defining polynomial f. This can be performed in a bounded number of field operations in \mathbb{F}_q . We note that there not more than (d-1)(d-2)/2 singular points in $C_a(\mathbb{F}_q)$ (cf. [17, I,7 Ex. 7.2; IV.1 Ex. 1.8]).

In Step 2 of the algorithm, the representatives of αa and βb can be calculated in randomized polynomial time by Proposition 2. The probability that D splits completely is on heuristic grounds 1/g! and thus constant for varying q. As $g \leq (d-1)(d-2)/2$, this indicates that this step can also be performed in a bounded number of field operations in \mathbb{F}_q .

The factor base can be determined in an expected time of O(q) as in the previous algorithm.

Each iteration in Step 4 can be performed in randomized polynomial time in $\log(q)$. As there are not more than q^{2r} unordered pairs of elements in the factor base, the total running time of Step 4 is in $\tilde{O}(q^{2r})$.

The linear algebra in Step 5 can also be performed in time $O(q^{2r})$, and the final step takes a negligible amount of time.

All in all, we have the following heuristic result:

The algorithm terminates in a time of $\tilde{O}(q^{2r})$.

We now come to the important question under which conditions on g, q and r we should expect that the algorithm in fact solves the DLP.

In order that one can solve the DLP with linear algebra in Step 5, it is necessary and sufficient that the second row linearly depends on the other rows of the matrix. Because it seems difficult to obtain a theoretical results on this question, we studied it with experiments.

A modification for experiments

For an experimental study, we modified the relation search slightly: The relation search is terminated if a predefined number of *different* relations has been obtained. (For this one should sort the relations while generating them, for example with a binary search tree ([25, 6.2.2]), and only later build the relation matrix.)

An experimental observation

We have conducted experiments in order to see how many relations we have to generate in order that the second row of the relation matrix is linearly dependent on the other rows. In our experiments we made the observation that with very few exceptions this is the case if we have generated as many different relations as there are elements in the factor base. (See the "Experimental results" at the end of this section for further information.)

Based on these experiments, we conclude that it should in practice suffice to generate as many relations as there are elements in the factor base for "most" curves. For our heuristic analysis below, we make the following assumption.

Heuristic Assumption 12 There exists a function c = c(d, q, r) which for fixed d is bounded by a polynomial function in $\log(q^r)$ such that for essentially all inputs it holds that for at least half the possible choices of the factor base, whenever the matrix R has $c(d, q, r) \cdot q^r$ rows, the second row is linearly dependent on the other rows.

Terminology

We call a line $L : \lambda X + \mu Y + \nu Z = 0$ completely split if $\operatorname{div}_{\mathcal{C}}(\lambda X + \mu Y + \nu Z)$, the divisor of zeros of $\lambda X + \mu Y + \nu Z$ on \mathcal{C} , is completely split. (Following [33] one could call this "potentially smooth".)

Note that a line L is completely spit if and only if principal divisor of the associated function $\lambda X_a + \mu Y_a + \nu \in \mathbb{F}_q(\mathcal{C})$ has the form $\sum_k P_k - D_\infty$ for some $P_k \in \mathcal{C}(\mathbb{F}_q)$.

We note that if a line intersects C in non-singular points, then it is completely split if and only if the intersection consists of points in $C(\mathbb{F}_q)$. More generally, a line is completely split if and only if the scheme-theoretic preimage of its scheme-theoretic intersection with C_{pm} to C consists of points in $C(\mathbb{F}_q)$.

Moreover, if $\mu = 1$ and the intersection lies in $\mathcal{C}_{a,ns}$, then the function / line is completely split if and only if $f(X_a, \lambda X_a + \mu)$ has degree d and splits completely into elements of $\mathcal{C}_{a,ns}(\mathbb{F}_q)$.

We call a function or the corresponding line *smooth* if additionally, all intersection points lie in the factor base (which by definition is a subset of $C_{a,ns}(\mathbb{F}_q)$).

The optimal size of the factor base

We now come to the determination of the optimal value of r (and thus of the optimal size of \mathcal{F}) under the condition that we can expect to generate enough linearly independent relations.

In Step 4 of the algorithm, we consider all smooth lines passing through at least two points of the factor base, and each such line gives rise to a different relation which is included as a row in the matrix R. The goal is now to give a heuristic estimate of the number of such lines (i.e. the number of different relations generated).

Below, we give a heuristic estimate of the total number of smooth lines. The difference between these numbers is equal to the number of smooth lines passing through exactly one point of $\mathcal{C}_{pm}(\mathbb{F}_q)$ (which is then an element of the factor base). This number is asymptotically negligible, as shown in the following lemma.

Lemma 13 The number of lines over \mathbb{F}_q which intersect \mathcal{C}_{pm} in one point of $\mathcal{C}_{a,ns}(\mathbb{F}_q)$ and no other point of $\mathcal{C}_{pm}(\overline{\mathbb{F}}_q)$ is bounded by $\#\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)[d]$, the order of the d-torsion subgroup of $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$, which in turn is bounded by d^{2g} .

Proof. Let $P \in \mathcal{C}_{a,ns}(\mathbb{F}_q)$. Then a line intersects \mathcal{C}_{pm} exactly in P (and no further point of $\mathcal{C}_{pm}(\overline{\mathbb{F}}_q)$) if and only if the corresponding function has the divisor $dP - D_{\infty}$. This implies that first, there is at most one line passing only through P (and no further point of $\mathcal{C}_{pm}(\overline{\mathbb{F}}_q)$), and second the points P with this property are exactly the points of $\mathcal{C}_{a,ns}(\mathbb{F}_q)$ for which dP is linearly equivalent to D_{∞} .

Let P_1, P_2 be two points with this property. We have $[d(P_1 - P_2)] = 0 \in Cl^0(\mathcal{C}/\mathbb{F}_q)$. But $P_1 \neq P_2$ and $g(\mathcal{C}) \geq 1$ implies that $[P_1] - [P_2] \neq 0$. Thus $[P_1] - [P_0] \neq [P_2] - [P_0]$.

This means that there cannot be more lines passing through one point of $\mathcal{C}_{pm}(\mathbb{F}_q)$ and no further point of $\mathcal{C}_{pm}(\overline{\mathbb{F}}_q)$ as there are elements of *d*-torsion in $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$. It is well-known that $\#\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)[d]$ divides d^{2g} . \Box

We now try to estimate the total number of smooth lines (as always for fixed degree and varying q).

By the bounds of Hasse-Weil, the probability that an effective divisor of degree d splits completely is asymptotically equal to 1/d!.

Heuristically, the probability that a line $L : \lambda X + \mu Y + \nu Z = 0$ is completely split is approximately equal to this probability, i.e. to 1/d!.

The probability that a completely split line is smooth is on heuristic grounds approximately equal to the probability that all entries of a tuple of d elements of $\mathcal{C}(\mathbb{F}_q)$ lie in the factor base, and this is approximately

$$\left(\frac{\#\mathcal{F}}{q}\right)^d \approx q^{(r-1)\cdot d}$$
.

In total, there are approximately q^2 lines $L : \lambda X + \mu Y + \nu Z = 0$. This means that heuristically, one should expect that the number of smooth lines is approximately

$$q^{2} \cdot \frac{1}{d!} \cdot q^{(r-1) \cdot d} = \frac{1}{d!} \cdot q^{r \cdot d - (d-2)} .$$
⁽²⁾

Together with Lemma 13 we conclude that heuristically, (2) gives the approximate number of different relations generated in Step 4 of the algorithm.

This number should be larger or equal than $c(d, q, r) \cdot \#\mathcal{F} \approx (q, r, d) \cdot q^r$, which means that we should have

$$q^{r} \ge (c(d,q,r) \cdot d!)^{\frac{1}{d-1}} \cdot q^{\frac{d-2}{d-1}} = (c(d,q,r) \cdot d!)^{\frac{1}{d-1}} \cdot q^{1-\frac{1}{d-1}}.$$
 (3)

We remark that this means in particular that we should have

$$q \ge d! \tag{4}$$

if we apply our algorithm. For our asymptotic study for curves of fixed degree, this condition is satisfied for all but finitely many input values.

The minimal choice for r is given by

$$q^{r} = (c(d,q,r) \cdot d!)^{\frac{1}{d-1}} \cdot q^{1-\frac{1}{d-1}}, \text{ i.e.}$$
(5)

$$r = 1 - \frac{1}{d-1} \cdot \left(1 - \frac{\log(c(d,q,r)) + \log(d!)}{\log(q)} \right) .$$
 (6)

Under our heuristic analysis concerning the running time and our assumptions on c(d, q, r) and on the number of smooth lines, we conclude:

Heuristic Result 14 With our variant of the Gaudry-Harley algorithm, one can asymptotically for varying q solve essentially all instances of the DLP in degree 0 class groups of curves represented by plane models of a fixed degree d over \mathbb{F}_q in an expected time of $\tilde{O}(q^{2-\frac{2}{d-1}})$.

Remark 15 One can use the heuristic estimate (2) of the number of smooth lines to derive a heuristic probability that a line passing through two points of the factor base is smooth:

Heuristically and on the basis of (2), the number of pairs (F_i, F_j) of elements of the factor base with i < j defining smooth lines is

$$\frac{d(d-1)}{2} \cdot \frac{1}{d!} \cdot q^{r \cdot d - (d-2)} = \frac{1}{2(d-2)!} \cdot q^{r \cdot d - (d-2)}$$

Dividing this number by $q^{2r}/2$, the approximate total number of such pairs, we obtain that with a probability of approximately

$$\frac{1}{(d-2)!} \cdot q^{(r-1)\cdot(d-2)} , \qquad (7)$$

a pair of elements of the factor base defines a smooth line.

As one might expect, this probability is equal to the asymptotic probability that a divisor of degree d-2 is smooth.

Experimental results

We implemented the algorithm in Magma with the "modification for experiments" stated above. We performed experiments with curves of genus 3 to 5 given by equations of degree 4 to 6 over fields \mathbb{F}_{2^n} with n up to 12. To facilitate the calculation of the order of the class group, all curves were obtained by base change from curves over \mathbb{F}_2 . Various possible choices for the size of the factor base were tried. We thereby included more and more different rows into the relation matrix and calculated the rank modulo the prime divisors of N.

Let us fix some notation to describe our experimental results:

Let $R_{0,t}$ be the relation matrix obtained just from Step 4 (i.e. without the relations for αa and βb after having included t different relations, let $R_{1,t}$ be the relation matrix without the second row (i.e. without the relation for βb), and let $R_{2,t}$ be the full relation matrix.

Let us fix a prime divisor ℓ of $N = \#\langle a \rangle$ and $\#\langle b \rangle$ and regard all matrices modulo ℓ . (We keep the same notations.) Note that we always have $\operatorname{Rank}(R_{2,t}) \leq \operatorname{Rank}(R_{1,t}) + 1$, and we have $\operatorname{Rank}(R_{2,t}) = \operatorname{Rank}(R_{1,t})$ if and only if we can solve the DLP modulo ℓ by linear algebra via $R_{2,t}$.

Now, in all except in one case (in which we did calculations modulo 5), in our experiments, we made the following observation:

Experimental Observation 16 As long as $\operatorname{Rank}(R_{1,t}) < \operatorname{Rank}(R_{2,t})$, the rows of $R_{0,t+1}$ (and of $R_{1,t+1}$) are linearly independent. (That is, as long as the row for βb is linearly independent of the rows of $R_{1,t}$, a "new row" is linearly independent of the rows of $R_{1,t}$.)

This observation means in particular that if $t = \#\mathcal{F}$, we can expect to be able to calculate the DLP modulo ℓ by linear algebra via the matrix $R = R_{2,t}$.

An example

Here is a particular (typical) example: We considered the curve given by

$$Y^{5} + (X^{2} \cdot Z + X \cdot Z^{2} + Z^{3}) \cdot Y^{2} + (X^{3} \cdot Z + X \cdot Z^{4}) \cdot Y + X^{5} + X^{2} \cdot Z^{3} = 0$$

over $\mathbb{F}_{2^{12}}$. This curve has genus 5, and the order of its class group is $N = 5^4 \cdot 37 \cdot 61 \cdot 277 \cdot 337 \cdot 8419249$. The 5⁴-torsion subgroup of the class group has the structure $\mathbb{Z}/5^3\mathbb{Z} \times \mathbb{Z}/5\mathbb{Z}$, thus the 5-torsion subgroup has rank 2.

We deliberately did not use the degree 5 model but used instead several degree 6 models obtained with the algorithm in Section 5 for our experiments.

Our heuristic analysis for the number of smooth lines indicates that we would need to include 2893 elements into the factor base in order to obtain as many different relations as elements of the factor base.

We included 2884 elements into the factor base and obtained 2902 different relations in Step 4 of the algorithm. (This supports the heuristic analysis concerning smooth relations.)

For all prime factor of N except 5, we obtained:

The minimal t with $\operatorname{Rank}(R_{1,t}) = \operatorname{Rank}(R_{2,t})$ was 2877. For this t, all rows of the matrix $R_{0,t}$ were linearly independent and the number of non-zero columns was 2878.

Modulo 5, the minimal t with $\operatorname{Rank}(R_{1,t}) = \operatorname{Rank}(R_{2,t})$ was 2876, again all rows of $R_{0,t}$ where linearly independent and the number of non-zero columns was also 2878.

We note that modulo all primes, $\operatorname{Rank}(R_{0,t})$ was equal to the number of non-zero columns minus the rank of the ℓ -torsion subgroup of $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$. This relation held most of the times when we included roughly $d!^{1/(d-1)} \cdot q^{1-1/(d-1)}$ elements into the factor base.

Implicit constants and practical aspects

We give some information on the implicit constants in our heuristic result as well as some suggestions of variations of the algorithm for practical purposes. We thereby assume that c(d, q, r) can be set to be 1.

- 1. One should apply the algorithm only if q > d!.
- 2. For curves represented by plane models of degree roughly equal to the genus of the curve, the running time of Step 2 is negligible. However, in the extreme case, a plane model might itself be non-singular such that g = (d-1)(d-2)/2. In this situation, Step 2 might very well dominate the running time.
- 3. In every iteration of Step 4 one has to calculate the polynomial $f(X_a, l(X_a))$. For varying d, this can clearly be achieved in $O(d^3)$ field operations. There is however an easy variant of Step 4 which is more efficient: One first fixes an element F_i of the factor base and performs a coordinate transformation such that F_i has the coordinates (0, 0). (I.e. one calculates the polynomial $\tilde{f}(X_a, Y_a) := f(X_a + X_a(F_i), Y_a + Y_a(F_i))$.) This can be achieved in $O(d^3)$ field operations. Then all lines though (F_i, F_j) with j > i are given by $Y_a = l(X_a) = \lambda X_a$ for some λ . The polynomial $\tilde{f}(X_a, \lambda X_a)$ can be calculated with $O(d^2)$ field operations.
- 4. Let $k_{\text{test}}(d,q)$, $k_{\text{fac}}(d,q)$ be the times in field multiplications to test whether a polynomial of degree d over \mathbb{F}_q splits completely and the time to factorize a completely split polynomial of degree d over \mathbb{F}_q respectively. Under the condition that the time to access the storage

is negligible, for our choice of r, the running time of Step 4 can *roughly* be estimated as

$$\left(\binom{d+2}{2} + k_{\text{test}}(d,q) + \frac{k_{\text{fac}}(d,q)}{(d-2)!} \right) \cdot d!^{\frac{2}{d-1}} \cdot q^{2-\frac{2}{d-1}}$$
(8)

field multiplications in \mathbb{F}_q . If Step 5 is performed with Lanczos' algorithm as in [9], we can expect a time of roughly 2d + 4 multiplications modulo the factors of N plus the time to factor N. If N is not highly composite, this is considerably less than the time needed for Step 4. (That changes however if the access to the storage becomes more costly.) Step 2 is negligible provided that g and d are roughly equal.

We remark that the behavior in relation to d is better than one might expect. (It does not involve a "huge d!-term".) An analogous remark also applies to the algorithm presented in the previous section if one balances the relation generation and the linear algebra part optimally (and $d! \leq q$) (cf. Point 2 in the "practical aspects" in the previous section).

5. As the relation generation and the linear algebra part take up to logarithmic factors the same running time, one cannot improve the asymptotic running time (up to logarithmic factors) by enlarging the factor base and terminating the relation search early.

A practical improvement is however possible: If one iterates over all pairs of the factor base, most smooth lines are selected d(d-1)/2 times. By enlarging the factor base by a constant factor, one drastically reduce the average number of times a line is reselected.

Historical remark

As is apparent from the headline of this section, we view our algorithm as a variant of the algorithm by Gaudry and Harley. The algorithm can however also be viewed as an adaption of the algorithm by Adleman, DeMarrais and Huang ([1]) for index calculus in class groups of hyperelliptic curves of large genus. Indeed, our main variation of the algorithm by Gaudry and Harley, the finding of relations by considering principal divisors (instead of considering linear combinations of the input elements), is an essential ingredient of the algorithm by Adleman, DeMarrais and Huang. We note that the heuristic analysis in [1] relies on a very similar heuristic assumption as Heuristic Assumption 12.

4 Our variant of the algorithm by Gaudry and Harley with double large prime variation

We use [15] as reference for the Gaudry-Harley algorithm with double large prime variation. We assume that the reader is familiar with this work and freely use the same notations.

Following the description in [15], it is possible to combine our algorithm with a double large prime variation (thereby using all points in $C_{a,ns}(\mathbb{F}_q)$ as "large primes"). We note however that just as in our adaption of the Gaudry-Harley algorithm itself, we need "enough" elements in the factor base.

We recall that there are *two* double large prime algorithms in [15]: In the first one, for each relation involving two large primes (and otherwise elements from the factor base), an edge in the graph of large prime variation is inserted, provided that this edge does not lead to a cycle. In the second one (which is called "simplified algorithm" in [15]), only the connected component of the graph corresponding to relations with one large prime is considered. As the cycles which lead to relations over elements of the factor base are not constructed, this second graph is in fact a tree; we call it the *tree of large prime variation*. The analysis in [15] is carried out for the simplified algorithm whereas in practice one should use the first algorithm. We also recall that the analysis in [15] involves some heuristic considerations.

When trying to adapt this analysis to our situation, ones encounters a problem: Whereas the probabilities to generate an Full, FP or PP relation are constant during the execution of the algorithm, the probabilities to generate a *new* Full, FP and PP relation are not. Indeed, the probabilities to generate new Full and FP relations drop during the execution of the algorithm. This is due to the fact that every Full relation can be obtained via lines through d(d-1) different pairs from the factor base whereas an FP relation can be obtained via (d-1)(d-2) different pairs from the factor base and an PP relation can be obtained via (d-2)(d-3) different pairs of the factor base.

To encounter this problem and in order to avoid an heuristic assumption on the final matrix as Heuristic Assumption 12 in the previous section, we proceed as follows:

We first construct the tree of large prime variation using only relations with *one* large prime. After that we extend the tree with relations involving two large primes. In both cases we proceed similarly as in the previous algorithm. Finally, we use the tree of large prime variation to express linear combinations of the input elements over the factor base and to construct the relation matrix R. For exactitude, we can thereby use ideas from [10].

The algorithm is as follows:

Our variant of the "simplified" Gaudry-Harley algorithm with double large prime variation

Input. A curve \mathcal{C}/\mathbb{F}_q of genus ≥ 1 , represented by a plane model of degree ≥ 4 , with a fixed point $P_0 \in \mathcal{C}(\mathbb{F}_q)$ (used for the representation of the elements of $\mathrm{Cl}^0(C/\mathbb{F}_q)$), $a, b \in \mathrm{Cl}^0(C/\mathbb{F}_q)$ such that $b \in \langle a \rangle \ a, b \in \mathrm{Cl}^0(C/\mathbb{F}_q)$ such that $b \in \langle a \rangle$, the number $N := \#\langle a \rangle$, and a positive rational number r < 1. Output. An x such that $x \cdot a = b$.

Let F(X, Y, Z) be the defining homogeneous polynomial of the plane model.

- 1. Calculate the non-singular part $\mathcal{C}_{a,ns}$ of \mathcal{C}_a . If $P_0 \notin \mathcal{C}_{a,ns}(\mathbb{F}_q)$, substitute it by a point in $\mathcal{C}_{a,ns}(\mathbb{F}_q)$ (also called P_0) and change the representations of a and b.
- 2. Fix a "factor base" $\mathcal{F} \subset \mathcal{C}_{n,s}(\mathbb{F}_q)$ with $\#\mathcal{F} = \lceil q^r \rceil$, selected uniformly at random from all possible subsets. (If no such set exists, output "failure" and terminate.)
- 3. [Construction of the tree of large prime variation]
 - (a) Iterate over all pairs of the factor base as in the previous algorithm and construct a tree of large prime variation by only considering relations with *one* large prime.
 - (b) Iterate again over all pairs of the factor base and enlarge the tree of large prime variation (using relations with two large primes, but only including edges connected to the original connected component).
- 4. Construct a sparse matrix R over $\mathbb{Z}/N\mathbb{Z}$ as follows: Repeat {

Select randomly and uniformly $\alpha, \beta \in \mathbb{Z}/N\mathbb{Z}$. Calculate a divisor D such that $[D] - [gP_0] = \alpha a + \beta b$. If D splits completely, i.e. $D = \sum_k P_k$ for some $P_k \in \mathcal{C}(\mathbb{F}_q)$, and all P_k are elements of the factor base or vertices of the tree of large prime variation, determine the corresponding relation and include a corresponding new row in the matrix R.

- $\}$ until R has more rows than columns.
- 5. Try to find a non-trivial element $\gamma \in \ker(R^t)$ with algorithms from sparse linear algebra. If the procedure fails, return to Step 4 and include some more relations into R.
- 6. If $\sum_i \gamma_i \beta_i \in (\mathbb{Z}/N\mathbb{Z})^*$, output

$$x = -\frac{\sum_{i} \gamma_i \alpha_i}{\sum_{i} \gamma_i \beta_i} \in \mathbb{Z}/N\mathbb{Z} ,$$

otherwise return to Step 4.

Complexity

Generation of the sparse matrix and linear algebra

We have presented Step 4 in its most simple form. By using the procedure of Step 2 in the algorithm in [10], one can under certain conditions rigorously analyze the running time of this step. The corresponding theoretical result is:

Proposition 17 Let us (additionally to g and d) fix a function $f : \mathbb{N} \longrightarrow \mathbb{R}$, and a positive rational numbers e < 1. Then the following holds:

Consider curves \mathcal{C}/\mathbb{F}_q of genus g represented by plane curves of degree d such that $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ is cyclic. For these curves it holds that if

- a factor base $\mathcal{F} = \{F_1, F_2, \ldots\}$ of size f(r) has been selected and
- a tree of large prime variation involving more than $e \cdot q$ elements has already been constructed in which
 - all edges correspond to relations of the form

$$\left(\sum_{j} r_j[F_j]\right) + [P] - [D_\infty] = 0$$

or

$$(\sum_{j} r_{j}[F_{j}]) + [P] + [Q] - [D_{\infty}] = 0$$

for $P, Q \in \mathcal{C}(\mathbb{F}_q)$ and $r_j \in \mathbb{N}_0$,

- the average distance of an element from the root of the tree is in $\tilde{O}(1)$,

an instance of the DLP in $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ can with a randomized algorithm be solved in an expected time of $\tilde{O}(f(q)^2)$.

The *proof* of this result is a very easy adaption of arguments in [10, Section 4]. We note that for the case that the order of the class group is not square-free, the arguments rely on a key lemma ([10, Lemma 4]) which can already be found in [31].

Because of this result, we concentrate on the question what size \mathcal{F} should have in order that the tree of large prime variation has $\frac{1}{2} \cdot q$ elements, say. For this, we first discuss a minor difficulty which is particular to our algorithm, the problem of "repeated selections".

Repeated selections

Most completely split lines passing through points of the factor base and exactly two large primes can be defined by specifying (d-2)(d-3)/2 pairs of points of the factor base. Say such a line occurs after the k_1^{th} and the k_2^{th} iteration of Step 3 (b). Then it might happen in the k_1^{th} iteration, the corresponding relation is not included in the tree, but – because the tree has grown in the meantime – in the k_2^{th} iteration it is.

In the following analysis of the evolvement of the tree of large prime variation, we *disregard* repeated selections of the same line, that is we analyze the algorithm as if every line was only selected *once*. This means that if our analysis is correct, the final number of elements in the tree of large prime variation is *larger or equal* than the number we state. (Concretely, we will determine the size of the factor base so that on heuristic grounds we expect that the number of elements in the tree of large prime variation is at least $\frac{1}{2}q$.)

Method and assumptions

Similarly to the analysis in [15], we use differential equations to study the evolvement of the tree of large prime variation.

Our heuristic analysis relies on several assumptions: First, similarly to the previous section, we estimate the number of completely split lines passing through one or two large primes and otherwise elements of the factor base. Second, we assume that the large primes occur "equidistributed" such that the evolvement of the tree of large prime variation can be modeled by equations of random variables. Third, we switch from discrete to continuous time (which involves mild heuristic assumptions). Under the first two assumptions, the evolvement of the tree of large prime variation is then modeled by a stochastic differential equation. Assuming that the variances are small (which is the forth assumption), we can study the evolvement by corresponding usual differential equations.

We note that the last three heuristic assumptions we just stated are also present in [15] (sometimes in implicit form). However, our assumption on equidistributiveness on the occurrence of large primes is stronger than the corresponding assumption in [15]. Indeed, if the algorithm in [15] is applied to a hyperelliptic curve \mathcal{C}/\mathbb{F}_q , the factor base is invariant under the hyperelliptic involution and contains all Weierstraß points and $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q) = \langle a \rangle$, the large primes are selected uniformly and independently.

Some probabilities and numbers and an assumption

Again the probability that a line is completely split is heuristically approximately 1/d!. The probability that a completely split line passes

through points of the factor base and one large prime is heuristically approximately

$$d \cdot \left(\frac{\#\mathcal{F}}{q}\right)^{d-1} \approx d \cdot q^{(r-1) \cdot (d-1)} .$$
(9)

This suggests that the number of relations involving one large prime one can obtain with lines passing through two points of the factor base is approximately

$$B := \frac{1}{(d-1)!} \cdot q^{r \cdot (d-1) - (d-3)} .$$
(10)

The probability that a completely split line passes through points of the factor base and two large primes is heuristically approximately

$$\frac{d\cdot(d-1)}{2}\cdot\left(\frac{\#\mathcal{F}}{q}\right)^{d-2}\approx\frac{d\cdot(d-1)}{2}\cdot q^{(1-r)\cdot(d-2)},\qquad(11)$$

and this suggests that the number of relations with two large primes one can obtain is approximately

$$C := \frac{1}{2 \cdot (d-2)!} \cdot q^{r \cdot (d-2) - (d-4)} .$$
(12)

As stated in "Method and assumption", it is not obvious that large primes hit with the lines we construct occur "equidistributed". This is the essence of the following heuristic assumption.

Heuristic Assumption 18 The final number of elements occurring in the tree of large prime variation constructed is typically close to the expect number of primes in the tree of large prime variation one obtains by applying the same procedure and the same number of iterations to the corresponding formalized model in Section 4.1 of [15].

Constructing the tree of large prime variation

Step 3 (a)

As for our choice of parameters the number of relations with one large prime is always negligible with respect to q, we can expect that after Step 3 (a) the tree of large prime variation contains as many elements as there are relations with one large prime (see (10)). (Under our assumptions, one can also use the approach via differential equations as in Section 4.1. to derive this result.)

Step 3 (b)

We use the approach via differential equations from [15]. As in [15], let u(t) be the portion of large primes occurring in the tree at time t (where one "unit of time" corresponds to the section of one line passing through at most 2 large primes).

As a line passing through two large primes leads to a new vertex of the tree if and only if exactly one of the two large primes is in the tree, we conclude that we can describe the process *roughly* via the differential equation

$$d(q u) = 2 u(t) \cdot (1 - u(t)) \cdot dt$$
(13)

with starting value u(0) = B/q given by (10). The equation is equivalent to

$$dt = \frac{q}{2} \frac{du}{u(1-u)} \,. \tag{14}$$

For our purposes it is convenient to express the time in terms of u. If we do so, the general solution to (13) is

$$t = \frac{q}{2}\log(\frac{u}{1-u}) + D$$
 (15)

with some constant D. Setting t = 0 gives in our case $D = -\frac{q}{2}\log((\frac{B/q}{1-B/q}))$ which is essentially $-\frac{q}{2}\log(\frac{B}{q}) = \frac{q}{2}\log(\frac{q}{B})$. We obtain

$$t = \frac{q}{2} \log\left(\frac{q}{B} \cdot \frac{u}{1-u}\right). \tag{16}$$

We want to determine T such that

$$u(T) = \frac{1}{2} \; .$$

This leads to

$$T = \frac{q}{2} \cdot \log\left(\frac{q}{B}\right) \,. \tag{17}$$

We simplify this to

$$T \le \frac{q}{2} \cdot \log(q) \ . \tag{18}$$

Remark 19 This simplification means that the size of the tree of large prime variation after Step 3 (a) (i.e. after having merely considered the relations with one large prime) is irrelevant for our analysis.

We now have to guarantee that the factor base has enough elements such that we can generate T relations with large primes. The condition is that

$$C \ge T$$
 .

Using (18) (and the definition of C in (12)), we see that this is satisfied if

$$q^{r(d-2)-(d-3)} \ge (d-2)! \cdot \log(q)$$
 (19)

Taking logarithms, we see that this is satisfied if

$$r(d-2) - (d-3) \ge \frac{\log(d-2!) + \log\log(q)}{\log(q)} .$$
⁽²⁰⁾

The minimal choice for r is

$$r = 1 - \frac{1}{d-2} \cdot \left(1 - \frac{\log(d-2!) + \log\log(q)}{\log(q)} \right).$$
(21)

(Provided that this is ≤ 1 , which is the case asymptotically.)

Let us assume for a moment that the average distance of an element in the tree from the root grows polynomially in $\log(q)$ (this is justified below). Then with the choice of r as in (21), we obtain a total running time of $\tilde{O}(r^{2-\frac{2}{d-2}})$. This leads to the following heuristic result (which is based on the assumptions summarized in "Method and assumptions", in particular Heuristic Assumption 18).

Heuristic Result 20 With the randomized algorithm presented above, one can asymptotically for varying q solve essentially all instances of the DLP in degree 0 class groups of curves represented by plane models of a fixed degree $d \ge 4$ over \mathbb{F}_q in an expected time of $\tilde{O}(q^{2-\frac{2}{d-2}})$.

The average distance

We now justify that we can expect that the average distance of an element in the tree from the root grows polynomially in $\log(q)$. The following (partially heuristic) argument applies to the tree which is actually constructed in Step 3 (b) of the algorithm as well as to the smaller tree which would be constructed without "repeated selections".

The corresponding question in [15] is studied in [15, Section 4.4]. The argument (in our situation as well as in the situation of [15]) can however be vastly simplified.

We study the average distance of the elements of the tree of large prime variation as a function of u, the number of elements in the tree. (This is more natural than to study it as a function of t because it is irrelevant how many tries one needs to introduce one element into the tree; it is just important what happens if one succeeds.)

Let $\omega = \omega(u)$ be the average distance from the root of the elements of the tree of large prime variation of size u. (As in [15], we assume that the variances are small.) The basic observation is: If an element is introduced into the tree, it has in average the distance of $\omega + 1$ from the root.

We note that $q \cdot u \cdot \omega$ is the sum of the distances to the root of all elements in the tree, and we recall that $q \cdot u$ is the number of elements in the tree. This leads to the differential equation.

$$d(q u \omega) = (\omega + 1) \cdot d(q u), \quad \text{i.e.}$$
(22)

$$\omega \, du + u \, d\omega = (\omega + 1) \cdot du, \quad \text{i.e.} \tag{23}$$

$$u\,d\omega = du.\tag{24}$$

We obtain

$$\omega = \log(u) + E \tag{25}$$

for some constant E. As for t = 0 we have $\omega = 1$ and u = B/q, we obtain

$$\omega = \log(u) + 1 - \log(\frac{B}{q}) = \log(u) + 1 + \log(\frac{q}{B}), \quad \text{i.e.}$$
 (26)

$$\omega \le \log(\frac{q}{B}) \le \log(q) \tag{27}$$

independent of the size of the tree. We remark that again the number B is irrelevant for our analysis.

This justifies the claim.

Experimental results

We implemented the algorithm of the previous section for non-hyperelliptic genus 3 curves, represented as plane curves of degree 4 with a double large prime variation as the "simplified algorithm" in [15]. (That is, the generation of the tree of does not follow the two-step procedure described here, and the relation matrix is built during the tree is constructed.)

With our Magma implementation (run on a PC clocked with 1.3 GHz) applied to curves over $\mathbb{F}_{2^{25}}$ we can construct roughly 2000 lines per second through the factor base and test them for smoothness.

By comparison, with the C/C++ implementation for arithmetic in class groups of hyperelliptic curves used for the experimental results in [15] one can perform roughly 200 000 additions in the class group per second for curves over $\mathbb{F}_{2^{25}}$ run on a PC clocked with 1.7 GHz (as reported in [15]). This means that over $\mathbb{F}_{2^{25}}$ the implementation for [15] cannot perform more than 100 times more tries for the relation search than our algorithm.

It was unclear to us when exactly the relation matrix is large enough such that the second row (for βb) linearly depends on the other rows. We terminated the relation generation when the relation matrix had twice as many rows as columns, and this did always suffice.

An example and a problem

The genus 3 curve with the largest class group to which we applied the algorithm was the curve C given by $Y^3 + Y + X^4 + X^3 + X^2 + X + 1 = 0$ over $\mathbb{F}_{2^{19}}$. The order of $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_{2^{19}})$ is $2 \cdot \ell$ for a prime number $\ell \approx 2^{56} = \frac{1}{2} \cdot 2^{19 \cdot 3}$.

We included $2442 \approx 3.3 \cdot 2^{19/2}$ elements in the factor base. We needed to consider 1 188 769 tuples of elements of the factor base giving rise to (not necessarily different) lines passing through at least two points of the factor base and otherwise large primes to construct a sparse matrix with doubly as many rows as columns, and the computation needed roughly 2 hours of CPU time. (The total number of tuples giving rise to such lines was 1 494 568 which is quite close to the heuristic prediction of $2442^2/4 = 1490\,841$.)

We tried to use the implementation of Lanczos' algorithm in Magma (ModularSolution(Transpose(R), ℓ : Lanczos := true)), but the algorithm crashed (twice) with an internal error. Doing calculations with other matrices, we noticed that the algorithm consumes much more storage than one might expect, and due to storage problems or an internal error, we not able to calculations with (extremely sparse) matrixes R with more than 5000 columns and rows.

To circumvent this problem, we used the structured Gaußian elimination algorithm in ModularSolution. The algorithm terminated after roughly 3 hours of CPU time using about 1 GByte of RAM. The calculated solution to the DLP was accurate.

It is reported in [15] that for $q \approx 2^{19}$, 76 709 007 completely split divisors had to be considered. This number is *roughly 64 times larger* than the number of completely split lines we had to consider.

Implicit constants and practical aspects

It is not necessary to follow the 2-step procedure of the relation generation, and one should construct the full graph of large prime variation and not just the tree of large prime variation.

Depending on g, g and the time one needs to perform the arithmetic in $\operatorname{Cl}^0(\mathcal{C}/\mathbb{F}_q)$, it might be advisable to combine the algorithm of the previous section with a double large prime variation as described in [15] (as we did in our experiments).

Concerning constants and further tricks, mutatis mutandis the remarks in the corresponding part of the previous section also apply here.

5 Finding plane models of small degree

In this section, we discuss methods to find a plane model of small degree of an explicitly given curve. We note that according to our general philosophy stated in the introduction, we assume that the curve given is also represented via a plane model, but that initial model might have a higher degree than the final model. All complexity estimates are given for curves represented *initially* by plane models of *bounded degree*.

We assume that the reader has some familiarity with divisors, Riemann-Roch spaces and linear systems.

Definitions and Basic Facts

Definition 21 Let D be a divisor on a curve C (as always a curve is assumed to be proper, smooth and geometrically irreducible). Then the Riemann-Roch space of D is denoted by L(D). The complete linear system defined by D is denoted by |D|. The *dimension* of D is the (projective) dimension of the associated complete linear system |D|, i.e.

$$\dim(D) = \dim(|D|) = \dim(L(D)) - 1.$$

Let us recall some basic facts we need.

Let \mathcal{C} be a curve a some field K, let D be a divisor on \mathcal{C} , let $\varphi_1, \ldots, \varphi_{n+1} \in L(D)$, not all zero, and let $V := \langle \varphi_1, \ldots, \varphi_{n+1} \rangle_K$.

Then we have a rational map

$$\varphi: \mathcal{C} \longrightarrow \mathbb{P}^n$$

given by

$$P \mapsto (\varphi_1(P) : \cdots : \varphi_{n+1}(P))$$

(where P is not a pole of any of the φ_i and P is not a common zero of the φ_i). The rational map φ extends uniquely to a morphism $\varphi : \mathcal{C} \longrightarrow \mathbb{P}^n$. This morphism satisfies

$$\deg(\varphi) \cdot \deg(\mathcal{C}) \le \deg(D)$$

If furthermore the linear system determined by (D, V) is base point free (i.e. if for every $P \in \mathcal{C}(\overline{K})$, there exists some $\psi \in V \otimes_{\mathbb{F}_q} \overline{\mathbb{F}}_q \subset \overline{\mathbb{F}}_q(\mathcal{C})$ such that P is not contained in the support of $\operatorname{div}(\psi) + D$), this inequality is satisfied with equality.

Here deg($\varphi(\mathcal{C})$), the *degree* of $\varphi(\mathcal{C})$ in \mathbb{P}^n , is degree of the intersection cycle of $\varphi(\mathcal{C})$ with any hyperplane which does not contain $\varphi(\mathcal{C})$.

We also note a fact we have already used implicitly: Any plane model is defined by a single homogeneous equation. This can for example be seen as follows: If $\mathcal{C}_{pm} \subset \mathbb{P}^2$ is a plane model, we can view \mathcal{C}_{pm} as a divisor in \mathbb{P}^2 , and \mathcal{C}_{pm} is the divisor of zeros of the global section $1 \in \Gamma(\mathbb{P}^2, \mathcal{O}(\mathcal{C}_{pm}))$. We have $\mathcal{O}(\mathcal{C}_{pm}) \approx \mathcal{O}(d)$ for some $d \geq 1$. If s is the global section of $\mathcal{O}(d)$ corresponding to $1 \in \Gamma(\mathbb{P}^2, \mathcal{O}(\mathcal{C}_{pm})), \mathcal{C}_{pm}$ is defined by s. The number d is then the on the one hand the degree of \mathcal{C}_{pm} and on the other hand the degree of the homogeneous equation s.

We also recall the Riemann-Roch Theorem (see [17, Theorem 1.3]):

Proposition 22 (Riemann-Roch) Let C be a curve of genus g with canonical divisor K, and let D be a divisor on C. Then

$$\dim(L(D)) - \dim(L(K - D)) = \deg(D) + 1 - g,$$

in particular

$$g - \deg(D) + \dim(D) \ge 0.$$

Basic Algorithms

All computations are concerned with curves over finite fields. We need algorithms for some basic tasks. We need:

- 1. An algorithm to compute bases of Riemann-Roch spaces of divisors.
- 2. An algorithm to compute canonical divisors.
- 3. An algorithm which given a curve \mathcal{C}/\mathbb{F}_q and $\varphi_1, \varphi_2, \varphi_3 \in \mathbb{F}_q(\mathcal{C})$, not all zero, computes an equation of $\varphi(\mathcal{C})$ in \mathbb{P}^2 .

The first task was already present several times in this work; we again refer to [20]. Given an effective divisor D, the number of field operations required (including calculating the required orders) is polynomial in deg(D) · log(q) (as always for curves given by plane models of *bounded degree*). More generally given any divisor, the number of field operations is polynomial in height $(D) \cdot \log(q)$. This means that for divisors of a bounded height, the running time is polynomial in $\log(q)$.

The second task is for example discussed in Section 9 of [20]. Again the algorithm terminates in polynomial time in $\log(q)$.

The third task can for example be performed via Gröbner base algorithms. Or one can apply the following easy method which works very efficiently in practice if q is not "too small":

Let us assume that $\varphi(\mathcal{C})$ has degree d, and let $\sum_{i,j \leq d} \alpha_{i,j} X^i Y^j Z^{d-i-j}$ be the (unknown) homogeneous equation defining $\varphi(\mathcal{C})$. Then for any point $P \in \varphi(\mathcal{C}(\mathbb{F}_q))$, we have

$$\sum_{i,j \le d} \alpha_{i,j} \cdot \varphi_1(P)^i \, \varphi_2(P)^j \, \varphi_3(P)^{d-i-j} = 0.$$

This gives a linear equation on the $\alpha_{i,j}$. If in this way one has generated $\geq \binom{d+2}{2} - 1$ linear independent equations, one obtains $(\alpha_{i,j})_{i,j}$ as the unique solution (up to multiples) of the system. We did experiments, and in these experiments it always sufficed to generate twice as many equations than unknowns.

We conclude that heuristically, given $\varphi_1, \varphi_2, \varphi_3$, represented as rational functions involving a bounded number of terms such that $\varphi(\mathcal{C})$ has degree d, own can with this method one can find an equation of $\varphi(\mathcal{C})$ in a bounded number of field operations.

By combining 1. and 3., we obtain:

Heuristic Result 23 Given a curve C/\mathbb{F}_q (represented by a plane model of bounded degree) and a divisor D on C of dimension 2 and bounded height, one can calculate an equation of the image of C in \mathbb{P}^2 given by a basis $\varphi_1, \varphi_2, \varphi_3$ of L(D) in a bounded number of field operations.

A first approach

Now let \mathcal{C}/\mathbb{F}_q be an explicitly given curve of genus g.

As a first idea to find a model of small degree, one might proceed as follows: One chooses a "random" effective divisor D of degree g + 2 on C. One might expect that if \mathbb{F}_q is "large enough", such a divisor will be non-special (i.e. $\dim(L(K - D)) = 0$) with overwhelming probability. The dimension of |D| will then be 2. Any basis in L(D) defines (as laid out above) a morphism to \mathbb{P}^2 . As D is "random" one further might expect that |D| does not have base points and that the morphisms defined by it are birational unto their image (which then is a plane model of the curve). Below, we will argue that one should expect that this method works as described for "sufficiently general" curves. Let us first fix some (classical and quite general) terminology.

Terminology

By general curve of genus $g \ge 2$ we mean a curve obtained by base-change from the curve corresponding to the generic point of the (coarse) moduli space \mathcal{M}_g of curves of genus g over $\operatorname{Spec}(\mathbb{Z})$. (This space exists by [27, Corollary 7.14.].)

As in the previous parts of this work, we only work with curves over fields and not over more general base schemes. Let \mathcal{C} be a curve over a field K. We note that for all field extensions L|K, the effective divisors of degree d on \mathcal{C}_L are classified by the L-valued points of $\mathcal{C}^{(d)}$, the d-fold symmetric power of \mathcal{C} . The effective divisors of dimension $\geq n$ thereby correspond to points of a closed subscheme which we call the *locus of effective divisors of degree d and dimension* $\geq n$ *in the symmetric power*, denoted by $\mathcal{C}_n^{(d)}$ (this notation is slightly different from the one in [16]).

Analogously, for all field extensions L|K, the complete linear systems of degree d on C_L are classified by the *L*-valued points of the d^{th} component of the Picard scheme (of all degrees) of C. This component is a twist of the

Jacobian of \mathcal{C} (i.e. it becomes isomorphic to the Jacobian of \mathcal{C} after a base change $\overline{\mathbb{F}}_q|\mathbb{F}_q$). The complete linear systems of dimension $\geq n$ again correspond to points of a closed subscheme which we call the *locus of complete linear systems d and dimension* $\geq n$ on a twist of the Jacobian, denoted by W_n^d (this notation is again not in accordance with [16]).

Let \mathcal{C}/K be a general curve, let $K(\mathcal{C}^{(d)})$ be the function field of $\mathcal{C}^{(d)}$. Then the effective divisor corresponding to the generic point of $\mathcal{C}^{(d)}$ on $\mathcal{C}_{K(\mathcal{C}^{(d)})}$, the curve obtained by base change $K(\mathcal{C}_{K(\mathcal{C}^{(d)})})|K$ from \mathcal{C} , is called a *general* effective divisor on a general curve.

We also talk about general effective divisors of a fixed degree d and dimension $\geq n$ on a general curve. This can also be made precise by considering generic points of the appropriate loci in $\mathcal{C}^{(d)}$.

Below we talk about certain properties of general curves and general effective divisors of a fixed degree d (and dimension $\geq n$) on general curves. These properties in fact hold in an open part of an appropriate moduli space. Because of this, it makes sense to speak of *sufficiently general* curves or about *sufficiently general* effective divisors of degree d (and dimension $\geq n$) on *sufficiently general* curves.

Complete linear systems on general curves

We are interested in the dimensions of (sufficiently) general effective divisors of degree d and dimension $\geq n$ on (sufficiently) general curves and the properties of the corresponding morphisms to \mathbb{P}^n . Work on these questions was already pioneered by Brill and Noether and continued among others by Kleiman and Laskov ([24]) and Griffiths and Harris ([16]). Part a) of the following proposition is proven in [24] and Part b) in [16]. (The proposition itself is an excerpt of the "Main Theorem" in [16].)

Proposition 24 Let $g, d, n \in \mathbb{N}$ and assume that $g - d + n \ge 0$. Let

$$\rho = \rho(g, d, n) := g - (n+1) \cdot (g - d + n).$$

Then

- a) for any curve of genus g, $C_n^{(d)}$ has dimension $\geq \rho + n$ and W_n^d has dimension ρ .
- b) For a general curve of genus g, $C_n^{(d)}$ is equidimensional of dimension $\rho + n$, and W_n^d is equidimensional of dimension ρ .

Part b) of this proposition implies (see Point b) in the introduction of [16]):

Proposition 25 Let $g, d, n \in \mathbb{N}$ and assume that $g - d + n \ge 0$. Let ρ be as above.

Let C be a general curve. Let |D| be a general complete linear system of degree d and dimension $\geq n$ on C. Then

- |D| has dimension exactly n,
- |D| has no base points,
- if $n \geq 2$, a corresponding morphism $\mathcal{C} \longrightarrow \mathbb{P}^n$ (i.e. a morphism given by a basis of L(D) as described above) is birational onto its image.

The statements also hold for sufficiently general effective divisors of degree d and dimension $\geq n$.

All these statements follow from simple dimension-arguments. For example, the first statement follows from $\rho(g, d, n') > \rho(g, d, n)$ for n' > n. We note that it is also true that for $\rho \ge 0$ a general effective divisor of degree d and dimension $\ge n$ on C has dimension exactly n, but this statement is weaker than the one in the proposition (it corresponds to $\rho(g, d, n') + n' > \rho(g, d, n) + n$ for n' > n).

Heuristically, Proposition 25 has the following interpretation:

Heuristic Interpretation 26 Let us fix g, d, n with $g - d + n \ge 0$. For a prime power q, let P(q) be the portion of isomorphism classes of tuples $(\mathcal{C}, |D|)$, where \mathcal{C}/\mathbb{F}_q is a curve of genus g and |D| is a complete linear system of degree d and dimension $\ge n$ on \mathcal{C} , for which the statements of Proposition 25 are satisfied. (If for some q no such tuples exist, we set P(q) := 1.) Then P(q) is in $\Omega(1 - \frac{1}{q})$ for $q \longrightarrow \infty$.

Back to the "first approach"

The statement in the "first approach" that one should expect that an effective divisor of degree g + 2 on a curve of genus g has no base points and defines a morphisms to \mathbb{P}^2 which are birational unto their image can now be made precise: Any sufficiently general divisor of degree g+2 on a sufficiently general curve of genus g has the properties.

Using special divisors

For sufficiently general curves of genus ≥ 3 , there is an easy method to find a plane model of degree g+1 instead of g+2: First one fixes a canonical divisor K. Then one chooses an effective divisor D of degree g-3 and considers the divisor K-D. By Riemann-Roch this divisor will have dimension ≥ 2 . As is shown in the following proposition, for any sufficiently general divisor D on a sufficiently general curve C, K-D has dimension 2, no base points and defines morphisms $\mathcal{C} \longrightarrow \mathbb{P}^2$ which are birational unto their image. **Proposition 27** Let C be any curve of genus $g \ge 3$. Then we have a birational morphism

In particular, for any (sufficiently) general divisor D of degree g-3 on a (sufficiently) general curve C, |K - D| has no base points and defines a morphism to \mathbb{P}^2 which is birational onto its image.

Proof. The morphism $\mathcal{C}_0^{(g-3)} = \mathcal{C}^{(g-3)} \longrightarrow W_0^{g-3}$, $D \mapsto |D|$ is birational because on any curve, any general divisor of degree $g - \delta$ for some $\delta > 0$ has dimension 1.

By the Riemann-Roch theorem and the fact that $\deg(K) = 2g - 2$, we have an isomorphism $W_0^{g-3} \longrightarrow W_2^{g+1}$, $|D| \mapsto |K - D|$. \Box

We note that whereas with this method one can find plane models of degree g + 1 of sufficiently general curves of genus $g \ge 3$, the method does *not* work for hyperelliptic curves. (Every special divisor of dimension ≥ 1 on a hyperelliptic curve defines morphism to some projective space whose image is birational to \mathbb{P}^1 .)

Non-hyperelliptic curves of genus 3

Non-hyperelliptic curves of genus 3 have the nice property that the canonical linear system itself has dimension 2. One obtains in this way morphisms into \mathbb{P}^2 which are in fact embeddings, the so-called *canonical embeddings*. The image is a non-singular curve of degree 4, a so-called *canonical curve* (see [17, IV,5]).

A problem

One might ask if it there exist divisors on sufficiently general curves of dimension 2 of smaller degree than g+1. One might also ask for an algorithm to determine if divisors of a prescribed degree and dimension exist on some curve, and if that is the case to find them.

Here are some thoughts which might be useful to answer the first question. Propositions 24 and 25 imply:

Proposition 28 Let $g, d, n \in \mathbb{N}$ such that $g - d + n \ge 0$ and the number ρ defined in Proposition 24 is ≥ 0 . Then on any sufficiently general curve over an algebraically closed field, there exists a divisor of degree d and dimension n.

In particular, if C is a sufficiently general curve over an algebraically closed field, there exists a morphism $\mathcal{C} \longrightarrow \mathbb{P}^n$ which is birational towards its image such that the image has degree d.

As a special case we obtain:

Proposition 29 Let $\frac{2}{3}g+2 \le d \le g+2$. Then any sufficiently general curve of genus g over an algebraically closed field has a plane model of degree d.

We are unaware of a similar result over finite fields. We note that it would be interesting to know if for the general curve C over the generic point of \mathcal{M}_g over $\operatorname{Spec}(\mathbb{Z})$, the locus of divisors of dimension 2 and degree d in $C^{(d)}$ has an irreducible component which is geometrically irreducible. Heuristically, this would mean that for a fixed genus and degree and for large q, an overwhelming portion of curves over \mathbb{F}_q would have a divisor of degree d and dimension 2.

We also asked for an algorithm to determine whether divisors of a prescribed degree and dimension exist on some curve, and if that is the case, to find such a divisor.

As was noticed by Brill and Noether, the loci $C_n^{(d)}$ in $C^{(d)}$ are (essentially) defined by subdeterminants of certain matrices, the so-called Brill-Noether matrices (cf. [24, Remark 6]).

This original method to describe $C_n^{(d)}$ can rather easily be turned into an algorithm. We plan to present the algorithm and experimental results obtained with it in a future work.

Experimental results

We implemented

- 1. the method to find models of degree g + 2 (the "first approach"),
- 2. the method to find models of degree g + 1 ("using special divisors").

The experiments were conducted with curves of genus 3 to 8. The derivation of the defining polynomial was carried out as described in Point 3 of "Basic Algorithms" in the previous section. In our experiments we were always able to derive the equation of $\varphi(\mathcal{C})$ after having chosen $2 \cdot \binom{d+2}{2}$ different points of $\mathcal{C}(\mathbb{F}_q)$.

In the first method, the divisors were selected with the RandomPlace function. The results were as follows:

- 1. We applied the first method to a large variety of curves of genus ≥ 3 (including hyperelliptic curves) and in *every* case, we obtained a plane model of the curve of degree g + 2.
- 2. We tried the second method for various classes of non-hyperelliptic curves over finite fields of size 1009 and larger (all in all trying several thousands of curves), and in *all cases*, we obtained a plane model of

degree g + 1. In particular, we tried the method for the explicit curve of genus 7 at the end of [7] and 100 other genus 7 curves generated with the particular instances of the GHS attack in odd characteristic described in [7, 6.2, 7.3], and we obtained plane models of degree 8.

We view the methods and heuristics presented in this section as being experimentally confirmed.

6 Discussion and outlook

On the security of systems based on non-hyperelliptic curves of genus 3

Given our heuristic result and our experimental evidence, one might conclude that non-hyperelliptic (canonical) genus 3 curves are cryptographically weak. Indeed, by considering the respective heuristic running times, one would be tempted to conclude that a non-hyperelliptic curve of genus 3 over a field \mathbb{F}_q does not provide a better security level than a genus 2 curve over the same field (but in the former case the key length is 1.5 times as large).

This conclusion would however be to hasty. As pointed out by Amirazizi and Hellman ([2]), Bernstein ([5]), Wiener ([35]) and other cryptographers, a better cost estimate than the mere running time is the *full cost*. According to Wiener ([35]), "the full cost of an algorithm run on a collection of hardware is number of components multiplied by the duration of their use".

Let us fix the following notation: If $f, g : \mathbb{N} \longrightarrow \mathbb{R}$ are two functions, then $f = \tilde{\Theta}(g)$ means that $f(n) \leq p(\log(n)) \cdot g(n)$ and $g(n) \leq q(\log(n)) \cdot f(n)$ for all $n \gg 0$ with two polynomial functions p, q.

Whereas our double large prime variation algorithm has a heuristic running time of $\tilde{\Theta}(q)$, due to the memory requirements of $\Theta(q)$, the full cost is $\tilde{\Theta}(q^2)$. This is clearly worse than the full cost $\tilde{O}(q^{3/2})$ of Pollard's rho method which only requires very little memory. It is however intuitively obvious that our (unparallelized) algorithm cannot be optimal for full cost – one might describe the situation by saying "A huge bulk of memory just sits around and does nothing". It is very intuitive that some kind of parallelization should lead to a reduction in the full cost.

As a first step towards a parallelized variant with a better full cost, we note that the linear algebra part of the algorithm (the solution of an inhomogeneous sparse linear system of size $q^{1/2}$) can (heuristically) via a 2dimensional mesh of size $q^{1/4} \times q^{1/4}$ be performed with a full cost of $\tilde{\Theta}(q^{5/4})$ instead of $\tilde{\Theta}(q^{3/2})$. The heuristic running time is thereby $\tilde{\Theta}(q^{3/4})$ (see [5] and [26]).

Given the definition of full cost cited above, it seems however to be difficult to derive a parallelized variant of our algorithm whose full cost is below $\tilde{O}(q^{3/2})$. Indeed, just by the facts that our graph of large prime variation needs a storage of $\Theta(q)$ and the running time of the linear algebra part is $\tilde{\Theta}(q^{3/4})$, we obtain a running time of $\Omega(q^{7/4})$. (One could now decrease the number of large primes but that would course other problems.)

Given that the relation collection and the linear algebra part are already today often performed on different hardware, we now ask whether it is possible to perform these two tasks *separately* with a smaller full cost than $O(q^{3/2})$.

A suitable factor base can be constructed in an expected time of $\Theta(q^{1/2})$, leading to a full cost of just $\Theta(q)$, even without parallelization.

To analyze a parallelization the relation generation step, we use the model of [35]: P different processors access a common storage in 3-dimensional space. Using [35, Corollary 2] (and the heuristic analysis in Section 4), we obtain a heuristic full cost of

$$\tilde{O}(q^{\frac{4}{3}})$$

for an optimally parallelized construction of the tree of large prime variation (using $\Theta(q^{2/3})$ processors). Again with $\Theta(q^{2/3})$ processors, the Step 4 of the algorithm, the relation generation, has a negligible full cost of

$$\tilde{O}(q^{\frac{5}{6}}).$$

The sum of the costs for the relation search and the linear algebra is thus – according to our heuristic analysis – asymptotically bounded by the cost of $\tilde{O}(q^{4/3})$ of the relation search.

We conclude that in a weakened sense and on the basis of a heuristic analysis, the full cost of an attack on the DLP in degree 0 class groups of non-hyperelliptic curves with the variant of our algorithm described here is indeed lower than the full cost of Pollard's rho method (if the order of the degree 0 class group is "nearly prime").

Concluding remarks

We conclude with some general remarks on the feasibility of index calculus in class groups non-hyperelliptic curves in comparison with hyperelliptic curves.

From time to time we heard cryptographers or mathematicians express the opinion that a certain cryptographic system based on non-hyperelliptic curves should be more secure than a system based on hyperelliptic curves of the same genus or that certain kinds of covering attacks are "just theoretical" because the resulting curves are not hyperelliptic. We hope that with this work we have clarified this issue in two aspects:

First, it does not provide any difficulties to adapt the index calculus algorithm by Gaudry as well as the essential publicly known improvements of this algorithm (reduced factor base, (double) large prime variation) to more general curves. We stress that the asymptotic results stay the same if one assumes that the curves are represented by a suitable model of small degree. (The results are however only heuristic whereas in the hyperelliptic case and under some conditions on the class groups a rigorous analysis can be given for a variant of the Gaudry-Harley algorithm.) Moreover, also from a practical point of view, the algorithms are not much slower than the algorithms for hyperelliptic curves. Indeed, the arithmetic in hyperelliptic curves (in imaginary quadratic representation) can be performed in $O(q^2)$ field operations whereas the arithmetic in general curves (represented by a suitable model of small degree) can be performed in $O(g^4)$ field operations ([20, 18]).

Second, independent of the perspective one has in mind (be it theoretical or practical and be it running time or full cost), our index calculus attack on the DLP in class groups of *non-hyperelliptic* genus 3 curves is *more efficient* than the currently best publicly known index calculus algorithm for class groups of hyperelliptic genus 3 curves. Further, again on the basis of current publicly available knowledge, asymptotically, the DLP in degree 0 class groups of "sufficiently general" curves of other genera can on heuristic grounds also be solved *faster* than the DLP in degree 0 class groups of hyperelliptic curves of the same genus.

Future research topics

There are several natural research topics related to this work.

- Conduct more experiments, in particular try to calculate instances of the DLP in class groups of non-hyperelliptic genus 3 curves over fields of size $> 2^{27}$ (which is the largest example in [15]).
- Analyze a "large degree" variant of our algorithm (which essentially generalizes the algorithm by Adleman, DeMarrais and Huang to more general than hyperelliptic curves) with particular emphasis on the case that the genus is much larger than the degree. (Such an algorithm was formulated by Hess in [19].)
- Implement the algorithm to find the locus of divisors of degree d and dimension $\geq n$ on $\mathcal{C}^{(d)}$ for some curve \mathcal{C} , and try to obtain plane models of small degree with this algorithm.

We plan to address these questions in the near future.

Acknowledgments

It is a great pleasure to thank G. Frey, P. Gaudry, F. Hess and E. Viehweg for discussions and helpful comments.

References

- L. Adleman, J. DeMarrais, and M.-D. Huang. A Subexponential Algorithm for Discrete Logarithms over the Rational Subgroup of the Jacobians of Large Genus Hyperelliptic Curves over Finte Fields. In M.-D. Huang L. Adleman, editor, *Algebraic Number Theory – ANTS I*, volume 877 of *LNCS*, pages 28–40, Berlin, 1994. Springer-Verlag.
- H. Amirazizi and M. Hellman. Time-Memory-Processor Trade-Offs. IEEE Transactions on Information Theory, IT-34(3):505-512, 1988.
- [3] A. Basiri, A. Enge, J.-C. Faugère, and N. Gürel. Implementing the Arithmetic of C_{3,4}-Curves. In Algebraic Number Theory – ANTS VI, LNCS, pages 87–101, Berlin, 2004. Springer-Verlag.
- [4] A. Basiri, A. Enge, J.-C. Faugère, and N. Gürel. The arithmetic of Jacobian groups of superelliptic cubics. *Math. Comp.*, 249:389–410, 2005.
- [5] D. Bernstein. Circuits for integer factorization: a proposal. available under cr.yp.to/papers/nfscircuit.pdf, Nov. 2001.
- [6] D. Cantor. Computing in the Jacobian of a hyperelliptic curve. Math. Comp., 177(95-101), 1987.
- [7] C. Diem. The GHS Attack in odd Characteristic. J. Ramanujan Math. Soc., 18:1–32, 2003.
- [8] C. Diem and J. Scholten. Cover attacks. A report for the AREHCC project, available under http://www.arehcc.com/documents.htm 2003.
- W. Eberly and E. Kaltofen. On randomized Lanczos algorithms. In W. Küchlin, editor, *Proceedings ISSAC 1997*, pages 176–183. ACM Press, 1997.
- [10] A. Enge and P. Gaudry. A general framework for subexponential discrete logarithm algorithms. Acta. Arith., 102:83–103, 2002.
- [11] S. Flon and R. Oyono. Fast arithmetic on Jacobians of Picard curves. In Advances in Cryptology — PKC 2004, volume 2947 of LNCS, pages 55–68, Berlin, 2004. Springer-Verlag.
- [12] S. Flon, R. Oyono, and C. Ritzenthaler. Fast addition on nonhyperelliptic genus 3 curves. preprint, 2004.
- [13] S. Galbraith and A. Menezes. Algebraic curves and cryptography. forthcomming.

- [14] P. Gaudry. An algorithm for solving the discrete log problem on hyperelliptic curves. In Advances in Cryptology — EUROCRYPT 2000, LNCS 1807, pages 19–34, New York and Berlin, 2000. Springer-Verlag.
- [15] P. Gaudry, N. Thériault, and E. Thomé. A double large prime variation for small genus hyperelliptic index calculus. Cryptology ePrint Archive, Report 2004/153, http://eprint.iacr.org/2004/153, Feb. 15, 2005.
- [16] P. Griffiths and J. Harris. On the variety of special linear systems on a general algebraic curve. Duke Math. J., 47(1):233–272, 1980.
- [17] R. Hartshorne. Algebraic Geometry. Springer-Verlag, New York, 1977.
- [18] F. Heß. personal communication.
- [19] F. Heß. Zur Divisorklassenberechnung in globalen Funktionenkörpern. PhD thesis, Technische Universität Berlin, 1999.
- [20] F. Heß. Computing Riemann-Roch spaces in algebraic function fields and related topics. J. Symbolic Computation, 11, 2001.
- [21] F. Heß. Weil descent attacks. In G. Seroussi I. Blake and N. Smart, editors, Advances in Elliptic Curve Cryptography. Cambridge University Press, 2004.
- [22] M.-D. Huang and D. Ierardi. Efficient Algorithms for the Riemann-Roch Problem and for Addition in the Jacobian of a Curve. J. Symbolic Computation, 18:519–539, 1994.
- [23] K. Khuri-Makdisi. Linear algebra algorithms for divisors on an algebraic curve. Math. Comp., 73:333–357, 2004.
- [24] S. Kleiman and D. Laksov. Another proof of the existence of special divisors. Acta. Math., 132:163–176, 1974.
- [25] D. Knuth. The Art of Programming, Vol. 3 (Sorting and Searching). Addison-Wesley, 1973.
- [26] A. Lenstra, A. Shamir, J. Tomlinson, and E. Trommer. Analysis of bernstein's factorization circuit. In Advances in Cryptography — ASI-ACRYPT 2002, LNCS. Springer-Verlag.
- [27] D. Mumford. Geometric Invariant Theory. Springer-Verlag, Berlin, 1965.
- [28] K. Nagao. Improvement of Thériault Algorithm of Index Calculus of Jacobian of Hyperelliptic Curves of Small Genus. Cryptology ePrint Archive, Report 2004/161, http://eprint.iacr.org/2004/161, 2004.

- [29] J. Pila. Frobenius maps of abelian varieties and fining roots of unity in finite fields. *Math. Comp.*, 55:745–763, 1990.
- [30] J. Pila. Counting points on curves over families in polynomial time. unpublished, 1991.
- [31] C. Pomerance. Fast, rigorous factorization and discrete logarithm algorithms. In D. Johnson, T. Nishizeki, A. Nozaki, and H. Wolf, editors, *Discrete Algorithms and Complexity, Proceedings of the Japan US Joint Seminar, June 4-6, 1986, Kyoto, Japan, pages 119–143, 1987.*
- [32] R. Schoof. Elliptic curves over finite fields and the compution of square roots mod *p. Math. Comp.*, 44:483–494, 1985.
- [33] N. Thériault. Index calculus attack for hyperelliptic curves of small genus. In Advances in Cryptology — ASIACRYPT 2003, volume 2894 of LNCS, pages 75–92, Berlin, 2003. Springer-Verlag.
- [34] E. Volcheck. Computing in the Jacobian of a Plane Algebraic Curve. In M.-D. Huang L. Adleman, editor, *Algebraic Number Theory – ANTS I*, volume 877 of *LNCS*, pages 28–40, Berlin, 1994. Springer-Verlag.
- [35] M. Wiener. The Full Cost of Cryptanalytic Attacks. J. Cryptology, 17(2):105–124, 2004.

Universität Duisburg-Essen, Institut für Experimentelle Mathematik, Ellernstr. 29, 45326 Essen, Germany. email: diem@iem.uni-due.de