An Algorithm for computing Isomorphisms of Algebraic Function Fields

F. Hess

Technical University of Berlin, Faculty II,
Institute of Mathematics, Secr. MA8-1,
Straße des 17. Juni 136, 10623 Berlin, Germany

Abstract. We develop an algorithm for computing isomorphisms and automorphisms of algebraic function fields of transcendence degree one in characteristic zero and positive characteristic.

1 Introduction

Let F_1/k and F_2/k denote algebraic function fields of transcendence degree one. An isomorphism ϕ of F_1/k and F_2/k is an isomorphism of fields $\phi : F_1 \to F_2$ whose restriction to k is the identity map. The main objective of this paper is to develop an algorithm to compute one or all isomorphisms of F_1/k and F_2/k if these algebraic function fields are isomorphic and have genus greater than one. We think of F_1/k and F_2/k as the function fields of some suitable, explicitly given and not necessarily non-singular irreducible algebraic curves such as plane curves defined over k, and describe ϕ by its action on the corresponding coordinates or field generators. For the special case $F = F_1 = F_2$ the algorithm computes the elements of the automorphism group $\mathrm{Aut}_k(F)$ of F/k. We restrict to genus greater than one since otherwise the number of isomorphisms or the automorphism groups may be infinite and the task is related to finding unique models by computing rational points on curves over k, a hard and from the techniques of this paper quite different problem.

Assume that F_1 and F_2 are given as finite and separable extensions of an algebraic function field K/k of transcendence degree one. Since F_1 and F_2 can be obtained by adjoining roots of monic, irreducible and separable polynomials $f_1, f_2 \in K[t]$ to K respectively the computation of isomorphisms ψ of F_1/k and F_2/k which restrict to the identity on K can essentially be reduced to the computation of the roots of f_1 in F_2 or of f_2 in F_1. This latter task can in turn be solved by special Hensel lifting and reconstruction techniques, or by general polynomial factorization algorithms. A concrete algorithm for computing such isomorphisms of F_1 and F_2 and for computing the elements of $\mathrm{Aut}_K(F)$ for finite extensions F_1/K, F_2/K, F/K and $K = \mathbb{Q}(x)$ is given in [6]. The computation of $\mathrm{Aut}_k(F)$ for example cannot be done this way since the fixed field of $\mathrm{Aut}_k(F)$, which would be a candidate for K, is a priori not known.

In the following we need to define and compare various entities for the function fields F_1 and F_2. To facilitate this we use the subscripts (1), (2) for either

of the two cases and (α) for both (hence $\alpha \in \{1, 2\}$). Entities with no subscript are related to both or completely independent of the two fields. We are thus considering the function fields $F_{(1)}$, $F_{(2)}$ and the isomorphisms ϕ.

Throughout the paper $F_{(\alpha)}/k$ denotes an algebraic function fields of transcendence degree one and genus greater than one. Unless otherwise stated, k is assumed to be the exact constant field of $F_{(\alpha)}/k$ (ie. algebraically closed in $F_{(\alpha)}$) and perfect. Our algorithms rely on the ability to compute in $F_{(\alpha)}/k$ as a field and $k(x_{(\alpha)})$-vector space for $x_{(\alpha)}$ a separating element of $F_{(\alpha)}/k$, and to compute with places, divisors and Riemann-Roch spaces $\mathcal{L}(D_{(\alpha)}) = \{a \in F_{(\alpha)}^{\times} \,|\, (a) + D_{(\alpha)} \geq 0\} \cup \{0\}$ for divisors $D_{(\alpha)}$. In particular an algorithm for computing Weierstrass places is required. We refer to [3, 4] for algorithmic aspects and to [8] for the theoretical background. Implementations of these algorithms are available in Kash [5] and Magma [1, 2].

2 The basic idea

Assume that ϕ is an isomorphism of $F_{(1)}/k$ and $F_{(2)}/k$. We derive a number of necessary conditions for the existence of ϕ, which altogether will also be sufficient. This also leads to a test for $F_{(1)}/k$ and $F_{(2)}/k$ being not isomorphic.

First, the genus of $F_{(1)}/k$ and $F_{(2)}/k$ is equal and ϕ maps places $P_{(1)}$ of $F_{(1)}/k$ to places $P_{(2)}$ of $F_{(2)}/k$, preserving the degree and gap numbers and inducing k-linear isomorphisms of the vector spaces $\mathcal{L}(nP_{(1)})$ and $\mathcal{L}(nP_{(2)})$. In particular, places of $F_{(1)}/k$ of smallest degree are mapped to places of $F_{(2)}/k$ of smallest degree and Weierstrass places of $F_{(1)}/k$ are mapped to Weierstrass places of $F_{(2)}/k$. For global function fields (that is k a finite field) there are only finitely many places of any fixed degree, and in general there are only finitely many Weierstrass places. This restricts the possibilities $P_{(2)} = \phi(P_{(1)})$ to a finite number. We remark that both sets of places (smallest degree, Weierstrass) can be computed efficiently.

Let $P_{(1)}$ and $P_{(2)}$ be places of degree one with $P_{(2)} = \phi(P_{(1)})$. Let g be the genus of $F_{(1)}$ and $F_{(2)}$ and let $n_i \in \mathbb{Z}^{>0}$ for $1 \leq i \leq r \leq g + 1$ be the first successive pole numbers greater than zero at $P_{(1)}$, equal to those at $P_{(2)}$, such that $\gcd\{n_i \,|\, 1 \leq i \leq r\} = 1$. Let $x_{(\alpha),i} \in \mathcal{L}(n_i P_{(\alpha)})$ be elements with $v_{P_{(\alpha)}}(x_{(\alpha),i}) = -n_i$. We define $x_{(\alpha),0} = 1$ and $n_0 = 0$. Obviously, $x_{(\alpha),i}$ is uniquely defined up to multiplication by elements from $k^{\times}$ and addition by k-linear combinations of the $x_{(\alpha),j}$ for $j < i$. As a result, $\phi(x_{(1),i}) = \sum_{j=0}^{i} \mu_{i,j} x_{(2),j}$ for suitable $\mu_{i,j} \in k$ and $\mu_{i,i} \neq 0$.

These observations are interesting because $F_{(\alpha)} = k(x_{(\alpha),1}, \ldots, x_{(\alpha),n_1})$, so ϕ is completely defined by the $\mu_{i,j}$. This equality holds because $P_{(\alpha)}$ is fully ramified and unramified over $k(x_{(\alpha),1}, \ldots, x_{(\alpha),n_1})$ due to the gcd condition. We let $I_{(\alpha)}$ be the kernel of the substitution homomorphism $k[t_1, \ldots, t_{n_1}] \to F_{(\alpha)}$, $t_i \mapsto x_{(\alpha),i}$. Thus generators of $I_{(1)}$ and $I_{(2)}$ define irreducible affine curves whose function fields are equal to $F_{(1)}$ and $F_{(2)}$ respectively. Since ϕ is an isomorphism, substituting $\sum_{j=0}^{i} \mu_{i,j} t_j$ for t_i yields an automorphism of $k[t_1, \ldots, t_{n_1}]$ which accordingly maps $I_{(1)}$ to $I_{(2)}$.

Turning things around assume now that the $\mu_{i,j}$ are indeterminates and $R = k[\{\mu_{i,j}|j \leq i\}]$. Then $I_{(1)}$ is also an ideal of $R[t_1, \ldots, t_{n_1}]$ and we define $I'_{(2)}$ to be the ideal of $R[t_1, \ldots, t_{n_1}]$ obtained from $I_{(1)}$ by substituting $\sum_{j=0}^{i} \mu_{i,j} t_j$ for t_i. Any solution $\mu_{i,j} \in k$ with $\mu_{i,i} \neq 0$ such that $I'_{(2)} = I_{(2)}$ yields an isomorphism of the coordinate rings $k[t_1, \ldots, t_{n_1}]/I_{(1)} \to k[t_1, \ldots, t_{n_1}]/I_{(2)}$ and thus an isomorphism ϕ of $F_{(1)}/k$ and $F_{(2)}/k$. It is actually sufficient to require only $I'_{(2)} \subseteq I_{(2)}$.

Summing up, our algorithm for computing one or all isomorphisms of $F_{(1)}/k$ and $F_{(2)}/k$ proceeds as follows. We first choose a suitable place of smallest degree or Weierstrass place $P_{(1)}$ of $F_{(1)}/k$ and compute all places $P_{(2)}$ which could correspond to $P_{(1)}$ under an isomorphism, respecting the condition of smallest degrees or Weierstrass with same gap sequences. Then we compute $I_{(1)}$, $I_{(2)}$ and $I'_{(2)}$ and solve for the $\mu_{i,j}$ with $\mu_{i,i} \neq 0$ such that $I'_{(2)} \subseteq I_{(2)}$. This gives all possible isomorphisms ϕ of $F_{(1)}/k$ and $F_{(2)}/k$.

This concludes the description of the basic idea. Of course, there are various problems to overcome. For one we need to find a convenient way to compute $I_{(1)}$ and $I_{(2)}$ and to check $I'_{(2)} \subseteq I_{(2)}$ in terms of the $\mu_{i,j}$. Another problem is that the number of the $\mu_{i,j}$ is roughly $g^2/2$ which makes the computation of $I'_{(2)}$ and finding a solution to $I'_{(2)} \subseteq I_{(2)}$ hard. To this end we essentially reduce the $\mu_{i,j}$ to two parameters. These issues are addressed in the following sections.

We finally remark that another way of computing isomorphisms is by using canonical curves. One advantage of this would be that the auxiliary places $P_{(1)}$ and $P_{(2)}$ could be avoided. Since these curves are uniquely determined up to linear transformation we could compare a canonical curve for $F_{(1)}/k$ with a generically linearly transformed canonical curve of $F_{(2)}/k$ similar as above. As is this would involve roughly g^2 indeterminates, and attempts to reduce this even larger number may again require the use of auxiliary places. Also, the computation of canonical curves is in general not very easy. Moreover, for our above strategy the number of the required $x_{(\alpha),i}$ can be considerably smaller than g, resulting in easier to handle affine curves, and it also works for hyperelliptic function fields. Therefore, we do not pursue the use of canonical curves in this paper.

3 Relating affine models

We now do not just choose the first r pole numbers n_i but let the n_i be special generators of the Weierstrass semigroup at $P_{(\alpha)}$ satisfying the following condition. We require that $n_j \neq n_i \bmod n_1$ for $1 \leq i < j$ and $1 \leq j \leq r$. We observe that the n_i are uniquely determined, because n_1 is the smallest pole number greater than zero at $P_{(\alpha)}$. The $x_{(\alpha),i}$ are again chosen in $\mathcal{L}(n_i P_{(\alpha)})$ such that $v_P(x_{(\alpha),i}) = -n_i$.

Given the congruence inequality it is not difficult to see that the elements $1, x_{(\alpha),2}, \ldots, x_{(\alpha),r}$ are a $k[x_{(\alpha),1}]$-basis of the integral closure $\mathrm{Cl}(k[x_{(\alpha),1}], F)$. Because $[F : k(x_{(\alpha),1})] = n_1$ this implies $r = n_1$. Since $\mathrm{Cl}(k[x_{(\alpha),1}], F)$ is a free

$k[x_{(\alpha),1}]$-module we know that there are uniquely determined $\lambda_{(\alpha),i,j,\nu} \in k[t]$ such that $x_{(\alpha),i}x_{(\alpha),j} = \lambda_{(\alpha),i,j,1}(x_{(\alpha),1}) + \sum_{\nu=2}^{n_1} \lambda_{(\alpha),i,j,\nu}(x_{(\alpha),1})x_{(\alpha),\nu}$ for $2 \leq i,j \leq n_1$. Looking at the valuations $v_{P_{(\alpha)}}$ we obtain $\deg(\lambda_{(\alpha),i,j,\nu}) \leq (n_i + n_j - n_\nu)/n_1$, and equality must hold for at least one ν because of the incongruence relations of the n_i. An algorithm to compute these $\lambda_{(\alpha),i,j,\nu}$ is given in [3].

Let $I_{(\alpha)}$ be the kernel of the substitution homomorphism $k[t_1,\ldots,t_{n_1}] \to F_{(\alpha)}$, $t_i \mapsto x_{(\alpha),i}$. From the previous discussion we see that $I_{(\alpha)}$ has a nice set of generators of the form $t_it_j - \lambda_{(\alpha),i,j,1}(t_1) - \sum_{\nu=2}^{n_1} \lambda_{(\alpha),i,j,\nu}(t_1)t_\nu$ for $2 \leq i,j \leq n_1$. In particular, given any polynomial $f \in k[t_1,\ldots,t_{n_1}]$ we can reduce it modulo $I_{(\alpha)}$ to a polynomial of degree at most one in $t_2,\ldots,t_{n_1}$ by substituting terms t_it_j with $\lambda_{(\alpha),i,j,1}(t_1) + \sum_{\nu=2}^{n_1} \lambda_{(\alpha),i,j,\nu}(t_1)t_\nu$.

Since we are working with a possibly smaller number of $x_{(\alpha),i}$ than in the previous section these elements do not necessarily realize all pole numbers anymore. Because of the congruence conditions on the n_i this can now be done using elements of the form $x_{(\alpha),1}^j x_{(\alpha),i}$, and a basis of $\mathcal{L}(n_rP)$ is given by $\{x_{(\alpha),1}^j x_{(\alpha),i} \mid jn_1 + n_i \leq n_r\}$. It follows that the transformation of Section 2 takes the form $\phi(x_{(1),i}) = \sum_{j=0}^{i} \mu_{i,j}(x_{(2),1})x_{(2),j}$ for suitable $\mu_{i,j} \in k[t]$ with $\deg(\mu_{i,j}) \leq (n_i - n_j)/n_1$ and $\mu_{i,i} \neq 0$. Reducing the generators of $I'_{(2)}$ modulo $I_{(2)}$ in $R[t_1,\ldots,t_{n_1}]$ and equating the coefficients of the t_i and 1 with zero finally yields equations for the coefficients of the $\mu_{i,j}$.

4 Relating expansions at a place depending on two parameters

We now discuss how to relate p-adic expansions at $P_{(1)}$ and $P_{(2)}$ in a meaningful way, depending on only two parameters.

We let $\pi_{(\alpha)}$ denote a local uniformizer at $P_{(\alpha)}$. Computing expansions in terms of $\pi_{(\alpha)}$ we may embed $F_{(\alpha)} \subseteq k((\pi_{(\alpha)}))$ and the isomorphism ϕ extends to an isomorphism of $k((\pi_{(1)}))$ and $k((\pi_{(2)}))$. Local uniformizers are not uniquely determined and we cannot hope for $\phi(\pi_{(1)}) = \pi_{(2)}$ if $\pi_{(1)}$ and $\pi_{(2)}$ are chosen independently. However, $\phi(\pi_{(1)})$ is a local uniformizer at $P_{(2)}$ and there are $c_i \in k$, $c_1 \neq 0$ such that $\phi(\pi_{(1)}) = \sum_{i=1}^{\infty} c_i\pi_{(2)}^i$. Of course, the c_i are unkown to us. Looking at $x_{(1),1}$ and $x_{(2),1}$ we see that $\phi(x_{(1),1}) = ax_{(2),1} + b$ for some $a,b \in k$, $a \neq 0$, which are also unknown to us. Furthermore, $x_{(\alpha),1} = \sum_{i=0}^{\infty} d_{(\alpha),i}\pi_{(\alpha)}^{i-n_1}$ for $d_{(\alpha),i} \in k$, $d_{(\alpha),0} \neq 0$. These coeffcients can be computed since $x_{(\alpha),1}, \pi_{(\alpha)}$ are known and lie in the same field respectively. Putting things together relates the c_i to a,b and gives the equation

$$\sum_{i=0}^{\infty} d_{(1),i}\phi(\pi_{(1)})^i = a\phi(\pi_{(1)})^{n_1}/\pi_{(2)}^{n_1} \sum_{i=0}^{\infty} d_{(2),i}\pi_{(2)}^i + b\phi(\pi_{(1)})^{n_1}. \tag{1}$$

We want to solve this equation for $\phi(\pi_{(1)})$ and recursively for c_1, c_2, ... Let $n = n_1$. Equating the coeffcients of $\pi_{(2)}^i$ for $i = 0$ gives $d_{(1),0} = ac_1^n d_{(2),0}$, hence $c_1^n = d_{(1),0}(d_{(2),0}a)^{-1}$. Write $n = p^{r_n}n'$ with $n' \neq 0 \bmod p$, where p denotes the

characteristic of k. Let $s \geq 1$ be such that $d_{(1),s} \neq 0$, and write $s = p^{r_s}s'$ with $s' \neq 0 \bmod p$. For $p = 0$ we let $n = n'$, $s = s'$ and $r_n = r_s = 0$. The terms of the smallest degree in $\pi_{(2)}$ involving c_j are of the form $d_{(1),s}(s'c_1^{s'-1}c_j)^{p^{r_s}}\pi_{(2)}^{s+(j-1)p^{r_s}}$, $a(n'c_1^{n'-1}c_j)^{p^{r_n}}\pi_{(2)}^{(j-1)p^{r_n}}$ and $b(n'c_1^{n'-1}c_j)^{p^{r_n}}\pi_{(2)}^{n+(j-1)p^{r_n}}$ for the three main expressions in equation (1) respectively, where $s+(j-1)p^{r_s}$ is minimal with $d_{(1),s} \neq 0$. As a result, using $a = d_{(1),0}(c_1^n d_{(2),0})^{-1}$, there are monic $f_j \in k[c_1, c_1^{-1}][t]$ consisting only of p-power terms in t and $g_j \in k[b, c_1, c_1^{-1}, c_2, c_3, \ldots, c_{j-1}]$ for every $j \geq 2$ such that equation (1) implies $f_j(c_j) = g_j$. Regarding a, b and the c_i as indeterminates we define $R_{a,b} = k[a, b, c_1, c_1^{-1}, c_2, c_3, \ldots]/I_{a,b}$ where $I_{a,b}$ is the ideal generated by $d_{(1),0} - ac_1^n d_{(2),0}$ and $f_j(c_j) - g_j$ for $j \geq 2$ and obtain the generic embedding $\phi_{a,b} : k((\pi_{(1)})) \to R_{a,b}((\pi_{(2)}))$ since the image of $\pi_{(1)}$ under $\phi_{a,b}$ is invertible in $R_{a,b}$. Now $\phi_{a,b}$ specializes to ϕ if we substitute the correct elements of k for a, b and the c_i corresponding to ϕ.

For $n \neq 0 \bmod p$ or $p = 0$ the terms of the smallest degree in $\pi_{(2)}$ involving c_j are $a(n'c_1^{n'-1}c_j)^{p^{r_n}}\pi_{(2)}^{(j-1)p^{r_n}} = anc_1^{n-1}c_j\pi_{(2)}^{j-1}$, and the f_j are hence all linear. This means that c_1 is an n-th root depending on a and that all other c_i are uniquely determined by c_1 and b. For $n = 0 \bmod p$ a more detailed analysis of the f_j shows that there can be at most n solution vectors $(c_i)_{i\geq 1}$ satisfying equation (1). The $(c_i)_{i\geq 1}$ thus depend also only on a, b up to finitely many possibilities. Looking at higher powers $\pi_{(2)}$ may give (usually gives) additional linear conditions on the c_j if the corresponding $d_{(1),i}$ are not zero. We remark that this strategy is not efficient if the series expansions of $x_{(1),1}$ and $x_{(2),1}$ are such that a larger number of the f_j do not have degree one ($n = 0 \bmod p$ necessarily).

A particularly easy form of the powers of $\phi(\pi_{(1)})$ in terms of $\pi_{(2)}$ can be achieved for $n \neq 0 \bmod p$ or $p = 0$, if we choose $\pi_{(1)}$ and $\pi_{(2)}$ in a special way. Let $h_{(\alpha)} = t^n - 1/x_{(\alpha),1}$. Since $v_{P_{(\alpha)}}(1/x_{(\alpha),1}) = n_1$ it follows that $h_{(\alpha)}$ has a root $\tilde{\pi}_{(\alpha)} \in k[[\pi_{(\alpha)}]]$, and $\tilde{\pi}_{(\alpha)}$ is a local uniformizer. If we require $\tilde{\pi}_{(\alpha)} = \pi_{(\alpha)} + O(\pi_{(\alpha)})$ then $\tilde{\pi}_{(\alpha)}$ is uniquely determined by $x_{(\alpha),1}$. Inverting the representation of $\tilde{\pi}_{(\alpha)}$ in terms of $\pi_{(\alpha)}$ gives a representation of $\pi_{(\alpha)}$ in terms of $\tilde{\pi}_{(\alpha)}$, and we may hence write everything in terms of $\tilde{\pi}_{(\alpha)}$ or equivalently embed $F_{(\alpha)} \subseteq k((\tilde{\pi}_{(\alpha)}))$. Also, ϕ extends to an isomorphism of $k((\tilde{\pi}_{(1)}))$ and $k((\tilde{\pi}_{(2)}))$ but again $\phi(\tilde{\pi}_{(1)}) = \tilde{\pi}_{(2)}$ is not necessarily true. Equation 1 then simplifies to the following equation,

$$1 = a\phi(\tilde{\pi}_{(1)})^n/\tilde{\pi}_{(2)}^n + b\phi(\tilde{\pi}_{(1)})^n. \tag{2}$$

With $\phi(\tilde{\pi}_{(1)}) = \tilde{c}_1\tilde{\pi}_{(2)} + O(\tilde{\pi}_{(2)}^2)$ and $\tilde{c}_1^n = 1/a$ equation (2) yields

$$\phi(\tilde{\pi}_{(1)})^r = \sum_{i=0}^{\infty} \rho_i b^i \tilde{c}_1^{r+in}\tilde{\pi}_{(2)}^{r+in} \tag{3}$$

where $r \neq 0$ and $\rho_i \in k$. This explicit form will be used in the next section.

5 Relating affine models at a place depending on two parameters

We now want to reduce the number of indeterminate coefficients in the polynomials $\mu_{i,j}$ to basically two indeterminates.

Using the notation of the previous section we have $\mu_{1,1} = a$, $\mu_{1,0} = b$ and $\mu_{0,0} = 1$. A local uniformizer at $P_{(\alpha)}$ is given by $\pi_{(\alpha)}$, and for $n_1 \neq 0 \bmod p$ or $p = 0$ we may assume that $\pi_{(\alpha)}$ is an n_1-th root of $1/x_{(\alpha),1}$.

For every pole number d_i of $P_{(\alpha)}$ let $w_{(\alpha),i} = x_{(\alpha),1}^{\nu} x_{(\alpha),j}$ where $d_i = \nu n_1 + n_j$. As mentioned the $w_{(\alpha),0}, \ldots, w_{(\alpha),i}$ form a basis of $\mathcal{L}(d_i P_{(\alpha)})$ because of the congruence relations of the n_j. We can compute the expansions $w_{(\alpha),j} = \sum_{i=-d_j}^{\infty} \rho_{(\alpha),i,j} \pi_{(\alpha)}^i$ up to any precision. However, precision $O(\pi_{(\alpha)})$ will be sufficient. Using a Gaussian elimination procedure we may assume that $\rho_{(\alpha),i,\nu} = 0$ for all $i = -d_j$ and $\nu > j \geq 0$ and $\rho_{(\alpha),-d_j,j} = 1$ for all $0 \leq j \leq d_1$: For decreasing j we simply eliminate the $-d_j$-th coefficients in the expansions of $w_{(\alpha),\nu}$ by replacing $w_{(\alpha),\nu}$ with $w_{(\alpha),\nu} - (\rho_{(\alpha),-d_j,\nu}/\rho_{(\alpha),-d_j,j}) w_{(\alpha),\nu}$ for all $\nu > j$. Finally, we replace $w_{(\alpha),j}$ by $(1/\rho_{(\alpha),-d_j,j}) w_{(\alpha),j}$. We now define $x_{(\alpha),j} = w_{(\alpha),i}$ where $n_j = d_i$. The resulting $x_{(\alpha),j}$ are then uniquely determined depending on the chosen local uniformizer.

In the two fields situation we assume that the $x_{(2),i}$ have been transformed using this Gaussian elimination procedure. We now know that the Gaussian elimination for $\phi(x_{(1),i})$ would yield the $x_{(2),i}$, but the $\phi(x_{(1),i})$ are unknown to us. We can however perform this Gaussian elimination for the $\phi_{a,b}(x_{(1),i})$ over $R_{a,b}$ in a generic way. Note that the "leading" coefficients of $\phi_{a,b}(x_{(1),1}^{\nu} x_{(1),i})$ are basically powers of c_1 and are hence invertible in $R_{a,b}$. The Gaussian elimination procedure thus precisely yields the $\mu_{i,j}$ as elements of $R_{a,b}[t]$. Since we can work with precision $O(\pi_{(2)})$, we have only to deal and to compute with elements of $R_{a,b}$ involving c_i for $i \leq n_r + 1$.

Taking up the strategy of Section 3 (see in particular its end) and clearing c_1 from the denominators in the coefficients of the $\mu_{i,j}$ and in $I_{a,b}$ yields equations in a, b and the c_i with $i \leq n_1 + 1$. The corresponding ideal in $k[a, b, c_1, \ldots, c_{n_1+1}]$ is zero dimensional, since there can only be finitely many solutions for a, b and the c_i can only assume finitely many different values given a, b. Thus the intersection with $k[a]$ and $k[a, b]$ also results in zero dimensional ideals, and the possible values of a and b and then c_i can be computed from this.

For $n_1 \neq 0 \bmod p$ or $p = 0$ things are much more explicit. Using the special local uniformizer we perform the Gaussian elimination on the $x_{(1),i}$ and the $x_{(2),i}$. Because of equation 2 this special form is almost preserved by ϕ, namely we have that $\phi(x_{(\alpha),i}) = c_1^{n_i} x_{(2),i}$ with $c_1^n = 1/a$ for $i \geq 2$ and $\phi(x_{(1),1}) = a x_{(2),1} + b$. Then

$$\phi(x_{(1),i} x_{(1),j}) = \lambda_{(1),i,j,1}(a x_{(2),1} + b) + \sum_{\nu=2}^{n_1} \lambda_{(1),i,j,\nu}(a x_{(2),1} + b) c_1^{n_\nu} x_{(2),\nu},$$

$$\phi(x_{(1),i})\phi(x_{(1),j}) = c_1^{n_i+n_j} \lambda_{(2),i,j,1}(x_{(2),1}) + \sum_{\nu=2}^{n_1} c_1^{n_i+n_j} \lambda_{(2),i,j,\nu}(x_{(2),1}) x_{(2),\nu}.$$

Combining the right hand sides via $\phi(x_{(1),i}x_{(1),j}) = \phi(x_{(1),i})\phi(x_{(1),j})$ yields easy equations for a, b and c_1.

6 Computing the isomorphisms

The choice of the sets of places of $F_{(1)}/k$ and $F_{(2)}/k$ which must be mapped to each other under any isomorphism ϕ depends mainly on the constant field and the genus. For constant fields other than finite fields we always consider subsets of the set of Weierstrass places of $F_{(1)}/k$ and $F_{(2)}/k$ of smallest cardinality, where the places have a particular common gap sequence. The number of such places can be quite high, a lower and upper bound in characteristic zero are given by $2(g+1)$ and $(g-1)g(g+1)$ and in general by $O(g^3)$. This leads to roughly g^2 up to g^6 comparisons of places $P_{(1)}$ and $P_{(2)}$. For finite fields we check whether the estimated number $q \pm 2gq^{1/2}$ of places of degree one is (considerably) smaller than an approximate expected number of Weierstrass places. This would lead to roughly q^2 comparisons of places and is in $O(g^6)$, but can also be much smaller. Note that places of a prescribed small degree can be computed efficiently for global function fields. A bound for the number of isomorphisms is given by the Hurwitz bound $84(g-1)$ in characteristic zero and by roughly $16g^4$ in positive characteristic (details can be found in [7]). Such large automorphism groups are only obtained for very special function fields.

It may happen that we cannot find suitable places of degree one, but this was essential for the strategy described above. A solution to this problem is to consider constant field extensions, over which we would obtain suitable places of degree one. For example, if $P_{(\alpha)}$ is a place of degree two over k which splits into Weierstrass places over an algebraic closure it is sufficient to consider the constant field extension $F_{(\alpha)}k(P_{(\alpha)})/k(P_{(\alpha)})$ by the residue class field of $P_{(\alpha)}$, since $P_{(\alpha)}$ already splits in $F_{(\alpha)}k(P_{(\alpha)})/k(P_{(\alpha)})$. We would then compute an isomorphism of $k(P_{(1)})$ and $k(P_{(2)})$ and isomorphisms ϕ of $F_{(1)}k(P_{(1)})/k(P_{(1)})$ and $F_{(2)}k(P_{(2)})/k(P_{(2)})$ which possibly do not come from isomorphisms of $F_{(1)}/k$ and $F_{(2)}/k$. This can be checked as follows. A generating system of $F_{(\alpha)}/k$ is also a generating system of $F_{(\alpha)}k(P_{(\alpha)})/k(P_{(\alpha)})$. If we compute the effect of ϕ in terms of these generating systems it is easy to check whether ϕ restricts to an isomorphism of $F_{(1)}/k$ and $F_{(2)}/k$. This operation requires the inversion of isomorphisms or inverting the representation of one set of generators in terms of another set of generators. A way of doing this computation is described in [3] and can also be achieved for our special elements $x_{(\alpha),i}$ by linear algebra over k involving algebraic functions of bounded degree. Since every isomorphism of $F_{(1)}/k$ and $F_{(2)}/k$ extends to an isomorphism of $F_{(1)}k(P_{(1)})/k(P_{(1)})$ and $F_{(2)}k(P_{(2)})/k(P_{(2)})$ this method yields all isomorphisms of $F_{(1)}/k$ and $F_{(2)}/k$.

We summarize the single steps for computing the isomorphisms between $F_{(1)}/k$ and $F_{(2)}/k$.

Algorithm 4 *(Isomorphisms)*

Input: *Function fields $F_{(1)}/k$ and $F_{(2)}/k$.*
Output: *A set L of all isomorphisms of $F_{(1)}/k$ and $F_{(2)}/k$.*

1. *Check whether $F_{(1)}/k$ and $F_{(2)}/k$ have the same genus g. If not then return the empty list L.*
2. *Let $p = \mathrm{char}(k)$, $q = \#k$ and $d = 1$.*
3. *If $g^3 < q^d$ compute lists $S_{(1)}$ and $S_{(2)}$ of all Weierstrass places of $F_{(1)}/k$ and of $F_{(2)}/k$ (degrees greater one allowed). Check whether numbers, degrees and gap sequences coincide. Determine suitable small subsets $S'_{(1)}$ and $S'_{(2)}$ as discussed above.*
4. *If $g^3 \geq q^d$ compute lists $S_{(1)}$ and $S_{(2)}$ of all places of degree d. If there are no such places let $d \leftarrow d + 1$ and go to step 3.*
5. *Choose a fixed $P_{(1)} \in S'_{(1)}$. If $\deg(P_{(1)}) > 1$ work with $F_{(1)}k(P_{(1)})/k(P_{(1)})$ in the following. Compute the n_i, $x_{(1),i}$ and d_j, $w_{(1),j}$ as in Section 3 and Section 5. If $n_1 \neq 0 \bmod p$ or $p = 0$ set $c = 1$, otherwise set $c = 0$. Choose a local uniformizer $\pi_{(1)}$ at $P_{(1)}$. If $c = 1$ then choose the special $\pi_{(1)}$ of Section 4. Apply the Gaussian elimination procedure of Section 5 to the $x_{(1),i}$.*
6. *The following three steps are done for every $P_{(2)} \in S'_{(2)}$.*
7. *Check that $k(P_{(2)})/k$ is isomorphic to $k(P_{(1)})/k$. If not, take the next $P_{(2)}$. If yes work with $F_{(2)}k(P_{(2)})/k(P_{(2)})$ in the following and identify $k(P_{(2)})$ and $k(P_{(2)})$.*
8. *Compute the $x_{(2),i}$ and $w_{(2),j}$ for $P_{(2)}$ as in Section 3 and Section 5. If $c = 1$ then choose the special $\pi_{(2)}$ of Section 4. Apply the Gaussian elimination procedure of Section 5 to the $x_{(2),i}$.*
9. *If $c = 1$ then solve for a, b, c_1 as described at the end of section 5. Otherwise solve for $a, b, c_1, c_2, \ldots$ as described in section 5. For any solution recover the $\mu_{i,j}$ and ϕ. If ϕ restricts to an isomorphism ϕ' of $F_{(1)}/k$ and $F_{(2)}/k$, then $L \leftarrow L \cup \{\phi'\}$.*
10. *Return L.*

We remark that we have implemented a prototype of Algorithm 4 in Magma [1, 2] which shows that the algorithm is quite practical in the global function field case.

References

1. W. Bosma, J. Cannon, and C. Playoust. The Magma algebra system I: The user language. *J. Symbolic Comp.*, 24, 3/4:235–265, 1997.
2. Comp. algebra group. Magma. `http://www.maths.usyd.edu.au:8000/u/magma/`, 2003.
3. F. Hess. An algorithm for computing Weierstrass points. In C. Fieker and D. R. Kohel, editors, *Proceedings of the Fifth Symposium on Algorithmic Number Theory, ANTS-V*, LNCS 2369, pages 357–371, Sydney, Australia, 2002. Springer-Verlag, Berlin-Heidelberg-New York.

4. F. Hess. Computing Riemann-Roch spaces in algebraic function fields and related topics. *J. Symbolic Comp.*, 33(4):425–445, 2002.
5. Kant group. Kash. `http://www.math.tu-berlin.de/~kant`, 2003.
6. J. Klüners. Algorithms for function fields. *Exp. Math.*, 11:171–181, 2002.
7. H. Stichtenoth. Über die Automorphismengruppe eines algebraischen Funktionenkörpers von Primzahlcharakteristik. *Arch. Math.*, 24:527–544, 1973.
8. H. Stichtenoth. *Algebraic Function Fields and Codes.* Springer-Verlag, Berlin-Heidelberg-New York, 1993.