Computing discrete logarithms with pencils

Claus Diem and Sebastian Kochinke

Augst 25, 2017

1 Introduction

Let us consider the discrete logarithm problem for curves of a fixed genus g at least 3. In [7] it is shown that the problem can be solved in an expected time of

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

where q is the cardinality of the base field.

The algorithm given follows the index calculus or relation generation and linear algebra strategy and as usual for such algorithms for the discrete logarithm problem for curves of a fixed genus, the so-called factor base consists of a subset of the set of rational points of the curve. Moreover, a so-called double large prime relation is used, where the set of "large primes" (which are not at all large here) consists of the remaining rational points.

The algorithm is randomized and – also as usual – the phrase "expected time" refers to the internal randomization of the algorithm; no randomization over input instances is considered. The same holds for all further statements.

For non-hyperelliptic curves this can be improved. Indeed, as was shown in [22], that for such curves, the problem can be solved in an expected time of

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

The algorithm for this result relies on the construction of a birational plane model of degree $g+1$ of the curve in question. In the case that $g = 3$ the canonical curve is used, and in higher genus, one can think of the plane model as being obtained by consecutive central projections through points on the curve. Then an index calculus algorithm (also with double large prime variation) is aplied to the resulting plane model. The central idea is here to generate relations by intersecting the plane model with lines. More precisely, the plane model is intersected with lines which already run through two points of the factor base. Note here that this means that one has to impose a condition on a "remaining" divisor of degree $g-1$ to generate relations. In contrast, in the original method, uniformly randomly given divisors where

considered, which are therefore with a high probability of degree g. Note here that the degrees of the divisors are directly reflected in the running times. This idea was first presented in [6], then a variant of the algorithm with a rigorous analysis was given in [9].

Let us note that one can interpret the approach via the birational plane model of degree $g+1$ as follows: The divisors defined by the lines form a (complete) linear system of (projective) dimension 2 and degree $g+1$. This linear system has index of speciality 1. For every fixed non-singular rational point of the plane model, the lines through the point define a *pencil*, that is, a 1-dimensional linear system of divisors. The center point is then what is called a base-point, and if one subtracts this base point from every divisor, one obtains a base-point free pencil of degree g. Again this system is complete and has index of speciality 1.

It is then natural to ask if one can use a family of pencils of higher index of speciality to attack the discrete logarithm problem. Let us note in passing that base-point free pencils correspond to equivalence classes of functions to $\mathbf{P}^1$ up to change of coordinates on the image, and that the use of functions to speed up the computation of discrete logarithms is classical. For example, for a hyperelliptic curve one uses the function of degree 2 to speed up the computation. The essential difference to our approch is however that this one function speeds up the computation by a constant factor whereas we want to obtain a better exponent in the expected running time. For this, we need much more than just one pencil.

The approach via the plane model of degree $g+1$ already gives an idea how to obtain such pencils: If one uses a singular rather than a non-singular rational point on the plane model, one obtains a base-point free pencil of degree at most $g-1$. (If, as we will show is typically the case, the singularity has order 2, the degree is exactly $g-1$.)

Now, for every curve of genus at least 4, every plane model of degree $g+1$ has a singularity. It turns out, however, that for curves of genus 4 the approach fails because every such curve has at most two pencils of degree 3. But for curves of genus at least 5, results on special linear systems on generic curves, the so-called Brill-Noether theory, suggests that the approach might be succesful.

With a suitable algorithm and an analysis of the underlying geometry we argue for the following conjecture:

Conjecture (short formulation). *For nearly all non-hyperelliptic curves of a fixed genus g at least 5, the discrete logarithm problem can be solved in an expected time of*

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

The formulation here contains the not yet defined phrase "for nearly all". A definition of this phrase and a completely precise formulation is given below under "Details". We note that we have deliberately written "Conjecture" and not "Heuristic Result" because we want to make a solid statement which exceeds "heuristic" statements for which it is not clear what the level of commitment to the exact claim made is.

Details

We now give a more detailed description of the statements, thereby fixing some terminology.

We follow the scheme-theoretic approach. We note that this means in particular that a morphism of varieties over a field k corresponds to what is classically called a morphism "defined over k".

In this work, a *curve* over some field is by definition always geometrically irreducible and geometrically reduced (the latter condition being automatic for curves over finite fields) but it need not be smooth. However, when we speak of a curve of a particular genus, we always demand that the curve be smooth without saying so explicitly. A *birational plane model* or simply a *plane model* of a smooth curve over a field k is a plane curve, that is, a curve in $\mathbf{P}^2_k$ which is birational to the given curve.

As said, we consider the discrete logarithm problem for curves of a fixed genus over finite fields. This means that upon input of a curve $\mathcal{C}$ of genus g over a finite field and two points $a, b \in \mathrm{Cl}^0(\mathcal{C})$, the degree 0 class group of $\mathcal{C}$ (as always over the field) with $b \in \langle a \rangle$, a natural number e with $e \cdot a = b$ shall be computed.

To represent the input instances, we follow the description in Section 2 of [7]: We represent the curve by a birational plane model of bounded degree, we represent the degree 0 divisor classes by divisors which are reduced along a fixed rational point P_0 (which always exists for q large enough), and we represent divisors in an ideal theoretic way; for details we refer to [7]. Let us note here that in the algorithm, the ideal theoretic representation is used to compute bases of the spaces of global sections (or Riemann-Roch spaces, L-spaces) $L(D) = \Gamma(\mathcal{C}, \mathcal{O}(D))$ for divisors D on the curve via Heß' algorithm ([19]).

As stated, the heuristally claimed result has not yet been formulated in a precise way, so let us do this now:

Considering curves of a fixed genus over finite fields, we define the phrase for "nearly all curves" as follows: Let us first fix a finite field $\mathbb{F}_q$. Then there are finitely many isomorphism classes of curves of the given genus over the field. We now consider a property P on curves over finite fields which is well-defined up to isomorphism, and we consider for each prime power q the

probability that the property holds if the isomorphism classes of curves are chosen uniformly randomly. We say that *P holds for nearly all curves* if this probability converges to 1 for $q \longrightarrow \infty$. To give an example, nearly all curves of a fixed genus over finite fields are non-hyperelliptic.

With this definition, the conjecture as given above is still not completely precise, mainly because of the unspecified logarithmic terms in the expected running time $\tilde{O}(q^{2-\frac{2}{g-2}})$. For an accurate statement, we need to go to down to a "technical" level which is usually ignored in complexity theoretic statements: We fix a random access machine model of computation with randomization and logarithmic cost function,[1] and we fix a representation of the objects involved via bit-strings (following the ideas stated above). Then we claim:

Conjecture (precise formulation). *For every fixed $g \geq 5$ there is a RAM Π, a function $f \in \tilde{O}(q^{2-\frac{2}{g-2}})$ and a constant $C > 0$ such that*

- *upon input of an instance of the discrete logarithm problem for curves of genus g, if Π terminates it outputs a solution to the problem,*
- *for some prime power q, let us consider some distribution of input instances for which the curve is distributed uniformly randomly over $\mathbb{F}_q$. Then with a probability of at least $(1 - \frac{C}{q})$ Π terminates in an expected time of $f(q)$.*

Indeed, we not only claim that there *exists* such a RAM but we also claim that one can obtain such a RAM by following the computation outlined in the next section. Also, we note again the expected time refers to every single curve. We also note that we have deleted the phrase "non-hyperelliptic" as it is unnecessary, but the algorithm definitely only operates for non-hyperelliptic curves.

Practical implications

The first computation performed in all the index calculus algorithms mentioned is the computation of the L-polynomial. For practical purposes only the group order is needed, but this computation can also be de facto impossible for curves for which the remaining parts of the algorithms can be performed without problems.

For this reason, for practical purposes, we restrict ourselves to instances where the group order is known (or can be computed easily). We then show

[1] For a straight-forward interpretation of the outlined algorithms, one might use a model with addition and subtraction. As shown in [8], it is then possible to transfer the result to a successor RAM model.

experimentally that the algorithm is practical and that indeed, for curves for which the order of the degree 0 class group is prime "up to a small cofactor", the problem for non-hyperelliptic curves of genus 5 and 6 should – by the current state of the art – be regarded as being equaly hard as the corresponding problem for curves of genus 4 and 5, respectively, *over the same field*. We note here that the computation of the plane mode in [22] is very quick and that with the index-calculus algorithm in [6] the problem for non-hyperelliptic curves of any particular fixed genus $g \geq 4$ should be regarded as being equaly hard as the problem for hyperelliptic curves of genus $g-1$ over the same field (again if the group order is prime "up to a small cofactor"). So, one can say that – under the restriction on the group order given – from a practical point of view, the discrete logarithm problem for non-hyperelliptic curves of genus 5 and 6 should be seen as being equaly hard as the problem for hyperelliptic curves of genus 3 and 4, respectively, over the same field.

This practical result has implications on the discrete logarithm problem in elliptic curves over finite extension fields and threfore also on the security of potential cryptographic schemes based on elliptic curves.

For example, in [5] it is shown that one can transfer certain instances of the discrete logarithm problem for elliptic curves over fields $\mathbb{F}_{q^5}$, q a prime power, to the corresponding problem for curves of genus 5 over $\mathbb{F}_q$. As there is no indication that the resulting curves are hyperelliptic, it suggests itself that the curves are non-hyperelliptic; experimentally this is the case. If one then applies the new algorithm to these curves, one obtains a practically relevant running time of $\tilde{O}(q^{\frac{4}{3}})$ instead of the time of $\tilde{O}(q^{\frac{5}{2}})$ for generic methods. This corresponds to a change of the bit-length by a factor of $\frac{8}{15} \approx 53\%$.

We note, however, that in contrast to, for example, the ρ-method, the storage requirements for an index calculus algorithm are always large, and they are particularly large if – as is the case here – a double large prime variation is used. Concretely, for all the mentioned index calculus algorithms, with a straight-forward implementation of the graph of large prime relations, one has to store about q relations. This can be reduced by directy iteratively constructing a tree of large prime variations. For the new algorithm, in practice the tree should have a size of about $q^{1-\frac{1}{g-2}+\frac{1}{(g-2)^2}}$, which of course can still considered to be enormous in comparison with the minimal storage requirements for the ρ-method. Also, again as usual for index calculus algorithms, there is the problem that it is difficult to parallize the linear algebra computation.

Notation and Terminology

We have already given many notations used. Generally, all the notation and terminology used follow [9] and [22]. Following [22] and in contrast to [9], a curve is not automatically smooth. In particular, we use the notation introduced in Definition 1 of [9] which might be called "asyptotically greater or equal". This is first used in Proposition 2.1 in subsection 2.2.

Outline

In the next section, we present the algorithm. For this, we first give some geometric background, then present "first ideas" for the algorithm. We recall from [7] that the task is to compute an appropriate tree of large prime relations and then give an algorithm to do so. At the end of the section, state the theoretical results on which the heuristic analysis relies. In Section 3 we then prove these results. For this, we make use of Brill-Noether theory for special linear systems. An important ingredient is here the use of linear systems on relative curves. In Section 4 we give the experimental results, and finally in Section 5 we briefly indicate how one might use even "more special" pencils to compute discrete logarithms. For the lack of a suitable reference, we give some general results on these in an appendix.

2 The algorithm

2.1 Geometric background

As already mentioned in the introduction, the algorithm relies on the consideration of special linear systems, in particular base-point free pencils. Here we briefly give some information related to this. Everything we need for the algorithm can be found in [18, Chapter II, §7]. Later we will also use the books [2] and [1] and sources cited therein. At this point, the reader might just consult the introduction to [2] as an additional source to [18].

Let $\mathcal{C}$ be a smooth curve over a field k. Let now $\mathcal{L}$ be a sheaf on $\mathcal{C}$. Then any global section s of $\mathcal{L}$ defines an effective divisor called the *divisor of zeroes* of s. There is an induced bijection from $\Gamma(\mathcal{C}, \mathcal{L})/k^*$ to this space of these divisors, thus the latter space, which we denote by $|\mathcal{L}|$ has the structure of a projective space. Such a space is called a *complete linear system* (or *series*), and a projective subspace of this a *linear system*; such a space is typically denoted by $\mathfrak{d}$. The projective space $\mathfrak{d}$ is then the image of a linear subspace of $\Gamma(\mathcal{C}, \mathcal{L})$. By definition, all divisors in a linear system are linearly equivalent.

If now a linear system $\mathfrak{d}$ is given, any divisor D of $\mathfrak{d}$ defines a sheaf $\mathcal{O}(D)$ and is then given as the divisor of zeroes of 1 on this sheaf. Then $\mathfrak{d}$ is a subspace of the complelte linear system $|D| = |\mathcal{O}(D)|$. In conclusion, there

is a bijection between linear systems and tuples $(\mathcal{L}, V)$, where V is a linear subspace of $\Gamma(\mathcal{C}, \mathcal{L})$.

Just as every projective space, a linear system has a dimension; it is usual to denote this by r. Note that the dimension of the complete linear system associated to a sheaf $\mathcal{L}$ is given as $\dim(\Gamma(\mathcal{C}, \mathcal{L})) - 1$.

Effective divisors can be identified with subschemes. Then the *base locus* of a linear system is the scheme-theoretic intersection of the divisors in the system. A system is *base point free* if its base locus is trivial.

If a morphism $\pi : \mathcal{C} \longrightarrow \mathbf{P}^r_k$ is given such that the image is not contained in a projective subspace, the pull-backs of the hyperplanes to $\mathcal{C}$ define a base point free linear system of dimension r, and the morphism is given up to change of coordinates by this system. Conversely, to every base point free system of dimension r, one can associate such a morphism such that the system is then given by the pull-backs of hyperplanes. This morphism is then also unique up to a linear change of the projective coordinates.

A 1-dimensional linear system is called a *pencil*. As a special case of what we just said, the preimages of any function define a base point free pencil, and conversely, a base point free pencil defines such a function up to change of coordinates. The degree of the function and of the pencil are then by definition identical.

2.2 First ideas for the algorithm

Let $\mathcal{C}$ be a curve of genus at least for over a finite field $\mathbb{F}_q$.

As already mentioned, we want to consider base-point free pencils of dimension $g - 1$ which are complete as linear systems.

Note the following interesting application of the Riemann-Roch theorem:

Let ω be the canonical sheaf. Then for every divisor D of degree $g - 1$, the systems $|D|$ and $|\omega(-D)|$ have the same dimension. In particular, the former is a pencil if and only if the latter is one.

To generate such pencils algorithmically, we use the ideas already presented in the introduction:

We consider some effective divisor D of degree $g - 3$. Then the system $|\omega(-D)|$ has degree $g+1$ and dimension at least 2. Heuristically, one expects that it defines a plane model of the curve. In [22] it is proven that for nearly all such divisors, this is the case. Concretely, there is a function f from the prime powers to the natural numbers converging to 1 such that for variable curves $\mathcal{C}$ over variable finite fields $\mathbb{F}_q$ the portion of divisors with the prescribed property is $\geq f(q)$.[2] Similarly, with the same framework it is

[2]Note that as usual in algorithmic considerations, we first said that the curve should be given and then varied the curve if this is appropriate. We shall proceed so also in the following.

not difficult to show that for nearly all divisors, the resulting plane models have only singularities of order 2.

Let us assume that we have such a divisor D and that the plane model $\mathcal{C}_{pm}$ and a birational morphism $\pi : \mathcal{C} \longrightarrow \mathcal{C}'_{pm}$ to it has been computed. (Note here that the curve given by a fixed plane model $\mathcal{C}_{pm}$, so π in fact is a birational map from the fixed plane model $\mathcal{C}_{pm}$ to the variable one.) is Suppose now that there is a rational singularity S of order 2. Then we can consider the base point free pencil defined by the lines through S. Let $\Delta := \pi^{-1}(S)$. Then the base point free pencil under consideration is then $|\omega(-D-\Delta)|$. Note that here $\deg(D+\Delta) = g+1$ and consequently $\deg(\omega(-D-\Delta)) = g+1$ as well.

Let us consider the relation generation, where for simplicity for the moment we avoid the double large prime variation: The factor base $\mathcal{F}$ is a subset of $\mathcal{C}(\mathbb{F}_q)$. Following [6] and [9] an evident idea is now as follows: One intersets the plane model with lines through the given singularity (through the images of points of the factor base and conviniently avoiding any lines which pass through further singular points). One considers all divisors which split completely over the factor base. If there are at least two such divisors, one mapps all these divisors back to the curve. One fixes one of them, say D_0. Then any other divisor, say D, defines a relation $[D] - [D_0] = 0$ which can be stored in the relation matrix.

For a theoretical analysis the following variant is however more convinient: Note that, as said above, the residual system of a complete one dimensional linear system of degree $g+1$ is again a complete one dimensional linear system of degree $g+1$. We can apply this to the systems $|\omega(-D-\Delta)|$ and $|D+\Delta|$. So, the system $|D+\Delta|$ is also a pencil of degree $g+1$.

In the algorithm, we then use the plane model only to compute S. From S we then compute Δ, and given this, we compute a function f in $L(D+\Delta)$, obtaining the basis $1, f$ of this space. From these, we can easily generate the relations.

For this approach, which might be called *implicit approach* in contrast to the *explicit approach* emphazising the different plane models it is easier to argue that the computations for each pencil can be performed in the desired expected time of $\tilde{O}(\#\mathcal{F}) \cdot \mathit{Poly}(\log(q))$. This is however not the main reason for this change. The reason is rather that it facilitates the overall analysis of the algorithm leading to the graph of large prime variation.

Concretely, one of the challenges is to give rigorous estimates for the number of completely split divisors in the pencils. Heuristically one expects that "usually", the number of completely split divisors in a base point free pencil of degree $g-1$ should be about $\frac{1}{(g-1)!}q$. For a rigorous estimate, the following proposition is handy:

Proposition 2.1 Let us consider base-point free pencils of a fixed degree d containing one divisor which splits completely into distinct points on curves of a fixed genus g over finite fields $\mathbb{F}_q$. Then the number of divisors in such a pencil which split completely into distinct points is $\gtrsim \frac{1}{d!} \cdot q$ and $\lesssim \frac{1}{d} \cdot q$.

Proof. Let such a curve $\mathcal{C}$ over $\mathbb{F}_q$ and such a pencil $\mathfrak{d}$ on $\mathcal{C}$ be given. Let $f : \mathcal{C} \longrightarrow \mathbf{P}^1_{\mathbb{F}_q}$ be a funtion corresponding to the pencil. Then there is a rational point of $\mathbf{P}^1_{\mathbb{F}_q}$, say Q, such that $f^{-1}(Q)$ splits completely into distinct points.

We consider the corresponding extension of function fields $\mathbb{F}_q(\mathcal{C})|\mathbb{F}_q(\mathbf{P}^1)$ and identify closed points of the curve with places. Let M be the Galois closure of this extension. Then a place of $\mathbb{F}_q(\mathbf{P}^1)$ is completely split in $\mathbb{F}_q(\mathcal{C})$ if and only if it is completely split in M; cf. Corollary III.8.4 of [27].

The place Q is by assumption completely split in $\mathbb{F}_q(\mathcal{C})$. (Note that by terminology of number / function field theory, a place which is completely split or completely decomposed in an extension is in particular unramified; this is the case here as we assume that the pencils contains a divisor which splits completely into distinct points.) As Q is completely split in $\mathbb{F}_q(\mathcal{C})$ it is also completely split in M. This in turn implies that $\mathbb{F}_q$ is the exact constant field of M.

With the effective Chebotaryov density theorem from [26], we conclude that the number of places of $\mathbb{F}_q(\mathbf{P}^1)$ of degree 1 which split completely in M (or in $\mathbb{F}_q(\mathcal{C})$) is $\frac{1}{[M:\mathbb{F}_q(\mathbf{P}^1)]} \cdot q + O(q^{\frac{1}{2}})$, which is $\gtrsim \frac{1}{d!} \cdot q$ and $\lesssim \frac{1}{d} \cdot q$.

This gives the proposition. □

Remark 2.2 A special case of this proposition for central projections via plane models was already given in [9]. We take the opportunity to point out two minor mistakes in the argument in [9]. These mistakes do not affect the result have been corrected above:

- A central projection with center P is considered. Then incorrectly the image of the point P is used where a new point, called Q above, should be used.
- It is incorrectly claimed that the number of completely split places is (in our notation) $\frac{1}{d!} \cdot q + O(q^{\frac{1}{2}})$.

We briefly explain why the considerations of pencils $|D + \Delta|$, where D and Δ are as above, is handy:

We want to consider pencils which contain by construction a divisor which splits completely into distinct points. In the analysis of the algorithm we then have to consider the number of such pencils which can be generated

with the given method and the probability that such one plane model leads to such a pencil.

For this, it is convenient to generate pencils as follows: We choose D to be completely split and then consider (assuming $|\omega(-D)|$ defines a plane model and if possible) a rational singularity of order 2 such that $D+\Delta$ splits completely into distinct points.

If we considered the residual system instead, in the analysis of the algorithm we would have to consider conditions both on the pencil itself (namely that it is defined via a singularity of a plane model, which in particular means that it contains a divisor which splits as $D + \Delta$, where D has degree $g - 3$ and Δ has degree 2) and its residual (namely that the residual has a completely split divisor).

2.3 Algorithmic background

As already stated, just as the previous algorithms, for example the ones in [7] and [6], we use a so-called double large prime relation. In Section 3.1 of [6] a general framework for the use of double large prime relation for curves of a fixed genus has been developed.

We want to consider instances of the discrete logarithm problem for curves of a fixed genus g. As already stated, the smooth curve in question, $\mathcal{C}$, shall be represented by a plane model $\mathcal{C}_{pm}$.

The factor base we want to use shall be a subset $\mathcal{F}$ of $\mathcal{C}(\mathbb{F}_q)$ of size $\lceil q^{1-\frac{1}{g-2}} \rceil$, and as usual the set of large primes is then $\mathcal{L} := \mathcal{C}(\mathbb{F}_q) - \mathcal{F}$. In our application, a *graph of large prime relations* is then an undirected labeled graph on $\mathcal{L} \dot\cup \{*\}$ with root $*$, where the labels are given in an evident way by relations involving one large prime (for edges to $*$) or two large primes (for edges connecting points of $\mathcal{L}$). A *tree of large prime relations* is defined analogously, only that it is a labeled rooted tree with root $*$. Moreover, in our application the large prime relations involve only a constant number of points of $\mathcal{F}$ (indeed at most $2(g-1)-1 = 2g-3$). Then by Proposition 11 of [6] we obtain:

Proposition 2.3 Given the data just described and a tree of large prime relations of size at least $q^{1-\frac{1}{g}+\frac{1}{(g-2)g}}$ and a depth which is polynomially bounded in $\log(q)$, one can compute any instance of the discrete logarithm problem for the particular curve in an expected time of

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

Without considering methods to find relations, there are different methods to construct such a tree:

- One method is to construct the tree directly by inserting a new relation if it connects to the tree. This is used to prove the result in [13].
- A slightly different method is to construct the tree in stages. This is used in [7] and in [9].
- The classic method is to first construct an approriate *graph of large prime relations* on $\mathcal{L} \dot{\cup} \{*\}$ and then to construct a tree from this graph. This is used for the practical result in [13] and the heuristic results in [6] and [10].

Here, we pursue the same approach as in [6] and [10], and we use similar heuristic assumptions. Explicitly, the goal is to construct a graph on $\mathcal{L} \dot{\cup} \{*\}$ (which is of size $\sim q$) with roughly q edges. From this graph we then construct a tree of large prime relations with a breath-first search starting with $*$. For the analysis, the graph constructed is then compared with a suitable random graph in a standard model.

One of the classes of well-studed random graphs are Bernoulli random graphs. Here, a natural number n and a real number p between 0 and 1 are fixed and then the random graph is considered were each edge occurs with a probability of p. Let $\mathbb{G}(n,p)$ be this random graph. Note that the expected value of edges in this graph is $p \cdot \frac{n(n-1)}{2}$. Then with [3] we have the following result:

Proposition 2.4 Let $c > 1$. Then there are constants $c_1, c_2 > 0$ such that with a probability converging to 1 for $n \longrightarrow \infty$, the graph $\mathbb{G}(n, \frac{c}{n})$ has a connected component of size at least $c_1 n$ and diameter at most $c_2 \log(n)$.

This result suggets that one is on the right track here.

2.4 Construction of the tree of large prime relations

We now describe the algorithm for the construction of a tree of large prime relations for conjectured results.

The algorithm starts by the construction of the factor base $\mathcal{F}$ by choosing a subset of size $\lceil q^{1-\frac{1}{g-2}} \rceil$ uniformly at random from the subsets of $\mathcal{C}(\mathbb{F}_q)$ of this size. For this, rational points are chosen uniformly at random until the disired set is constructed. The construction can be done in an expected time of $\tilde{O}(q^{1-\frac{1}{g-2}})$; see also [7].

The reason for choosing the factor base in this randomized way is that then nicely make use of probabilistic methods for the analysis of the relation generation. This idea was already applied fruitfully in [6] and is applied dirctly below in Lemma 2.5.

Then pencils of degree $g-1$ are constructed in the way outlined in subsection 2.2. Let us assume we have such a pencil. Now, we want to obtain relations with up to two large primes from this. Particularly, we want that every divisor which splits completely into factor base elements and one or two large primes leads to such a relation. There is however a problem here now: Relations are given by differences of divisors, and for this approach, we need at least one relation which splits completely into factor base elements. This is however not the case most of the time (under the randomized choice of the factor base and any choice of the pencils), as is shown by this lemma:

Lemma 2.5 For a pencil of degree $g-1$ with a divisor which splits completely into distinct points,

a) the expected number of divisors which split over the factor base is $\gtrsim \frac{1}{(g-1)!} \cdot q^{-\frac{1}{g-2}}$ and $\lesssim \frac{1}{g-1} \cdot q^{-\frac{1}{g-2}}$.

b) the expected number of divisors which split into points of the factor base and one large prime is $\gtrsim \frac{1}{(g-2)!}$.

c) the expected number of divisors which split into points of the factor base and two large primes is $\gtrsim \frac{1}{(g-3)!} \cdot q^{\frac{1}{g-2}}$.

Proof. Let D be a divisor which splits completely into distinct points. Then the probability that it splits completely over the factor base is $\sim q^{-\frac{g-1}{g-2}}$, the probability that it splits completely into points of the factor base and one large prime is $\sim (g-1) \cdot q^{-\frac{g-2}{g-2}} = (g-1) \cdot q^{-1}$, and the probability that it splits completely into points of the factor base and two large primes is $\sim (g-1) \cdot (g-2) \cdot q^{-\frac{g-3}{g-2}}$.

As there are $\gtrsim \frac{1}{(g-1)!} \cdot q$ (by Proposition 2.1) and $\lesssim \frac{1}{g-1} \cdot \frac{1}{q}$ such divisors, the result follows. □

For this reason, whenever we have constructed a pencil, we add the points of the divisor for its construction, that is the completely split divisor $D+\Delta$, to the factor base.

The third item of the lemma suggests that one should consider $\lceil (g-3)! \cdot q^{1-\frac{1}{g-2}} \rceil$ pencils, and this is what we do. For this, we choose the effective divisor D of degree $g-3$ uniformly at random, and if $|\omega(-D)|$ defines a plane model (which is then of degree $g+1$), we consider all rational singularities of order 2, compute the corresponding divisors Δ of order 2, and among the ones for which $D+\Delta$ splits completely into distinct points, we choose one divisor Δ uniformly at random. This then gives the pencil $|D+\Delta|$ we consider.

In this way, we generate $(g-3)! \cdot q^{1-\frac{1}{g-2}}$ distinct pencils of degree $g-1$ each containing a divisor which splits completely into distinct points. (We check that the pencils are distinct with the divisors $D+\Delta$ constructed.) For each of the pencils we generate relations with one or two large primes as described.

This gives the following overview of the algorithm.

Algorithm: Construction of the tree of large prime relations

Input: A curve $\mathcal{C}/\mathbb{F}_q$ of genus $g \geq 4$, represented by a plane model of bounded degree.

Output: A factor base and a tree of large prime relations.

1. Construct a list of appropriate pencils:

   Construct a list of $\lceil (g-3)! \cdot \lceil q^{1-\frac{1}{g-2}} \rceil$ pairwise non-linearly equivalent divisors D_i of degree $g-1$ splitting completely into distinct points and defining pencils as follows:

   For $i \longleftarrow 1$ to $\lceil (g-3)! \cdot \lceil q^{1-\frac{1}{g-2}} \rceil$ do:

   Choose a divisor D of degree $g-3$ splitting completely into distinct points uniformly at random among all such divisors.

   sCompute a basis of $L(K-D)$, where K is a canonical divisor.

   If the dimension of $L(K-D)$ is larger then 3, go back to the beginning of the loop.

   Compute the image of the curve $\mathcal{C}$ in $\mathbf{P}^2_{\mathbb{F}_q}$ defined by the computed basis of $L(K-D)$. If this image does not have degree $g+1$ go back the the beginning of the loop.

   (Now a plane model and morphism to it have been computed; denote this by $\pi : \mathcal{C} \longrightarrow \mathcal{C}'_{pm}$.)

   Compute the rational points of the singular locus of $\mathcal{C}'_{pm}$. For each such rational point, compute the corresponding divisor Δ. If there is no such divisor Δ of degree 2 such that $D+\Delta$ split completely into distinct rational points, go back to the beginning of the loop.

   Choose one divisor Δ from the ones constructed uniformly at random. If $D+\Delta$ is linearly equivalent to any of the divisors $D_1, \ldots, D_{i-1}$ go back to the beginning of the loop.

   Let $D_i \longleftarrow D+\Delta$

2. Construct the *factor base* $\mathcal{F}$:

   Choose a set $\mathcal{F}$ of size $\lceil q^{1-\frac{1}{g-2}} \rceil$ uniformly at random among the subsets of $\mathcal{C}(\mathbb{F}_q)$ of this size.

3. Construct a *graph of large prime relations* on $\mathcal{L} \dot{\cup} \{*\}$, where $\mathcal{L} := \mathcal{C}(\mathbb{F}_q) - \mathcal{F}$:

   For $i \longleftarrow 1$ to $\lceil (g-3)! \cdot q^{1-\frac{1}{g-2}} \rceil$ do:

   Insert the points in D_i into the factor base. For each point P of $\mathcal{F}$, compute the unique divisor D in $|D_i - P|$. If it splits into points of the factor base and one or two large primes, store the relation defined by $D - D_i$ in the tree of large prime relations, provided that there is not yet an edge between the vertices in question.

4. Construct a *tree of large prime relations*:

   Use a breadth-first search starting with $*$ to construct a tree of large prime relations $\mathcal{T}$ from the graph. Stop this construction if the tree has $\lceil q^{1-\frac{1}{g}+\frac{1}{(g-2)g}} \rceil$ points.

5. Check if the tree is appropriate: If the tree has a diameter $\geq \log^2(q)$, go back to the beginning.

6. Output $\mathcal{F}$ and $\mathcal{T}$.

Remarks 2.6

- As already stated, the purpose of this algorithm, with the more detailed information on its steps below, is to argue that the conjectures in the introduction are correct. For practical purposes, we propose an algorithm with an explicit use of the plane model for relation generation, following the first ideas presented in the introduction and in subsection 2.2.

- We have put Step 1 at the beginning rather than mixing it with Step 3 after the construction of the factor base to facilitate the understanding and analysis of the algorithm.

- One feature of the algorithm, convenient for the analysis, is (by Lemma 2.5): Conditionally to any outcome of Step 1, the expected value of (distinct) relations (with one or two large primes) generated in Step 3 is $\gtrsim q$. Note that this does however not mean by itself that the number of edges in the graph is then $\gtrsim q$. The reason is that there might be different relations for the same tuple of large primes.

- In Step 3, for each i, the points of the divisor D_i are inserted in order to be able to store the relations. There is also the following alternative approach for this: If D is a divisor in the pencil $|D_i|$ which splits completely into points of the factor base and one or two large primes, one can store the relation generated also by storing the information for D and the divisor D_i or also the class $[D_i]$. Now, the class $[D_i]$ can simply be represented by i. One can then say that with this approach the factor base is augmented with the divisors D_i instead of the points in it.

- The construction of the tree of large prime relations from the graph in Step 4 and the criterion in Step 5 are adapted from [6]. The exponent in the criterion that the diameter be $\leq \log^2(q)$ is quite arbitrary.

With the ideal arithmetic for divisors, up to factors polynomial in $\log(q)$, the computations with divisors in the algorithm can be carried out as fast as one can reasonably expect. We refer again to Section 2 of [7] for more information on the basic computations. Now we go through the different steps. When a computation can be perfomed in an (expected) time which is polynomially bounded in $\log(q)$ (and therefore in the input size), we say that it can be performed in expected polynomial time.

Step 1. To compute the divisor D, we compute uniformly distributed rational points (disregarding dublicates) until we have $g-3$ distinct points. This can be done in a expected polynomial time. The same is true for the computation of the canonical divisor K and the basis f_1, f_2, f_3 of the space $L(K-D)$. Then the image (i.e. the polynomial defining the image) can be computed in polynomial time as well,[3] and it is trivial to check the degree condition.

Let $F' \in \mathbb{F}_q[X,Y,Z]$ be the polynomial defining the plane model $\mathcal{C}'_{pm}$. Then the singular locus of $\mathcal{C}'_{pm}$ is defined by F and its partial derivatives. Now, the degrees of the polynomials in this system are bounded and the scheme defined is 0-dimensional. The rational points can then be computed in expected polynomial time with a Gröbner base computation.

Let now P be such a point. The task is now to compute the corresponding divisor Δ on $\mathcal{C}$. For this, we compute a basis g_1, g_2 of the subspace of $L(K-D)$ vanishing at P on $\mathcal{C}'_{pm}$. Explcitly, if, say, $P=(a:b:1)$, then $f_1 - af_3, f_2 - bf_3$ is such a basis.

From this, we compute Δ as $\Delta = K - D + \inf(\mathrm{div}(g_1), \mathrm{div}(g_2))$.

A crucial question is now if Step 1 can be completed at all, that is, if enough divisors can be computed, and also what the probability is that one

[3] Zitat auf Kochinke

iteration of the loop runs through. This will be addressed below.

Step 2. This can be performed in an expected time of $\tilde{O}(q^{1-\frac{1}{g-2}})$.

Step 3. For one divisor and one point of the factor base the computation can be performed in expected polynomial time. In total, the complete step can be performed in an expected time of $\tilde{O}(q^{2-\frac{2}{g-2}})$.

Step 4 to 6. These step can clearly be performed in a time of $\tilde{O}(q^{1-\frac{1}{g}+\frac{1}{(g-2)g}})$.

We have included curves of genus 4 only for completeness. In fact, for such curves, Step 1 never terminates because:

Proposition 2.7 On a non-hyperelliptic curve of genus 4, there are at most two pencils of degree 3. These are base pont free and complete.

It is easy to see that any such pencil is base point free and complete. The fact that there are only two such pencils is proven below the "Existence Theorem" in [2, V]. It is also shown that on a general curve there are two such systems, and that these are residual to each other, that is, if $|D|$ is the one system then $|K - D|$ is the other. Moreover, there are particular curves with only one such pencil, which is then autoresidual.

The proposition is in line with the following statement of the theory of Brill and Noether on special linear systems: For any non-hyperelliptic genus g at least 4 over an algebraically closed field, the pencils move in a family of dimension $(g - 4)$. The necessary definitions for to make this statement accurate and the accurate statement itself can be found in the next section. There are still a lot of obstracels for a precise analysis, but this gives a first hint on what to expect.

Indeed, as we shall show in the next section, we have:

Proposition 2.8 Let us consider curves of a fixed genus $g \geq 5$ over finite fields $\mathbb{F}_q$ such that $q \geq g$ or $g \leq 7$. Then for curves in all but a portion of $O(\frac{1}{q})$ isomorphism classes, the following holds.

a) There are $\sim \frac{1}{(g-1)!} \cdot q^{g-3}$ distinct effective divisors of degree $g - 1$ and dimension 1 which can be given as $D + \Delta$, where D and Δ are effective divisors, D has degree $g - 3$, Δ has degree 2, $|K - D|$ defines a plane model and Δ defines a singularity (of order 2) on the plane model.

b) The probability that an effective divisor D of degree $g - 3$ which splits comletely into distinct points gives rise to a base point free system $|K-D|$ defining a plane model with a rational singularity defined by a divisor Δ

of degree 2 such that $D + \Delta$ splits completely into distinct points is in $\Theta(1)$.

c) Let D be uniformly distributed among divisors of degree $(g-3)$ which split completely into distinct points such that the condition of b) holds. Now let (D, Δ) be distributed such that Δ is distributed uniformly for each value D_0 of D. Then $D + \Delta$ is distributed such that each value (which is attained at all with non-trivial probability) is obtained with a probability of $\Omega(q^{g-3})$.

Remark 2.9 We need the condition that g be $\geq q$ in one step of the proof. We conjecture that the corresponding statement is always valid and therefore that the statements in the proposition hold without this condition.

The proposition implies the following statement on Step 1 of the algorithm:

Proposition 2.10 Let us consider curves of a fixed genus $g \geq 5$ over finite fields $\mathbb{F}_q$ such that $q \geq g$ or $g \leq 7$. Then for curves in all but a portion of $O(\frac{1}{q})$ isomorphism classes, the following holds.

a) With a probability of $\Theta(1)$, a divisor D as in the algorithm leads to at least one tuple (D, Δ) as in the algorithm.

b) If a divisor $D + \Delta$ is chosen, $D + \Delta$ is distributed among $\sim \frac{1}{(g-1)!} \cdot q^{g-3}$ divisors in such a way that each value is obtained with a probability of $\Theta(q^{g-3})$. Furthermore, then $|D+\Delta|$ is distributed among $\Theta(q^{g-4})$ pencils in such a way that each pencil is obtained with a probability of $\Theta(q^{g-4})$.

c) Any iteration of the for-loop can be performed in expected polynomial time.

d) Step 1 can be performed in an expected time of $\tilde{O}(q^{1-\frac{1}{g-2}})$.

Here, item a) and the first sentence of b) are just applications of the previous proposition, and the further statements follow immediately.

Besides the gap to prove the result without the condition $g \geq q$ or $g \geq 7$, it remains to analyse the tree of large prime relations. Concretely, it remains to show that the condition in Step 5 is satisfied with a probability of $\frac{1}{q^{O(1)}}$.

We are not able to perform this anlysis, so we rely on a comparison with random graphs related to $\mathbb{G}(n, p)$.

As sated in Remark 2.6, the algorithm has the property that the expected number of relations generated in Step 3 is $\gtrsim q$. A first indication that the analysis is accurate is Proposition 2.4 applied with $n = q$ and $p = \frac{2}{q}$ (such that the expected number of edges is $\sim q$). The conclusion is that there are

constants $c_1, c_2 > 0$ such that with a probability converging to 1 for $q \longrightarrow \infty$, the graph $\mathbb{G}(q, \frac{2}{q})$ has a large connected component of size at least $c_1 n$ and diameter at most $c_2 \log(q)$.

Another kind of "standard random graph" are uniform random graphs. For two appropriate natural numbers n, m the uniform random graph $\mathbb{G}(n, m)$ is the random graph obtained by choosing a set of m edges uniformly at random from the set of edges of this size. One sees rather easily that for any $c > 1$, the conclusion of Proposition 2.4 also hold for the graphs $\mathbb{G}(n, \lceil \frac{c}{2} n \rceil)$; cf. Proposition 10 in [10]. We apply this with the graphs $\mathbb{G}(q, q)$ for $q \longrightarrow \infty$.

As a variant, we consider a random graph for natural numbers n, m where n edges are drawn uniformly but with replacement from the set of edges. Let us denote this random graph by $\mathbb{G}^*(n, m)$. For our applications we note that for any $c > 1$ and any $d < c$, for nearly all n, the graph $\mathbb{G}^*(n, \lceil \frac{c}{2} n \rceil)$ has $\geq \frac{d}{2} n$ edges with probability converging to 1 and therefore the conclusions of Proposition 2.4 also hold for $\mathbb{G}^*(n, \lceil \frac{c}{2} n \rceil)$. So in particular they hold for $\mathbb{G}^*(q, q)$ with $q \longrightarrow \infty$.

To model the situation considered here even better, we consider for some natural numbers n, m and a real number p between 0 and 1 the random graph obtained by the following procedure:

1. $\mathbb{G}$ is set to be the empty graph on the set $\{1, \ldots, n\}$.
2. For $i \longleftarrow 1$ to m:
   With a probablity of p:
   One edge is chosen uniformly at random
   if this edge is not yet contained in $\mathbb{G}$, it is inserted.

We denote this graph by $\mathbb{G}^*(n, m, p)$.

Ths graph can also be described as follows: It is the random graph $\mathbb{G}^*(n, m')$, where m' is distributed according to the binomial distribution for m tries and success probability p. In particular, the expected number of edges chosen edges is mp.

In the algorithm, the probability that one divisor which is completely split in one of the considered pencil splits into elements of the factor base and two large primes is $\sim (g-1) \cdot (g-2) \cdot q^{-\frac{g-3}{g-2}} = (g-1) \cdot (g-2) \cdot q^{\frac{1}{g-2}-1}$, and in total there are $\gtrsim \frac{1}{(g-1)\cdot(g-2)} \cdot q^{2-\frac{1}{g-2}}$ such divisors in all the considered pencils together. (In the computation, for efficiency reasons, we only consider divisors containing an element of the factor base, but we can think of the graph as being constructed as follows: First all completely split divisors in all the pencils $|D_i|$ are computed and then the factor base is chosen and it is determined which divisors lead to edges in the graph.) This suggests to model the situation with the graph $\mathbb{G}^*(q, \frac{1}{(g-1)\cdot(g-2)} \cdot q^{2-\frac{1}{g-2}}, (g-1)\cdot(g-2)\cdot q^{\frac{1}{g-2}-1})$.

With the Chebychev inequality we see that the conclusion is just as for the graph $\mathbb{G}^*(q,q)$ and then also for $\mathbb{G}(q,q)$.

This concludes our analysis of the algorithm. Based on it, we are confident to say that the conjecture raised in the introduction holds.

3 Geometric analysis

To prove Proposition 2.8, we make use of Brill-Noether theory for special linear systems on curves. We then combine these results with the results of Section 2 of [22] to get probabilistic estimates.

We follow closely [2] and [1] in notations. Concerning techniques and results, the last chapter of the second book, that is, Chapter 21 called "Brill-Noether theory on a moving curve") (which can be seen as a book for itself) is of particular importance.

There is, however, the problem that in these books analytic techniques are used: The first book is written from an analytic point of view (which also and in particular means that even if purely algebraic techniques are used it is assumed that the characteristic is 0). This is less so for the second book, but also there analytic techniques are sometimes used. One example is Lemma 2.12 in [1, Chapter 21] in which analytic spaces are used.

For this reason, we state all the results we wish to use (and some more for a more complete picture) and point to sources in which they are proven algebrico-geometrically.

We also make use of the definitions and resuls introduced in Section 1 of [22]. We would like to suggest to the reader to read Sections 1 and 2 of [22] before continuing.

In addition to the objects introduced in [22], for a smooth relative curve $\mathcal{C}/S$, we make use of the S-schemes $\mathcal{W}^r_d(\mathcal{C})$ and $\mathcal{G}^r_d(\mathcal{C})$. Let us recall the $\mathcal{W}^r_d(\mathcal{C})$ paramerizes complete linear systes of degree d and dimension at least r in the fibers of $\mathcal{C}$ over S and that $\mathcal{G}^r_d(\mathcal{C})$ parameterizes (not necessarily complete) linear systems of degree d and dimension exactly r in the fibers of $\mathcal{C}$ over S. These definitions can be found in [1, Chapter 21, §3]. The definitions for $\mathcal{G}^r_d(\mathcal{C})$ are recalled in subsection A.

3.1 Brill-Noether theory

Let us recall the definition if the *Brill-Noether number*:

For non-negative integers g, d, r this number is defined by

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

The Brill-Noether number plays an important role in several statements which are summerized below.

Proposition 3.1 Let $\mathcal{C}$ be a nonsingular curve of genus $g \geq 1$ over some algebraically closed field k. Fix integers d, r with $d \geq 1$, $r \geq 0$ and denote $\rho(g, d, r)$ by ρ.

a) If $\rho \geq 0$ then $\mathcal{G}^r_d(\mathcal{C})$, $\mathcal{W}^r_d(\mathcal{C})$ and $\mathcal{C}^r_d$ are non-empty.

b) Every component of $\mathcal{G}^r_d(\mathcal{C})$ has dimension at least ρ. Similarly, if $r \geq d-g$ every component of $\mathcal{W}^r_d(p)$ has dimension at least ρ and every component of $\mathcal{C}^r_d$ has dimension at least $\rho + r$.

c) If $\rho \geq 1$ then $\mathcal{G}^r_d(\mathcal{C})$ and $\mathcal{W}^r_d(\mathcal{C})$ are connected.

Proof. All statements are summarized in [2, V], Theorem 1.1 and Theorem 1.4 for complex curves. The original sources [21] and [20] for a) and b), and [12] for c) provide proofs valid over algebraically closed fields of any characteristic. For a) and b) the authors do not restrict the characteristic from the start, whereas c) is proven for complex curves but is also valid in any characteristic by Remark 2.8 in [12]. □

Proposition 3.2 Fix an algebraically closed field k and let $\mathcal{C}$ be a general nonsingular curve of genus $g \geq 1$ over k. Furthermore, fix integers d, r, g with $d, g \geq 1$, $r \geq 0$ and denote $\rho(g, d, r)$ by ρ.

a) The scheme $\mathcal{G}^r_d(\mathcal{C})$ is smooth of dimension ρ. If $r \geq d - g$ then $\mathcal{W}^r_d(\mathcal{C})$ is of dimension ρ, as well. In particular, if $\rho < 0$ then $\mathcal{G}^r_d(\mathcal{C})$ and $\mathcal{W}^r_d(\mathcal{C})$ are empty.

b) If $\rho \geq 1$ then $\mathcal{G}^r_d(\mathcal{C})$ and $\mathcal{W}^r_d(\mathcal{C})$ are irreducible.

c) Suppose that $r \geq d - g$ and $\rho \geq 1$. Then $\mathcal{C}^r_d - \mathcal{C}^{r+1}_d$ is irreducible of dimension $\rho + r$.

Proof. The statements a) and b) are summarized in [2, V] for complex curves. Statemnt a) for $\mathcal{G}^r_d(\mathcal{C})$ follows from Gieseker's theorem ([14]) which states: For a general curve $\mathcal{C}$ of genus g and an effective divisor D on $\mathcal{C}$ the cup-product homomorphism

$$H^0(\mathcal{C}, \mathcal{O}(D)) \otimes H^0(\mathcal{C}, \omega(-D)) \longrightarrow H^0(\mathcal{C}, \omega)$$

is injective. As stated in [2, V] the result for $\mathcal{W}^r_d(\mathcal{C})$ then follows easily.

With Proposition 3.1 c) this implies b) for $\mathcal{G}^r_d(\mathcal{C})$. The result for $\mathcal{W}^r_d(\mathcal{C})$ then follows immediately as the image of an irreducible space is irreducible.

On c): The scheme $\mathcal{C}^r_d - \mathcal{C}^{r+1}_d$ is non-empty by Lemma [2, IV], Lemma 1.7. Moreover by Lemma 1.6 (which holds in arbitrary chacteristic) and Gieseker's result the scheme is smooth. The fiber of $(\mathcal{C}^r_d - \mathcal{C}^{r+1}_d) \longrightarrow (\mathcal{W}^r_d(\mathcal{C}) -$

$\mathcal{W}_d^{r+1}(\mathcal{C})$ are projective spaces of dimension d. By applying this to the generic point of $\mathcal{W}_d^r()$, we obtain that $\mathcal{C}_d^r - \mathcal{C}_d^{r+1}$ has an irreducible component V which maps domimently to $\mathcal{W}_d^r(\mathcal{C}) - \mathcal{W}_d^{r+1}(\mathcal{C})$ and has dimension $\rho + d$. It remains to show that $V = \mathcal{C}_d^r - \mathcal{C}_d^{r+1}$. As $\mathcal{C}_d^r - \mathcal{C}_d^{r+1}$ is smooth for this it suffices to show that $\mathcal{C}_d^r - \mathcal{C}_d^{r+1}$ is connected.

The morphism $(\mathcal{C}_d^r - \mathcal{C}_d^{r+1}) \longrightarrow (\mathcal{W}_d^r(\mathcal{C}) - \mathcal{W}_d^{r+1}(\mathcal{C}))$, which is obtained by base change from $\mathcal{C}_d^r \longrightarrow \mathcal{C}_d^{r+1}$, is proper. Therefore, V maps surjectively to $\mathcal{W}_d^r(\mathcal{C}) - \mathcal{W}_d^{r+1}(\mathcal{C})$. Suppose that there is another component V' of $\mathcal{C}_d^r - \mathcal{C}_d^{r+1}$. Let $P \in V'$. Then the fiber containing P is a projective space, thus connected. This fiber has non-trivial intersection with V and V', a contradition. □

Additionally we have the following theorem due to H. Martens in [24]. Here we cite its version as in [2, IV], Theorem 5.1.

Proposition 3.3 Let $\mathcal{C}$ be a smooth curve of genus $g \geq 3$ over some algebraically closed field k. Let d be an integer such that $2 \leq d \leq g-1$ and let r be an integer such that $0 < 2r \leq d$. If $\mathcal{C}$ is not hyperelliptic then

$$\dim\left(\mathcal{W}_d^r(\mathcal{C})\right) \leq d - 2r - 1 \ .$$

If $\mathcal{C}$ is hyperelliptic then

$$\dim\left(\mathcal{W}_d^r(\mathcal{C})\right) = d - 2r \ .$$

3.2 Pencils of degree $g-1$

We are interested in pencils of degree $g-1$. The following lemma shows that we are on the right track by considering plane models with singularities of order 2 to generate pencils of order $g-1$:

Lemma 3.4 Let k be an algebraically closed field.

a) Let $\mathcal{C}$ be a non-hyperelliptic curve of genus g over k. Then each effective divisor D on $\mathcal{C}$ of degree $g-3$ is part of a divisor in a $\mathfrak{g}^1_{g-1}$.

b) Let $\mathcal{C}$ be a general curve of genus g over k. Then a general divisor $D \in \mathcal{C}_{g-3}$ leads to a plane model such that all its singularities are of order 2.

Proof. a) Let $D \in \mathcal{C}_{g-3}$ and denote a canonical divisor on $\mathcal{C}$ by K. We distinguish three cases. First, assume $\dim(|\omega(-D)|) \geq 3$. Then for any effective divisor D' of degree 2 we have $\dim(|\omega(-(D+D'))|) \geq 1$ and thus $\dim(|D+D'|) \geq 1$ as well.

Now, suppose $\dim(|\omega(-D)|) = 2$ and $|\omega(-D)|$ possesses a base point P. Then $\dim(|\omega(-(D+P))|) = 2$ and thus for any point P' $\dim(|D+P+P'|) = \dim(|\omega(-(D+P+P'))|) \geq 1$.

Finally, suppose $|K-D|$ defines a morphism $\varphi : \mathcal{C} \to \mathbf{P}^2_k$. This morphism cannot be an embedding. So there is a point $P \in \varphi(\mathcal{C})$ such that $\varphi^{-1}(P)$ has degree at least 2, say $\varphi^{-1}(P) = P_1 + P_2 + \cdots + P_k$. Denote by $\mathfrak{g}^1_{g+1-r}$ the pencil given by lines through P. Then $|D + P_1 + P_2|$ is a $\mathfrak{g}^1_{g-1}$ on $\mathcal{C}$ since its residual is $\mathfrak{g}^1_{g+1-r} + p_3 + \ldots + p_r$.

b) One sees easily that any other possiblity contradicts the conditions given by the theory of Brill and Noether. □

As already stated in Proposition 2.7, there are only at most two pencils of order 2 on a curve of genus 4. We now consider curves of some genus $g \geq 5$. By Proposition 3.1 and Martens' theorem we see that the $\mathcal{W}^1_{g-1}(\mathcal{C})$ and $\mathcal{G}^1_{g-1}(\mathcal{C})$ are connected and $(g-4)$-dimensional.

More important is for us to study the space $\mathcal{C}^1_{g-1}$. We start with the results for general curves:

Proposition 3.5 Let $\mathcal{C}$ be a general curve of genus ≥ 5. Then the space $\mathcal{C}^1_{g-1} - \mathcal{C}^2_{g-2}$ is irreducible and $(g-3)$-dimensional.

The task is now to establish that for curves of genus at least 5 over finite fields there are enough pencils suitable for the algorithm.

For this, just as in [22] we study the univeral family of three-canonically embedded curves $\mathcal{Z}_g \longrightarrow \mathcal{H}_g$. We note that the reason we study particularly this family is that for any prime power q, every isomorphism class of curves is represented by an equal number of elements in $\mathcal{H}_g(\mathbb{F}_q)$.

As already said, the following results are for $g \geq 5$.

Lemma 3.6 There is an open subscheme U_0 of $\mathcal{H}_g$ such that for any geometric point u of U_0, the fiber $(\mathcal{Z}^1_{g,g-1} - \mathcal{Z}^2_{g,g-1})_u$ is non-empty and irreducible and such that $U_0 \longrightarrow \mathrm{Spec}(\mathbb{Z})$ is surjective (or equivalently such that for any field k, $(U_0)_k$ is non-trivial).

Proof. This follows from the general statement in [17, IV, Theorem 9.7.7] and the fact that in any characteristic the generic fiber is irreducible. □

Lemma 3.7 There is a unique open subscheme U_1 of $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ such that a geometric point u of $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ lies in U_1 if and only if the following holds: Let s be the image of u on $\mathcal{H}_g$. Then the pencil on $(\mathcal{Z}_g)_s$ corresponding to u is base point free. This scheme maps surjectively to $\mathrm{Spec}(\mathbb{Z})$.

Proof. Note that the fiber $(\mathcal{Z}_g)_s$ of $\mathcal{Z}_{g,g-1}/\mathcal{H}_g$ at s is equal to the fiber of $(\mathcal{Z}_g \times_{\mathcal{H}_g} \mathcal{G}^1_{g-1}(\mathcal{Z}_g))/\mathcal{G}^1_{g-1}(\mathcal{Z}_g))$ at u.

Now the statement that U_1 is open is a special case of Lemma A.7. Moreover, in every characteristic a generic divisor on a generic curve is base point free. □

In order to show that there are "enough" divisors which split completely into distinct points, we consider base point free pencils corresponding to functions defining what we call an *ordinary covering* (see also [22, Definition 11]):

Definition 3.8 Let $f : \mathcal{C} \longrightarrow \mathbf{P}^1_k$ be a function of degree d on a smooth curve $\mathcal{C}$ over a field k. An *ordinary branch point* of f is a $\overline{k}$-valued point of $\mathbf{P}^1_k$ which has exactly $d-1$ preimages in $\mathcal{C}(\overline{k})$. A

Proposition 13 from [22] then says:

Proposition 3.9 Let us consider base point free pencils of a fixed degree $d > 2$ defining an ordinary covering on curves of a fixed genus over finite fields $\mathbb{F}_q$. Then the number of divisors in such a pencil which split completely into distinct points is $\sim \frac{1}{d!} \cdot q$.

Let us note that this notion is closely related to the notion of a *simple covering* as defined in [11]. There is however a slight difference. Let us first recall the definition from [11]:

Definition 3.10 A *simple divisor* on a smooth curve $\mathcal{C}$ over a field k is a divisor on $\mathcal{C}$ such that over $\overline{k}$ it splits into distinct points. (If k is perfect this is equivalent to splitting into distinct topological points already over k.) A function $f : \mathcal{C} \longrightarrow \mathbf{P}^1_k$ of degree d on a curve $\mathcal{C}$ is a *simple covering* if it is generically unramified and its discriminant divisor is simple.

The two conditions are equivalent if $\mathrm{char}(k) \neq 2$. However the condition for simple coverings is never satisfied in characteristic 2. This is why we introduce the notion of an ordinary covering.

The following proposition, which was already used in [9], is well known:

Lemma 3.11 There is a unique open subscheme U_2 of $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ such that a geometric point u of $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ lies in U_2 if and only if the following holds: Let s be the image point on $\mathcal{H}_g$. Then the pencil on $(\mathcal{Z}_{g,g-1})_s$ corresponding to u is base point free amd defines an ordinary covering.

Proof. Let Δ_2 be the diagonal in $(\mathcal{Z}_g)^2 := \mathcal{Z}_g \times_{\mathcal{H}_g} \mathcal{Z}_g$ and Δ_3 the diagonal in $(\mathcal{Z}_g)^3$. Note Δ_2 and Δ_3 are closed in the surrounding spaces.

Let Δ be the image of $((\mathcal{Z}_g)^{g-5} \times \Delta_2 \times \Delta_2) \cup ((\mathcal{Z}_g)^{g-4} \times \Delta_3)$ in $\mathcal{Z}_{g-1}$ under $(\mathcal{Z}_g)^{g-1} \longrightarrow \mathcal{Z}_{g,g-1}$. Note that as the morphism is proper Δ is closed in $\mathcal{Z}_{g,g-1}$.

As shown in subsection A.3, particularly Proposition A.17, there is a projective line bundle $\mathbf{P} \longrightarrow \mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ such that for a geometric point u of

$\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$ lying over a geometric point s of S, the fiber of $\mathbf{P}$ over s parameterizes the divisors D in the pencil given by u on $\mathcal{C}_s$. Moreover, there is a morphism $\mathbf{P} \longrightarrow \mathcal{C}^1_{g-1}$ which corresponds to the assignment $(\mathfrak{d}, D) \mapsto D$.

Let $\tilde{V}$ be the preimage of Δ in $\mathbf{P}$. Finally, let V be the image of $\tilde{V}$ under the projection from $\mathbf{P}$ to $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$. Now V is yet again closed as $\mathbf{P}$ is proper over $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$.

Now for any geometric point u of $\mathcal{G}^1_{g-1}(\mathcal{Z}_g)$, u lies in V if and only if the following holds: Let s be the image of u on $\mathcal{H}_g$. Then u defines a complete linear system which contains a divisor on $(\mathcal{Z}_g)_s$ of the form $2P_1 + 2P_2 + P_3 + \cdots + P_{g-3}$ or $3P_1 + P_2 + \cdots + P_{g-3}$ for points P_i.

Let $U := \mathcal{G}^1_{g-1}(\mathcal{Z}_g) - V$ and $U_2 := U \cap U_1$. Then U_2 has the desired property. □

We would now desire to prove that U_2 maps surjectively to Spec($\mathbb{Z}$). For this it would suffice to show that in any characteristic there is *at least one* ordinary covering $\mathcal{C} \longrightarrow \mathbf{P}^1_k$ with k an algebraically closed field of characteristc p with $g(\mathcal{C}) = g$.

For characteristic 0 or characteristic at least g one can refer to W. Fulton's work [11] on Hurwitz schemes (for simple coverings). We therefore habe:

Lemma 3.12 The image of the scheme U_2 as the previous lemma in Spec($\mathbb{Z}$) contains all primes $\geq g$.

Of course, we would like to prove that the scheme U_2 in the lemma maps surjectively to Spec($\mathbb{Z}$). This seems to be a difficult question (see also the remark following the next lemma), but for a particular genus and particular characteristic, one can try to construct such curves algorithmically. For this, one can proceed as in the algorithm: One fixes an effective divisors D of degree $g-3$ until one has a plane model with a rational singularity of order 2. Then one computes, via the polar curve, the tangents to points of the curve running trough such a singularity.

We performed the necessary computations for curves of genus ≤ 7. We therefore have:

Lemma 3.13 Let U_2 be as in the previous lemma. Then for $g \leq 7$, $U_2 \longrightarrow$ Spec($\mathbb{Z}$) is surjective.

Remark 3.14 For characteristic $\neq 2$ the problem of finding a curve of genus g and simple covering of degree $g-1$ can be stated in the following way via fundamental groups, following [16]:

Let first k be an algebraically closed field and $f : \mathcal{C} \longrightarrow \mathbf{P}^1_k$ a covering of degree $g-1$. We then have $4g-4$ branch points; let these be $a_1, \ldots, a_{4g-4}$. We fix a yet different point b of $\mathbf{P}^1(k)$. Now, the étale fundamental group

$\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_k\}; b)$ operates transitively on the fiber $f^{-1}(b)$. The isomorphism class of the covering $f : \mathcal{C} \longrightarrow \mathbf{P}^1_k$ is uniquely given by the isomorphism class of this group theoretic operation, and conversely, any transitive operation of $\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_k\}; b)$ on a set of size $g-1$ gives such a covering.

We now omit the base point because it is not relevant for the following. In characteristic 0 we have

$$\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\}) \simeq \langle \gamma_1, \ldots, \gamma_{4g-4} \mid \gamma_1 \cdots \gamma_{4g-4} = 1 \rangle\widehat{\ },$$

where the γ_i are inertia elements. A simple covering is then given by a transitive operation of $\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$ on $\{1, \ldots, n\}$, where every γ_i operates as a transposition. Clearly this is possible, so there is such a covering.

For characteristic $p > 0$, we have the so-called *tame fundamental group* $\pi^t_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$, which operates on the fibers of tame coverings. This group is a quotient of $\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$. The images of the γ_i are again inertia elements, and now a simple covering is given by a transitive operation of this group on $\{1, \ldots, n\}$, where every γ_i operates as a transposition.

The pro-p-prime quotient of $\pi_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$ (and thus $\pi^t_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$) is isomorphic to the p-prime quotient of $\pi_1(\mathbf{P}^1_K - \{a_1, \ldots, a_{4g-4}\})$, where K is an algebraically closed field of characteristik 0. It follows that for $g \geq p$ there is a simple covering of degree $g-1$ ramified above $a_1, \ldots, a_{4g-4}$.

Above we stated the problem to find to prove that there are simple coverings of degree $g-1$ if $p < g$. This can now be formulated as a question on the tame fundamental groups $\pi^t_1(\mathbf{P}^1_k - \{a_1, \ldots, a_{4g-4}\})$ as stated. Unfortunately, this reformulation does not provide an answer to the question either.

We now introduce the condition that the pencils are obtained via singularities of plane models:

Lemma 3.15 There is an unique open subscheme U_3 of $\mathcal{Z}^1_{g,g-1}$ such that for any geometric point D of U_3, if s is the image of this point on $\mathcal{H}_g$, there is a decomposition $D = \tilde{D} + \Delta$, where $\tilde{D}$ has degree $g-3$ and Δ has degree 2, $|\omega(-\tilde{D})|$ is base point free and defines a birational plane model of $(\mathcal{Z}^1_{g,g-1})_s$ and Δ is the divisor contained in the divisor for a rational singularity of $(\mathcal{Z}^1_{g,g-1})_s$. In particular, $|\omega(-D)|$ and $|D|$ are pencils, and therefore $U_3 \subseteq (\mathcal{Z}^1_{g,g-1} - \mathcal{Z}^2_{g,g-1})$.

The scheme U_3 maps surjectively to $\mathbb{Z}$.

Proof. We consider the diagram

$$\begin{array}{ccccc} (\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2})^1 & \longrightarrow & \mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2} & \xrightarrow{\pi} & \mathcal{Z}_{g,g-3} \\ \downarrow & & \downarrow & & \\ \mathcal{Z}^1_{g,g-1} & \longrightarrow & \mathcal{Z}_{g,g-1} \ , & & \end{array}$$

where $(\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2})^1$ is the preimage of $\mathcal{Z}^1_{g,g-1}$ in $\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2}$ and π is the projection. The morphism $(\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2})^1 \longrightarrow \mathcal{Z}_{g,g-3}$ is proper because it is a composition of a closed and therefore proper morphism and a proper morphism.

In $\mathcal{Z}_{g,g-3}$ there is an open subscheme V whose geometric points correspond to divisors on curves over algebraically closed fields defining plane models with singularities of order 2. (In [22] we showed that V maps surjectively to $\mathcal{H}_g$, but we do not need this here.) Let U be the preimage of V in $(\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2})^1$. As a preimage of an open scheme it is open in $(\mathcal{Z}_{g,g-3} \times \mathcal{Z}_{g,2})^1$. Let U_3 be the image of U in $\mathcal{Z}^1_{g,g-1}$. As the projection to $\mathcal{Z}^1_{g,g-1}$ is flat, U_3 is open in $\mathcal{Z}^1_{g,g-1}$.

We already know by Lemma 3.4 that U_3 maps surjectively to $\mathrm{Spec}(\mathbb{Z})$. □

The space. Based on U_0, U_2 and U_3 as in the lemmata, we now define a scheme whose rational points correspond to divisors which can be obtained in the algorithm.

For U_2, note that the morphism $\mathcal{G}^1_{g-1}(\mathcal{C}) \longrightarrow \mathcal{W}^1_{g-1}(\mathcal{C})$ induces an isomorphism over $\mathcal{W}^1_{g-1}(\mathcal{C}) - \mathcal{W}^2_{g-1}(\mathcal{C})$. We restrict U_2 to this space; let the resulting space be U_2'. Now let U be the intersection of U_0 with the intersection of the images of U_2' and U_3 in $\mathcal{H}_g$. As the morphisms are smooth, U is open in $\mathcal{H}_g$. Let A be the complement of U in $\mathcal{H}_g$ and let B be the complement of of U_3 in $\mathcal{Z}^1_{g,g-1} - cZ^2_{g,g-1}$.

If we now apply [22, Corollary 6] and Proposition 3.9 to this setting, we obtain:

Lemma 3.16 Let $g \geq 5$.

a) There is a constant $C > 0$ such that for any prime power q with $q \geq g$ if $g > 7$,
$$\frac{\#A(\mathbb{F}_q)}{\#\mathcal{H}_g(\mathbb{F}_q)} \leq C \cdot \frac{1}{q} \ .$$

b) For curves corresponding to $\mathbb{F}_q$ points of U, there are $\sim q^{g-4}$ pencils corresponding to points in $U_2'(\mathbb{F}_q)$. These pencils define $\frac{1}{(g-1)!} \cdot q^{g-3}$ base point free completely split divisors of degree $g-1$.

c) There is a constant $C' > 0$ such that for any $\mathbb{F}_q$-rational point s of $\mathcal{H}_g$, the number of effective divisors given on $(\mathcal{H}_g)_s$ by points in $B_s(\mathbb{F}_q)$ is $\leq C' \cdot q^{g-3}$.

d) For curves corresponding to $\mathbb{F}_q$-rational points of U, there are $\sim q^{g-4}$ pencils corresponding to points in $U_3(\mathbb{F}_q)$. In these pencils there are in total $\sim \frac{1}{(g-1)!} \cdot q^{g-3}$ completely split divisors of degree $g-1$ which come from the plane model to a system $|K-D|$ and a singularity of order two on the curve.

Proof of Proposition 2.8. We have just proven part a) of the proposition.

For b) we just need to note that every plane model of degree $g+1$ has $\leq c := \frac{g(g-1)}{2} - g$ singularities. Thus the $\sim \frac{1}{(g-1)!} \cdot q^{g-3}$ divisors fulfilling the deired conditions come from at least $\sim \frac{1}{c \cdot (g-1)!} \cdot q^{g-3}$.

Item c) then follows.

4 Experiments

We implemented what we called, in subsection 2.2, the "explicit approach" in the computer algebra system Magma. For a comparison, we also implemented the approach based on plane models of degree $(g+1)$ as described in [6]. We note that a function based on these ideas is already available in Magma under the name IndexCalculus. Nevertheless, we also implemented a new version of this algorithm in order to vary different parameters like the size of the matrix of relations and the number of vertices in the graph of large prime relations.

Let us call the algorithm the algorithm based on [6] the "plane model based algorithm" and the new algorithm the "pencil based algorithm". We now briefly describe the specifications made for both algorithms. The plane model based algorithm proceeds as follows: First, a plane model of degree $(g+1)$ is computed and the instance discrete logarithm problem is transfered to this. Given such a plane model of $\mathcal{C}$ and two divisor classes in $\mathrm{Pic}^0_{\mathcal{C}}(\mathbb{F}_q)$, in a first step we apply the method from Algorithm 3 to generate q relations of type FP or PP. From this, we build the graph $\mathcal{G}$ of large prime relations using a factor base $\mathcal{F}$ of size $\lceil \kappa \cdot q^{1-\frac{1}{g-1}} \rceil$, where $\kappa := (4 \cdot (g-1)!)^{1/(g-1)}$ is chosen as indicated in [9]. So for $g = 4, 5$ the constant κ is approximately 2.8 and 3.1, respectively. We note that by [9] asymptotically any κ with $\kappa \geq (2 \cdot (g-1)!)^{1/(g-1)}$ should be sufficient. In a second step we use a breadth-first search to construct a shortest path tree $\mathcal{T}$ on the graph. We then create further relations factoring over $\mathcal{F} \cup \mathcal{V}$, where $\mathcal{V}$ is the vertex set of $\mathcal{T}$. These relations are generated in the same way as before, that is, by

intersecting the plane model with lines. Whenever an appropriate relation is created, we check if it has already been constructed before. This can easily be done using the aggregate "set" in Magma. Substituting elements of $\mathcal{V}$ with the help of $\mathcal{T}$ we generate a matrix of relations R with slightly more rows than columns. Then we apply the Lanczos algorithm to find a non-trivial row vector γ in the left kernel of R.

The pencil based algorithm on the other hand starts with the same input and a factor base $\mathcal{F}$ of size $\lceil \kappa \cdot q^{1-\frac{1}{g-2}} \rceil$. We want the running times for the relation generation and the linear algebra step to be similar so experimentally we decided to set $\kappa := 1$. In this case, we can not guarantee there are lines in $\mathbf{P}^2_{\mathbb{F}_q}$ factoring over $\mathcal{F}$ for any plane model of $\mathcal{C}$. So we increase $\mathcal{F}$ by up to $(g-3)$ new elements for each of the plane models used. Again, we always check if a relation has already been constructed. Furthermore, in order to decrease the number of duplicates during the relation generation process, for each pencil given by a singularity p of a plane model $\pi : \mathcal{C} \to \mathcal{C}_{pm}$ of degree $(g+1)$, we only consider lines through p and $\pi(Q)$ where $Q \in \{P_1, \ldots, P_{\lceil \frac{k}{2} \rceil}\}$ with $k = q^{1-\frac{1}{g-2}}$. As before, we create a graph $\mathcal{G}$ with q vertices in this way and consider the tree $\mathcal{T}$ constructed from $\mathcal{G}$ by breadth-first search.

Similarly to the plane model based algorithm we use the identical relation generation method as in the construction of $\mathcal{G}$ to generate a matrix of relations R' from T. Again we stop the construction of R' if it has slightly more columns than rows. As before we solve the corresponding system of equations by applying the Lanczos algorithm.

We intend to compare the plane model based algorithm for genus 4 or 5 to the pencil based algorithm for genus 5 or 6, respectively. For this we generated curves of the desired genera over $\mathbb{F}_3$, $\mathbb{F}_5$ and $\mathbb{F}_7$ by the Magma function RandomCurveByGenus and made a base change to the fields used in the tables below. We had to proceed this way so we could calculate the order N of the degree-0 Picard groups using the L-polynomial. We then picked the biggest prime p in the factorization of N and generated a random divisor class a of degree 0 on the corresponding curve which had order p. We set $b := n \cdot a$ where the integer n was chosen from $\{1, \ldots, p-1\}$ uniformly at random and computed the discrete logarithm of b with respect to a using the indicated algorithms.

The rounded results for varying fields can now be found in the tables below. T_{rel} and T_{la} stand for the time (in hours) needed to create R and solving $\gamma R = 0$, respectively.

$\mathbb{F}_{5^7} = \mathbb{F}_{78125}$	genus	size of $\mathcal{F}$	T_{rel}	T_{la}
Plane model	4	5000	0.3	0.2
Pencil	5	5000	0.6	0.4
Plane model	5	15000	2.0	3.2
Pencil	6	15000	5.2	6.1

$\mathbb{F}_{3^{11}} = \mathbb{F}_{177147}$	genus	size of $\mathcal{F}$	T_{rel}	T_{la}
Plane model	4	9000	1.0	1.2
Pencil	5	9000	2.0	2.1
Plane model	5	27000	7.7	12.5
Pencil	6	26000	20.9	23.3

$\mathbb{F}_{7^7} = \mathbb{F}_{823543}$	genus	size of $\mathcal{F}$	T_{rel}	T_{la}
Plane model	4	25000	8.2	10.6
Pencil	5	24000	11.6	17.6

$\mathbb{F}_{3^{13}} = \mathbb{F}_{1594323}$	genus	size of $\mathcal{F}$	T_{rel}	T_{la}
Plane model	4	39000	50.9	42.3
Pencil	5	38000	71.8	78.8

The theoretical results indicate that for $q \to \infty$, by applying the pencil based method instead of the plane model based method one can, at least for most curves, obtain an improvement which corresponds to the drop of the genus by one. The experiments show that this also holds from a practical point of view. The running times of the pencil based algorithm for genus $g = 5, 6$ differ from those of the plane model based algorithm for genus $(g-1)$ only by a factor between 1.5 and 2.0 for $g = 5$ and between 1.7 and 1.9 for $g = 6$. This difference comes from the way relations are generated. The plane model based algorithm only uses one plane model whereas the pencil based algorithm varies the plane model and maps the factor base and large primes back and forth between these models.

Changing the plane model various times also has an effect on $\mathcal{F}$. During the operation of the pencil based algorithm the original factor base of size $\lceil q^{1-\frac{1}{g-2}} \rceil$ is increased by a factor of about 2.8 for $g = 5$ and 3.1 for $g = 6$. This also indicates why we did not choose κ considerably smaller than 1, because this meant a bigger increase of $\mathcal{F}$ in each step. So the advantage of a smaller factor base at the beginning would be almost canceled during the process of relation generation. A bigger constant on the other hand would mean bigger matrices and hence the Lanczos algorithm would take considerably longer. It is also worth noting that the factor by which $\mathcal{F}$ is increased is similar to the one we chose for κ in the plane model based algorithm if the genus is

dropped by one. So the sizes of the sparse matrices R' generated by the pencil based algorithm are similar to those of the corresponding matrices R created by the plane model based algorithm for a genus that is decreased by 1.

On the other hand, we observed that matrices of type R' had 3.1 to 3.5 times the density of matrices of type R for which the genus is dropped by one. This can be explained by the more complicated way relations are generated in the pencil based algorithm.

Additionally, we note that the testing we did was primarily designed to check the behavior of the pencil based algorithm in comparison to the plane model based algorithm. In particular, we intended to show that the relation generation via different plane models and lines through singular points is practical. So far, this could only be tested in full for curves which are generated by base change. Nevertheless, we also did some testing of the relation generation step for fields of bigger characteristic and the corresponding results did not differ significantly from the ones above.

As expected, for q sufficiently large and $g = 5, 6$ in the considered cases the pencil based algorithm never failed. So first of all there always existed enough linear systems to create the graph of large prime relations. Secondly, the relations generated by different plane models in the way described above are sufficiently independent. However, in order to get matrices of appropriate rank we had to create them from slightly more relations than factor base elements. In fact, the Lanczos algorithm almost always succeeded with matrices of size $(\#\mathcal{F} + 10) \times \#\mathcal{F}$.

The main difficulties while implementing the pencil based algorithm were caused by the linear algebra step. The function ModularSolution specified for index calculus and relying on either structured Gaussian elimination or the Lanczos algorithm is available in Magma but failed to work reliably for large finite fields. The Gaussian elimination got extremely slow whenever the group order was prime, and often ModularSolution did not succeed independently from the chosen option. So we implemented the Lanczos algorithm based on [23] ourselves in Magma and in the C++ library LinBox. It turned out that Magma was about twice as fast multiplying dense vectors by sparse matrices and hence our implementation in Magma is considerably faster than the one in LinBox. So in order to get the results above we chose to use our Magma version of the Lanczos algorithm.

We also encountered failures in Magma in some other cases. Whenever there is a large set of tuples initiated, “for" loops are executed more slowly depending on the set’s size. However, when the set only consists of field elements or elements in an affine or projective space this was not the case. Hence we adjusted the implementation so that only sets of the above form are used to store the generated relations. Furthermore, we got an error

whenever we tried to check if a given divisor on a plane curve is principal. On the other hand, a similar function worked fine for divisors of function fields. So in contrast to the Magma function IndexCalculus we decided to represent the input curve by the function field of a corresponding plane model and not by the plane model itself. Finally, in certain cases, we had to delete the vertex set of a given graph before ending the function as otherwise we got an internal error.

5 Further ideas

We showed how one can, for curves of a fixed genus ≥ 5, efficiently use pencils to solve the discrete logarithm problem. The pencils we considered have index of speciality 2. It is then natural to ask if one can further improve the asymptotic running time by considering "even more special" pencils.

Indeed, Brill-Noether theory itself gives an explicit description of the spaces $\mathcal{C}_d^r$. We outline here how one might use this in principle.

We first give the essential theoretical observations we need. Let for this any smooth curve $\mathcal{C}$ of genus $g \geq 3$ over a field k be given, and let D be an effective divisor of degree d on $\mathcal{C}$. Then the Riemann-Roch Theorem states that

$$\dim(|D|) = \dim(\Gamma(\mathcal{C}, \omega(-D))) + d - g \ .$$

We have the short exact sequence

$$0 \longrightarrow \omega(-D) \longrightarrow \omega \longrightarrow \omega/\omega(-D) \longrightarrow 0$$

Now let $\kappa(D) := \Gamma(\mathcal{C}, \omega/\omega(-D))$. Then we have the exact sequence

$$0 \longrightarrow \Gamma(\mathcal{C}, \omega(-D)) \longrightarrow \Gamma(\mathcal{C}, \omega) \longrightarrow \kappa(D) \ .$$

Let n be the dimension of the image of $\Gamma(\mathcal{C}, \omega)$ in $\kappa(D)$. Then by Riemann-Roch

$$\dim(|D|) = d - n \ .$$

We make this more explicit now. Let $\omega_1, \ldots, \omega_g$ be a basis of $\Gamma(\mathcal{C}, \omega)$ and let $D = P_1 + \cdots + P_d$ with distinct k-rational points P_i. Then $\kappa(D) = \bigoplus_{i=1}^{d} \kappa(P_i)$ and $\kappa(P_i) = \omega_{P_i}/(\mathfrak{m}_{P_i}\omega_{P_i})$, which is a 1-dimensional k-vector space. Now n is the rank of the matrix $((\omega_i(P_j)))_{i,j}$, whose entries in the j^{th} column lie in $\kappa(P_j)$. Explicitly, we can fix one non-trivial differential ω, express every ω_j in the form $\omega_j = f_j \cdot \omega$ for a function f_j and consider the rank of the matrix $((f_i(P_j)))_{i,j}$. This matrix is called *Brill-Noether matrix* by Griffiths and Harris in [15]. Note that with $K := \operatorname{div}(\omega)$, we have an isomorphism $L(K) \longrightarrow \Gamma(\mathcal{C}, \omega)$, $f \mapsto f \cdot \omega$.

Let now d and r be given. Then $|D|$ has dimension at least r if and only if the rank of the matrix is $\leq d-r$. This means that the determinants of all minors of size $d-r+1$ of the matrix vanish.

So very briefly, we have the following method to compute such linear systems: We first compute an effective canonical divisor K and a basis $f_1, \ldots, f_g$ of $L(K)$. Then we regard the P_i as unknown points on the curve and express the condition that the determinants of all minors of size $d-r+1$ vanish by a non-linear system of equations in the coordinates of the P_i.

Note also that if $D = P_1 + \cdots + P_d$ is one solution, then so is any effective divisor which is linearly equivalent to D and which splits into distinct rational points. Because of this it is reasonable to fix $\rho + r$ of the points, where ρ is the Brill-Noether number defined above.

Concerning the systems of equations, there are in fact two obvious variants here: The first variant is that one starts off with a plane model. In this case, obviously one only has one curve equation for each point P_i but the determinants of the matrices are rather complicated. A second variant is to consider the canonical embedding of the curve. In this case, each point has g coordinates, and one has various equations for each point, but the functions $f_1, \ldots, f_g$ are now just the coordinates $x_1, \ldots, x_g$ and the determinants are as easy as possible.

In any case, for a fixed genus, degree and dimension, the number of variables and the number of equations is fixed and the degrees of the equations are bounded. It is therefore reasonable to expect that the computation can be performed in expected polynomial time. With a suitable variant of the algorithm it should be possible to prove this.

Now by Proposition 3.1, for any curve of genus g and for $d = \lceil \frac{g}{2} + \frac{3}{2} \rceil$ the space W_d^1 has dimension at least 1. This suggests that one might use these pencils to construct the graph of large prime relations.

This suggests that one might obtain an algorithm to compute discrete logarithms for nearly all curves of a fixed genus g with an expected running time of

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

The indicated index calculus algorithm seems however to be of little practical value. Indeed, it seems to be very difficult to perform any practical and non-trivial computation with the indicated method to find special divisors.

Given that we cannot perform any computation, we also have no way to see if experimentally the algorithm performs as one might expect heuristically. We therefore also do not want to make a claim that there is an algorithm with the indicated running time for nearly all curves.

References

[1] E. Arbarello, M. Cornalba, and P. Griffiths. *Geometry of Algebraic Curves II*. Springer-Verlag, 2011.

[2] E. Arbarello, M. Cornalba, P. Griffiths, and J. Harris. *Geometry of Algebraic Curves*. Springer-Verlag, 1985.

[3] F. Chung and L. Lu. The diameter of sparse random graphs. *Adv. in Appl. Math.*, 26:257–279, 2001.

[4] A.J. de Jong. The stacks project. http://stacks.math.columbia.edu, 2017.

[5] C. Diem. The GHS Attack in odd Characteristic. *J. Ramanujan Math. Soc.*, 18:1 – 32, 2003.

[6] C. Diem. An Index Calculus Algorithm for Plane Curves of Small Degree. In F. Hess, S. Pauli, and M. Pohst, editors, *Algorithmic Number Theory — ANTS VII*, LNCS 4076, pages 543 – 557, Berlin, 2006. Springer.

[7] C. Diem. On the discrete logarithm problem in class groups of curves. *Math.Comp.*, 80:443 – 475, 2011.

[8] C. Diem. On the notion of bit complexity. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS*, 103:35–52, 2011. In the "Complexity Column".

[9] C. Diem. On the discrete logarithm problem for plane curves. *Journal de Théorie des Nombres de Bordeaux*, 24:639–667, 2012.

[10] C. Diem and E. Thomé. Index calculus in class groups of non-hyperelliptic curves of genus three. *J. Cryptology*, 21:593–611, 2008.

[11] W. Fulton. Hurwitz schemes and irreducibility of moduli of algebraic curves. *Annals of Mathematics*, 90(3):542–575, 1969.

[12] W. Fulton and R. Lazarsfeld. On the connectedness of degeneracy loci and special divisors. *Acta Mathematica*, 146(1):271–283, 1981.

[13] P. Gaudry, E. Thomé, N. Thériault, and C. Diem. A double large prime variation for small genus hyperelliptic index calculus. *Math. Comp.*, 76:475–492, 2007.

[14] D. Gieseker. Stable curves and special divisors: Petri's conjecture. *Inventiones mathematicae*, 66(2):251–275, 1982.

[15] 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.

[16] A. Grothendieck. *Séminaire de Geometrie Algebrique 1960-61: Revêtements Etales et Groupe Fondamentale (SGA I).* Institut des Hautes Etudes Scientifiques, 1961.

[17] A. Grothendieck with J. Dieudonné. Eléments de Géométrie Algébrique (I-IV). ch. I: Springer-Verlag, Berlin, 1971; ch. II-IV: Publ. Math. Inst. Hautes Etud. Sci. 8,11, 17, 20,24, 28, 32, 1961-68.

[18] R. Hartshorne. *Algebraic Geometry.* Springer-Verlag, 1977.

[19] F. Heß. Computing Riemann-Roch spaces in algebraic function fields and related topics. *J. Symb. Comput.*, 11, 2001.

[20] G. Kempf. *Schubert Methods with an Application to Algebraic Curves.* Mathematisch Centrum, Amsterdam. Afdeling Zuivere Wiskunde [Publications] 71-6. 1971.

[21] S. Kleiman and D. Laksov. On the existence of special divisors. *American Journal of Mathematics*, 94(2):431–436, 1972.

[22] S. Kochinke. Ordinary plane models and completely split divisors. submitted.

[23] B. LaMacchia and A. Odlyzko. Solving large sparse linear systems over finite fields. In *Proceedings of the 10th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '90, pages 109–133, London, UK, 1991. Springer-Verlag.

[24] H. Martens. On the varieties of special divisors on a curve. *Journal für Reine und Angewandte Mathematik*, 227:111–120, 1967.

[25] J. Milne. Jacobian Varieties. In G. Cornell and J. Silverman, editors, *Arithemtic Geometry*, pages 167–212. Springer-Verlag, 1998.

[26] V.K. Murty and J. Scherk. Effective versions of the Chebotarev density theorem for funciton fields. *C. R. Acad. Sci.*, 319:523–528, 1994.

[27] H. Stichtenoth. *Algebraic Function Fields and Codes.* Springer-Verlag, 1993.

A Relative linear systems and relative effective divisors

For the lack of a suitable reference we discuss some foundational questions related to linear systems and divisors for smooth relative curves. We take the opportunity to give an exposition which is a bit more general than what is needed for the article.

A.1 Definition and basics

Let us first state some statements which are implicitly used in [1] without further elaboration:

Remark and Definition A.1 Let X be an S-scheme, and let $p : X \longrightarrow S$ be the structure morphism. Let $\mathcal{L}$ be a sheaf on X and $s \in S$. Moreover, let $X_s = p^{-1}(s)$ be the fiber over s and $\iota_s : X_s \longrightarrow X$ the inclusion. Then we have a natural homomorphism

$$(p_*\mathcal{L})_s = \lim_{s \in U \text{ open in } X} \Gamma(p^{-1}(U), \mathcal{L}) \longrightarrow \Gamma(X_s, \iota_s^*\mathcal{L})$$

Now $\Gamma(X_s, \iota_s^*\mathcal{L})$ is a $\kappa(s)$-module and the homomorphism induces a homomorphism

$$(p_*\mathcal{L})_s \otimes_{\mathcal{O}_s} \kappa(s) \longrightarrow \Gamma(X_s, \iota_s^*\mathcal{L})$$

Let now $\mathcal{H}$ be a subsheaf of $p_*\mathcal{L}$. Then we have a homomorphism

$$\mathcal{H}_s \otimes_{\mathcal{O}_s} \kappa(s) \longrightarrow \Gamma(X_s, \iota_s^*\mathcal{L}) \ .$$

In [1] this homomorphism is called *fiber homomorphism* and $\mathcal{H}_s \otimes_{\mathcal{O}_P} \kappa(s)$ is denoted by $\mathcal{H} \otimes \kappa(s)$, and $\iota_P^*\mathcal{L}$ is denoted by $\mathcal{L} \otimes \kappa(s)$. We do not follow the latter notation and rather write $\mathcal{L}_{|X_s}$.

We also recall the so-called *projection formula*, given in [17, I, (5.4.10)], which says:

Lemma A.2 Let $f : X \longrightarrow S$ be a morphism of ringed spaces, let $\mathcal{Q}$ be a locally free sheaf of finite rank on S and let $\mathcal{F}$ be any sheaf on X. Then there is a natural isomorphism

$$f_*(\mathcal{F}) \otimes_{\mathcal{O}_S} \mathcal{Q} \simeq f_*(\mathcal{F} \otimes_{\mathcal{O}_X} f^*\mathcal{Q}) \ .$$

Recall that a smooth (relative) curve over a scheme S is a proper S-scheme whose fibers are all smooth curves of the same genus (which itself are by definition geometrically reduced and geometrically irreducible). We then have:

Lemma A.3 Let $\mathcal{C}/S$ be a smooth relative curve with structure morphism $p : \mathcal{C} \longrightarrow S$.

a) For an invertible sheaf $\mathcal{Q}$ on S, the canonical morphism $\mathcal{Q} \longrightarrow p_*p^*\mathcal{Q}$ is an isomorphism.

b) For an invertible sheaf $\mathcal{L}$ on $\mathcal{C}$, $p_*\mathcal{L}$ is locally free.

Proof. On a). The question being local on S, we can assume that $\mathcal{Q}$ is free and then that $\mathcal{Q} = \mathcal{O}_S$. Now $p_*p^*\mathcal{O}_S = p_*\mathcal{O}_\mathcal{C} = \mathcal{O}_S$, the last equation as $\mathcal{C}/S$ is proper with geometrically connected fibers (Stein factorization).

On b). For this statement see [1, Chapter 21, §3]. □

Let now S be a locally noetherian scheme and let $\mathcal{C}/S$ be a smooth relative curve with structure morphism $p : \mathcal{C} \longrightarrow S$. The following definition on "relative linear systems" is inspired by [1, Chapter 21, Definition 3.12]. The underlying mathematical ideas are identical, but we changed the terminology in two ways: First, we emphasize the relative point of view; this is a minor, linguistic change. Second, we define these systems via isomorphism classes in such a way that in the classical case (i.e. for curves over fields), the objects are in natural bijection with what is classically called "linear system". Third, we again change the terminology in order to be able to distinguish this sheaf theoretic description of linear systems and a more classical description with spaces of divisors.

Definition A.4 A *representative of a relative linear system of degree d and dimension r* is a pair $(\mathcal{L}, \mathcal{H})$, where $\mathcal{L}$ is an invertible sheaf on $\mathcal{C}$ whose restriction to each fiber of $p : \mathcal{C} \longrightarrow S$ has degree d, and $\mathcal{H}$ is a locally free subscheaf of $p_*\mathcal{L}$ of rank $r + 1$ such that, for all $s \in S$, the fiber homomorphism

$$\mathcal{H} \otimes \kappa(s) \longrightarrow \Gamma(X_s, \mathcal{L}_{|\mathcal{C}_s})$$

is injective. Two such tuples $(\mathcal{L}, \mathcal{H})$, $(\mathcal{L}', \mathcal{H}')$ are *equivlent* if there is an invertible sheaf $\mathcal{Q}$ on S and an isomorphism $\mathcal{L}' \xrightarrow{\sim} \mathcal{L} \otimes p^*\mathcal{Q}$ inducing an isomorphism (via p_* and Lemma A.2) an isomorphism $\mathcal{H}' \xrightarrow{\sim} \mathcal{H} \otimes \mathcal{Q}$. A *sheaf theoretic relative linear system of degree d and dimension r* is an equivalence class of such tuples. If $(\mathcal{L}, \mathcal{H})$ is a representative of a relative linear system, the corresponding linear system, that is, equivalence class of $(\mathcal{L}, \mathcal{H})$, is denoted by $[(\mathcal{L}, \mathcal{H})]$.

Theorem 3.13 in [1, Chapter 21] states: If $\mathcal{C} \longrightarrow S$ is a smooth curve of genus > 1 which has a section then the functor assigning to T the set of relative linear systems of degree d and dimension r on $\mathcal{C}_T$ is representable. The definition of $\mathcal{G}^r_d(\mathcal{C})$ under the condition that $\mathcal{C} \longrightarrow S$ has a section is that it is a representing object of this functor. If it does not have a section, $\mathcal{G}^r_d(\mathcal{C})$ becomes a representing object of such a functor after an étale base change. One can then say that $\mathcal{G}^r_d(\mathcal{C})$ represents the sheaf associated to the functor in the étale topology, similar to the corresponding statement for the Jacobian; cf. [25].

A.2 The base locus

For curves over fields, base point free linear systems of dimension r correspond to morphisms to $\mathbf{P}^r$ up to automorphism. Moreover, for curves over algebraically closed fields, a base point free linear system given by a tuple $(\mathcal{L}, v)$ on a curve $\mathcal{C}$ defines a morphism $\mathcal{C} \longrightarrow \mathbf{P}(V^\vee)$ given by $P \mapsto \{g \in V \mid g(P) = 0\}$.

We now discuss how the base locus should be defined in order that one again obtains such a result.

Before we continue we would like to clarify an aspect of locally free sheafs and projective vector bundles: A. Grothendieck showed in [17, II, §4] how one can associate to any coherent sheaf $\mathcal{E}$ on a scheme X a space $\mathbf{P}(\mathcal{E})$ over X whose fibers are projective spaces. Applied to a locally free sheaf one obtains then a projective vector bundle. As a special case of a vector space V over a field k one obtains a projective space with a coordinate ring whose homogeneous elements of degree 1 are the elements from V. This means that the thus defined space $\mathbf{P}(V)(k)$ is then canonically isomorphic to the classical space $\mathbf{P}^\vee(V) \simeq \mathbf{P}(V^\vee)$. The reader should keep this in mind in the following.

We have by [18, Proposition 7.12]:

Proposition A.5 Let $\mathcal{E}$ be a locally free coherent sheaf on S and let $\mathcal{L}$ be an invertible sheaf on $\mathcal{C}$. Then morphisms $f : \mathcal{C} \longrightarrow \mathbf{P}(\mathcal{E})$ over S correspond to isomorphism classes of tuples $(\mathcal{L}, c)$ of invertible sheaves $\mathcal{L}$ and surjections $c : p^*\mathcal{E} \longrightarrow \mathcal{L}$ with respect to the following notion of isomorphism:

An isomorphism from $(\mathcal{L}, c)$ to $(\mathcal{L}', c')$ is an isomorpism $\psi : \mathcal{L} \longrightarrow \mathcal{L}'$ with $\psi \circ c = c'$.

Note furthermore that morphisms $p^*\mathcal{E} \longrightarrow \mathcal{L}$ correspond to morphisms $\mathcal{E} \longrightarrow p_*\mathcal{L}$; cf. the beginning of [18, Chapter II, §5]. This inspires the following definition.

Definition A.6

a) The *set theoretic base locus* of $(\mathcal{L}, \mathcal{H})$ is the set of points $P \in \mathcal{C}$ such that $\mathcal{H}_P \longrightarrow \mathcal{L}_P$ is not surjective.

b) $(\mathcal{L}, \mathcal{H})$ is *base point free* if $p^*\mathcal{H} \longrightarrow \mathcal{L}$ is surjective.

These definitions are invariant under isomorphism; we can therefore also say that a relative linear system is base point free etc.

Lemma A.7 Let $(\mathcal{L}, \mathcal{H})$ be a representative of a relative linear system on $\mathcal{C}/S$.

a) The set theoretic base locus of $(\mathcal{L}, \mathcal{H})$ is closed in $\mathcal{C}$.

b) For $P \in \mathcal{C}$ lying over $s \in S$, P is a base point of $(\mathcal{L}, \mathcal{H})$ on $\mathcal{C}$ if and only if P is a base point of the system $(\mathcal{L}_{|\mathcal{C}_s}, \mathcal{H} \otimes \kappa(s))$ obtained by restricting $(\mathcal{L}, \mathcal{H})$ to the fiber $\mathcal{C}_s$ over s

c) The formation of the set theoretic base locus commutes with base change.

d) The set of $s \in S$ such that the system $(\mathcal{L}_{|\mathcal{C}_s}, \mathcal{H} \otimes \kappa(s))$ obtained by restricting $(\mathcal{L}, \mathcal{H})$ to the fiber $\mathcal{C}_s$ over s is base point free is closed.

Proof. As $\mathcal{L}$ is locally generated by a single element, a) holds.

Let P and s be as in b). The statement is that $p^* : \mathcal{H}_P \longrightarrow \mathcal{L}_P$ is surjective if and only if $p_s^* : (\mathcal{H} \otimes \kappa(s))_P \longrightarrow (\mathcal{L}_{|\mathcal{C}_s})_P$ is surjective.

Now, the former can be restated by saying that $(\mathcal{L}/p^*\mathcal{H})_P = 0$ and the latter by saying that $(\mathcal{L}/p^*\mathcal{H})_P \otimes_{\mathcal{O}_{S,s}} \kappa(s) = 0$. It is clear that the former implies the latter. Conversely, the latter implies that $(\mathcal{L}/p^*\mathcal{H})_P \otimes_{\mathcal{O}_{\mathcal{C},P}} \kappa(P) = 0$. Now, Nakayama's lemma implies the former condition.

The statement in c) follows immediately. For the statemederwestenhernederwestenhernederwestenhernent in d) we consider the image of the base locus in S. The image of this closed set is closed because $\mathcal{C}$ is proper over S. □

A.3 Functorial divisor theoretic relative linear systems

We recall some basic statements on divisors.

Remark A.8 The following statements in a),b),d) and e) are explained in [25, Section 3]; see also [18, II, Section 6] and [17, IV, §21]. The statement in c) is immediate.

Let X be any scheme.

a) One can define an *effective divisor* on X as a closed subscheme on X given locally by a single element ("equation") which is not a zero-divisor; we follows this definition.

b) There is a canonical bijection between effective divisors and the isomorphism classes of tuples $(\mathcal{L}, g)$, where $\mathcal{L}$ is an invertible sheaf and $g \in \Gamma(X, \mathcal{L})$ which are non-zero divisors, that is, which induces an inclusion $\mathcal{O}_X \hookrightarrow \mathcal{L}$. It is given by $D \mapsto [(\mathcal{O}(D), 1)]$ with $\mathcal{O}(D) := \mathcal{I}(D)^{-1}$. Conversely, to $(\mathcal{L}, g)$ one associates (via the inclusion $\mathcal{O}_X \longrightarrow \mathcal{L}$) the sheaf $\mathcal{L}^{-1}$ as an ideal seaf of $\mathcal{O}_X$, and to this the corresponding closed subscheme of X. This subscheme is then called the *divisor of zeros* of g.

c) Let $\mathcal{L}$ be an invertible sheaf on X. Then the effective divisors D on X with $\mathcal{O}(D) \approx \mathcal{L}$ correspond to equivalence classes of elements g as in b) with respect to multiplication by elements of $\Gamma(X, \mathcal{O}_X^*)$.

d) Let now X be an S-scheme. Then a *relative effective divisor* on X/S is by definition an effective divisor on X which is flat over S.

e) Let X be flat over S. Then there is a bijection between the relative effective divisors on X/S and the isomorphism classes $(\mathcal{L}, g)$ as in b) such that for all $s \in S$, $g|_{X_s} \in \mathcal{L}_{|X_s}$ induces an inclusion $\mathcal{O}_{X_s} \hookrightarrow \mathcal{L}_{|X_s}$.

The following lemma contains the key statement for the study of relative effective divisors for a possibly non-reduced base. (For a reduced base it is immediate.)

Lemma A.9 Let S be locally noetherian and let X be a flat S-scheme. Let $\mathcal{L}$ be an invertible sheaf on X and $g \in \Gamma(\mathcal{C}, \mathcal{L})$ such that for all $s \in S$, $g|_{X_s} \in \mathcal{L}_{|X_s}$ induces an inclusion $\mathcal{O}_{X_s} \hookrightarrow \mathcal{L}_{|X_s}$. Then $g|_{X_s} \in \mathcal{L}_{|X_s}$ induces in inclusion $\mathcal{O}_X \hookrightarrow \mathcal{L}$.

Proof. The proof relies on associated points; for the necessary background information we refer to [4, Chapter 30,§1 – §3]. Note first that (essentially by definition), the associated points of $\mathcal{L}$ are equal to the associated points ot X. By [4, Chapter 30, Lemma 2.10] we have to check $g_P \in \mathcal{L}_P$ is a non-zero divisor for all associated points P of X. Note that as $\mathcal{L}_P \approx \mathcal{O}_{X,P}$, g_P is a non-zero-divisor of $\mathcal{L}_P$ if and only if it is not contained in $\mathfrak{m}_{\mathcal{L},P}$, that is, if $g(P) \in \mathcal{L}_P/\mathfrak{m}_P$ is non-trivial.

By [4, Chapter 30, Lemma 3.1] the associated points of X are the associated points of the fibers defined by associated points of S.

Now, by assumption, for every point $s \in S$ and every associated point P of X_s, $g_{|X_s}$ is non-zero divisor of $(\mathcal{L}_{|X_s})_P$. This in turn can be expressed by saying that $g(P) \in \mathcal{L}_P/\mathfrak{m}_P$ is non-trivial. □

We now apply this to smooth relative curves over a locally noetherian base. So let for this S be locally noetherian and $\mathcal{C}/S$ a relative curve.

Lemma A.10 Let $\mathcal{L}$ be an invertible sheaf on $\mathcal{C}$ and $g \in \Gamma(\mathcal{C}, \mathcal{L}) = \Gamma(S, p_*\mathcal{L})$. Then g satisfies the condition in item f) of the remark if and only if for all $s \in S$, $g(s) \in p_*\mathcal{L} \otimes \kappa(s)$ is non-trivial. In this case g also defines an inclusion $\mathcal{O}_S \hookrightarrow p_*\mathcal{L}$.

Proof. We have $p_*\mathcal{L} \otimes \kappa(s) \simeq \Gamma(\mathcal{C}_s, \mathcal{L}_{|\mathcal{C}_s})$ for all $s \in S$ by "cohomology and base change" as explained at the beginning of [1, Chapter 21, §3]. For $g \in \Gamma(\mathcal{C}, \mathcal{L})$, here $g(s)$ corresponds to $g_{|\mathcal{C}_s}$. The result then follows from the previous lemma. □

Remark A.11 Let $\mathcal{L}$ be an invertible sheaf on $\mathcal{C}$. How can then the divisors defining the the class of $\mathcal{L}$ in $\mathrm{Pic}(\mathcal{C})$ be described?

Let for this $(\mathcal{L}', g)$ be a tuple defining a relative effective divisor on $\mathcal{C}/S$. Then $\mathcal{L}' \approx \mathcal{L} \otimes p^*\mathcal{Q}$ for an invertible sheaf $\mathcal{Q}$. Equivalently, $p_*(\mathcal{L}' \otimes \mathcal{L}^{-1})$ is an invertible sheaf, and this sheaf is then a sheaf $\mathcal{Q}$ as first stated. Such a tuple $(\mathcal{L}', g)$ then gives rise to an inclusion $\mathcal{O}_\mathcal{C} \hookrightarrow p_*\mathcal{L} \otimes \mathcal{Q}$ which induces inclusions in the fibers. This in turn corresponds to an inclusion $\mathcal{Q}^{-1} \hookrightarrow \mathcal{L}$ which induces inclusions in the fibers.

The result is that relative effective divisors on $\mathcal{C}/S$ defining the class of $\mathcal{L}$ in $\mathrm{Pic}(\mathcal{C})$ correspond to 1-dimensional subsheaves of $\mathcal{L}$ which induce inclusions in the fibers.

One can then say that the relative effective divisors on $\mathcal{C}/S$ correspond to the 0-dimensional relative linear systems as defined above. This shows that we are on the right track here.

We now have that relative effective divisors on $\mathcal{C}/S$ correspond to:

1. Classes of tuples $(\mathcal{L}, \mathcal{K})$, where $\mathcal{L}$ is an invertible sheaf on $\mathcal{C}$ and $\mathcal{K}$ is a 1-dimensional *free* subsheaf of $p_*\mathcal{L}$ which induces fiberwise inclusions, where two tuples $(\mathcal{L}, \mathcal{K})$ and $(\mathcal{L}', \mathcal{K}')$ are equivalent if there is an isomorphism $\mathcal{L} \longrightarrow \mathcal{L}'$ mapping (via p_*) $\mathcal{K}$ to $\mathcal{K}'$.

2. Classes of tuples $(\mathcal{L}, \mathcal{K})$, where $\mathcal{L}$ is an invertible sheaf on $\mathcal{C}$ and $\mathcal{K}$ is a 1-dimensional *locally free* subsheaf of $p_*\mathcal{L}$ which induces fiberwise inclusions, where two tuples $(\mathcal{L}, \mathcal{K})$ to $(\mathcal{L}', \mathcal{K}')$ are equivalent if there is an invertible sheaf $\mathcal{Q}$ on S and an isomorphism $\mathcal{L} \longrightarrow \mathcal{L}' \otimes p^*\mathcal{Q}$ mapping (via p_* and Lemma A.2) $\mathcal{K}$ to $\mathcal{K}'$.

The reader might want to check directly that there is a bijection between the classes in 1. and in 2.

This inspires:

Definition A.12 Let $(\mathcal{L}, \mathcal{H})$ be a representative of a relative linear system on $\mathcal{C}/S$. We call a relative effective divisor D on $\mathcal{C}/S$ given by a 1-dimensional locally free subsheaf of $(\mathcal{L}, \mathcal{H})$ a divisor *defined by* $(\mathcal{L}, \mathcal{H})$ (and also by $[(\mathcal{L}, \mathcal{H})]$).

Lemma A.13 Let $(\mathcal{C}, \mathcal{H})$ be a representative of a linear system on $\mathcal{C}/S$. Then relative effective divisors on $\mathcal{C}$ defined by $(\mathcal{L}, \mathcal{H})$ correspond to the sections of the projective vector bundle $\mathbf{P}(\mathcal{H}^\vee)$ over S.

Proof. A 1-dimensional locally free subsheaf of $\mathcal{L}$ corresponds to (or is depending on the definition) an equivalence class of inclusions $\iota : \mathcal{K} \hookrightarrow p_*\mathcal{L}$, where two such data $(\mathcal{K}, \iota), (\mathcal{K}', \iota')$ are equivalent if there is an isomorphism $\mathcal{K} \longrightarrow \mathcal{K}'$ respecting the inclusions.

By dualizing, subsheaves as in the lemma correspond to isomorphism classes of homomorphisms $\pi : \mathcal{H}^\vee \longrightarrow \mathcal{M}$ for which the morphisms $\mathcal{H}^\vee \otimes \kappa(s) \longrightarrow \mathcal{M}^\vee \otimes \kappa(s)$ are surjective. Here two such data $(\mathcal{M}, \pi), (\mathcal{M}', \pi')$ are equivalent if there is an isomorphism $\mathcal{M} \longrightarrow \mathcal{M}'$ respecting the surjections.

By Nakayama's lemma, the condition on such a datum $(\mathcal{M}, \pi)$ means that $\mathcal{H}^\vee \longrightarrow \mathcal{M}$ is surjective.

Thus the equivalence classes of g's correspond to the equivalence classes of surjections $\mathcal{H}^\vee \longrightarrow \mathcal{M}$. These in turn correspond to the sections of $\mathbf{P}(\mathcal{H}^\vee)$ over S. □

Definition A.14 We define the *functorial divisor theoretic linear system* associated to a representative of a linear system $(\mathcal{L}, \mathcal{H})$ on $\mathcal{C}/S$ as the following functor: To a locally noetherian S-scheme T we associate the set of relative effective divisors on $\mathcal{C}_T/T$ given by 1-dimensional locally free subsheaves of $\mathcal{H}$ (inducing fiberwise inclusions).

Remark A.15 We refrained from calling the set of relative effective divisors on $\mathcal{C}/S$ a linear system because this set might be unreasonably small, even empty.

On the other hand, a sheaf theoretic linear system on $\mathcal{C}/S$ is uniquely defined by the associated functorial divisor theoretic linear system. Let for this $(\mathcal{L}, \mathcal{H})$ be a sheaf theoretic linear system on $\mathcal{C}/S$.

First, it suffices to show this locally in the étale topology. We can therefore assume that $\mathcal{H}$ is free. This means in particular that there are 1-dimensional free subsheaves of $\mathcal{H}$. An implication of this is that there is an effective divisor D on $\mathcal{C}/S$. Let us fix such a divisor. We then have $\mathcal{L} \approx \mathcal{C}(D)$. Now we have to recover a space corresponding to $\mathcal{H}$ under such an isomorphism to $\mathcal{C}(D)$. For this we consider all relative effective divisors D on $\mathcal{C}/S$ whose image in the Picard group is defined by $\mathcal{L}$. These correspond now to 1-dimensional subsheaves of $p_*\mathcal{O}_\mathcal{C}(D)$. Let $\mathcal{H}'$ be the sheaf generated by all these subsheaves. Then under an isomorphism $\mathcal{O}(D) \approx \mathcal{L}$, $\mathcal{H}$ corresponds to $\mathcal{H}'$. Therefore $(\mathcal{O}(D), \mathcal{H}')$ is also a representative of the linear system defined by $(\mathcal{C}, \mathcal{H})$.

Let $(\mathcal{L}, \mathcal{H})$ have degree d. We then have an obvious homomorphism (natural transformation) from this functor to the functor assigning to an S-scheme T the set of relative effective divisors of degree d on $\mathcal{C}_T/T$.

The following lemma is now immediate.

Lemma A.16 Let $(\mathcal{L}, \mathcal{H})$ be a representative of a linear system of degree d on $\mathcal{C}/S$.

a) The functorial divisor theoretic linear system is represented by the projective space bundle $\mathbf{P}(\mathcal{H}^\vee)$ over S.

b) The natural transformation just mentioned corresponds to an S-morphism $\mathbf{P}(\mathcal{H}^\vee) \longrightarrow \mathcal{C}_d$.

We now consider the functor which assigns to a locally noetherian S-scheme T the set of equivalence classes of triples $(\mathcal{L}, \mathcal{H}, D)$, where $(\mathcal{L}, \mathcal{H})$ is a representative of a relative linear system of degree d and dimension r on $\mathcal{C}_T$ and D a relative effective divisor on $\mathcal{C}_T/T$ given by a 1-dimensional locally free subsheaf of $\mathcal{H}$.

Let $(\mathcal{L}_u, \mathcal{H}_u)$ be the universal relative linear system of degree d. We then have, by applying the lemma to this system:

Proposition A.17 If $\mathcal{C}/S$ has a section, the functor just described is represented by the bundle $\mathbf{P}(\mathcal{H}_u^\vee)$ over $\mathcal{G}_d^r(\mathcal{C})$. The (functorial) projection then corresponds to the assignment $(\mathcal{L}, \mathcal{H}, D) \mapsto (\mathcal{L}, \mathcal{H})$ and the (functorial) projection $(\mathcal{L}, \mathcal{H}, D) \mapsto D$ corresponds to an S-morphism $\mathbf{P}(\mathcal{H}_u^\vee) \longrightarrow \mathcal{C}_d$.

In full generality, $\mathcal{C}/S$ represents the associated sheaf of the functor in the étale topology. There is still an S-morphism $\mathbf{P}(\mathcal{H}_u^\vee) \longrightarrow \mathcal{C}_d$ which corresponds to the given functorial projection locally in the étale topology.