Efficient Resolution of Singularities of Plane Curves*

Dexter Kozen

Computer Science Department
Cornell University
Ithaca, New York 14853, USA
kozen@cs.cornell.edu

Abstract. We give a new algorithm for resolving singularities of plane curves. The algorithm is polynomial time in the bit complexity model, does not require factorization, and works over $\mathbb{Q}$ or finite fields.

1 Introduction

Resolving singularities is a central problem in computational algebraic geometry. In this paper we describe a new algorithm for resolving singularities of irreducible plane curves. The algorithm runs in polynomial time in the bit complexity model, does not require polynomial factorization, and works over $\mathbb{Q}$ or any finite field.

Classical algorithms for resolving singularities [2, 15, 7] use a combination of methods involving

- the *Newton polygon*, a polygon in $\mathbb{Z}^2$ whose vertices are the exponents of terms in f;
- *Puiseux series*, power series with fractional exponents.

These algorithms take polynomial time if we assume efficient factorization over algebraic extensions of the base field and unit-time arithmetic these extensions.

Teitelbaum [13] establishes bounds on the degree of field extensions containing the Puiseux coefficients, leading to an algorithm that is polynomial in the number of base field operations. However, the algorithm is not analyzed for its bit complexity. The problem of intermediate coefficient swell is not often addressed in theoretical analyses, but is a serious consideration in practical implementations [14, 17].

Chistov [4] and Walsh [16] give algorithms that are polynomial time in the bit complexity model, but require factorization of polynomials over $\mathbb{Q}$ and algebraic number fields. The best known algorithms for these problems [12, 11, 9] are theoretically polynomial time in the bit complexity model, but quite inefficient from a practical standpoint. There are currently no known deterministic

* Proc. 14th Conf. Foundations of Software Technology and Theoretical Computer Science, December 1994, Madras, India, to appear.

polynomial-time algorithms for factorization of polynomials over finite fields. For practical considerations, it is important to avoid factorization if at all possible.

A novel feature of our algorithm is that it does not calculate Puiseux series explicitly. Rather, given a curve $f(x,y) = 0$, it builds a tree whose vertices are labeled with approximations to local parameters of the exceptional places of the curve. These are not Puiseux series as in the classical algorithms, but rational functions of x and y obtained by direct reparametrizations based on information obtained from the Newton polygon.

The algebraic numbers involved are represented and manipulated symbolically, using the technique of *passive factorization* [5, 6, 13] in conjunction with the squarefree decomposition algorithm of [1]. All necessary calculations can be carried out with this implicit representation, thus avoiding factoring and the explicit construction of high-degree extensions.

As an immediate corollary, we obtain an efficient algorithm for calculating the genus of a plane curve using the Hurwitz formula [10, 7]. Trager has recently given a polynomial-time algorithm for the genus problem using different techniques [14].

2 Algebraic Preliminaries

Let R be an algebraic function field over an algebraically closed field K, *i.e.*, a finite extension of a purely transcendental extension $K(x)$ of K. A *place* of R is a valuation ring $\mathbf{P} \subsetneq R$, *i.e.*, a subring such that $K \subseteq \mathbf{P} \neq R$ and for all $u \in R$, either $u \in \mathbf{P}$ or $1/u \in \mathbf{P}$. $\mathbf{P}$ has a unique maximal ideal $\mathbf{I} = \{u \in \mathbf{P} \mid 1/u \notin \mathbf{P}\}$, and $\mathbf{P}/\mathbf{I} \cong K$. $\mathbf{I}$ is a principal ideal; a principal generator $t_{\mathbf{P}}$ is called a *local parameter*.

For $R = K(x)$, there is exactly one place for each $a \in K$ with local parameter $x - a$, denoted $x = a$, plus one place with local parameter $1/x$, denoted $x = \infty$.

If $\mathbf{P}$ is a place of R, then $\mathbf{p} = \mathbf{P} \cap K(x)$ is a place of $K(x)$ and $t_{\mathbf{p}} = t_{\mathbf{P}}^{c} u$ for some unit $u \in \mathbf{P}$ and positive integer c. The number c is called the *ramification index* of $\mathbf{P}$ over $\mathbf{p}$. Every place $\mathbf{p}$ of $K(x)$ extends to at least one place $\mathbf{P}$ of R.

Now let $f(x,y) = 0$ be an irreducible plane curve of degree n in y with coefficients in K, and let $R = K(x)[y]/f$. A place $\mathbf{p}$ of $K(x)$ is called *simple* if for any two places $\mathbf{P}, \mathbf{P}'$ of R over $\mathbf{p}$, $\mathbf{P} \cap K(y) \neq \mathbf{P}' \cap K(y)$, and all places of R over $\mathbf{p}$ have ramification index 1 (*i.e.*, $t_{\mathbf{p}}$ is a local parameter); otherwise it is called *exceptional*. There are only finitely many exceptional places of $K(x)$.

See [10, 7, 2, 3] for details.

2.1 Statement of the Problem

The problem we consider in this paper is as follows. Given an irreducible plane curve $f(x,y) = 0$ with coefficients in some subfield $k \subseteq K$, determine all the exceptional places $\mathbf{p}$ of $K(x)$. For each such $\mathbf{p}$, determine the set of places $\mathbf{P}$ of R over $\mathbf{p}$. For each such $\mathbf{P}$, determine its ramification index over $\mathbf{p}$, a local

parameter $t_{\mathbf{P}}$, and a birational reparametrization (change of variables) giving $t_{\mathbf{P}}$ in terms of x and y.

In the solution of this problem, all computation should take place in the subfield k. In practice, k will usually be $\mathbb{Q}$ or a finite field.

Example 1. Consider the hyperbola $y^2 = x^2 - 1$.

There are three exceptional places $x = \pm 1$ and $x = \infty$ of $K(x)$. Each of the first two places extends to one place of R with ramification index 2, with local parameters $(x \pm 1)/y$. The third extends to two places of R, each with ramification index 1 over $K(x)$, corresponding to the two asymptotes $x = \pm y$ of the hyperbola. The two local parameters are $x/y \pm 1$.

2.2 Efficient Birational Reparametrization

A critical subroutine is *birational reparametrization*, or change of variables. We will often wish to reparametrize an irreducible plane curve $f(x, y) = 0$ in terms of new parameters u, v that are rational functions of x, y. We will need to compute the defining equation $g(u, v) = 0$ of the curve in terms of the new parameters.

The problem of birational reparametrization of algebraic surfaces has an efficient solution in general, using multivariate gcds. For our application, all cases we will encounter will be of the following special form:

$$\begin{aligned} u &= (x-a)^s (y-b)^{-t} + c \\ v &= (x-a)^{-n} (y-b)^m + d \end{aligned} \tag{1}$$

where $a, b, c, d \in K$ and $s, t, m, n \in \mathbb{Z}$ such that $sm - tn = \pm 1$. Any such map is invertible, and its inverse is of the same form:

$$\begin{aligned} x &= (u-c)^m (v-d)^t + a \\ y &= (u-c)^n (v-d)^s + b \end{aligned}$$

We can reparametrize efficiently as follows:

1. Substitute $(u-c)^m(v-d)^t + a$ for x and $(u-c)^n(v-d)^s + b$ for y in f to obtain $f((u-c)^m(v-d)^t + a, (u-c)^n(v-d)^s + b)$.
2. Multiply or divide by appropriate powers of $u-c$ and $v-d$ as necessary to clear denominators and eliminate extraneous factors of the form $u-c$ and $v-d$.

The resulting polynomial $g(u,v)$ is the irreducible representation of the plane curve $f(x,y) = 0$ in terms of the parameters u, v.

Example 2. Consider the irreducible curve $y^2 = x^3 - x^4$ and the map $u = x^2/y$, $v = y^2/x^3$. This is a birational map of the form (1) with $(s,t,m,n) = (2,1,2,3)$ and $a = b = c = d = 0$. Its inverse is $x = u^2v$, $y = u^3v^2$. To reparametrize in terms of u, v, substitute u^2v and u^3v^2 for x and y, respectively, in the equation of the curve to get $u^6v^4 = u^6v^3 - u^8v^4$ and eliminate the factor u^6v^3, leaving the irreducible equation $v = 1 - u^2v$. Rewriting this as $v = 1/(1+u^2)$, we see that the curve is of genus zero and that u is a generator of the function field.

2.3 Squarefree Decomposition

Another key subroutine is *squarefree decomposition*. One wishes to transform a set A of polynomials into another set A' such that the elements of A' are squarefree and pairwise relatively prime, and every element of A is a product of powers of elements of A'. An efficient solution to this problem was given in [1].

3 First Algorithm

In this section, we describe the high-level structure of our algorithm. For this section only, we assume that K is algebraically closed, that arithmetic in K is unit-cost, and that we are provided with an efficient algorithm for factoring univariate polynomials over K. These strong assumptions are to simplify the high-level description. Once the algorithm is understood, we will describe in the next section how to remove these assumptions.

Most of the techniques used here, such as the *Newton polygon*, are well known [2, 7, 15]. However, as mentioned in the introduction, one novel feature of our algorithm is that we do not extend to the field of Puiseux series to obtain local parameters, but reparametrize directly to obtain local parameters in the function field itself. This approach allows us to give an explicit description of the coefficients occurring in the computation and analyze their bit complexity that will be useful in the analysis of the next section.

3.1 Overview

We are given an irreducible polynomial $f(x,y) \in K[x,y]$. The algorithm builds a rooted labeled tree. Each vertex τ in the tree is labeled with a reparametrization $f_\tau(x_\tau, y_\tau)$ of $f(x,y)$, along with some other information Δ_τ that specifies a set

of places $\mathbf{P}_\tau$ of R. The tree grows at the leaves as more information becomes available that allows us to split or refine the set $\mathbf{P}_\tau$. The information is calculated from the Newton polygon of $f_\tau(x_\tau, y_\tau)$.

The edges of the tree are labeled with positive integers giving ramification information, along with coefficients of birational transformations. If τ is a child of σ, then the integer labeling the edge (σ, τ) gives the ramification index of x_σ over x_τ, and the birational transformation gives x_τ and y_τ in terms of x_σ and y_σ.

The information Δ_τ is a Boolean combination of constraints of the form $h(x_\tau, y_\tau) \in \mathbf{P}$ or $h(x_\tau, y_\tau) \in \mathbf{I}$, where the symbols $\mathbf{P}, \mathbf{I}$ range over places of R and their unique maximal ideals, respectively. The set $\mathbf{P}_\tau$ is the set of places of R satisfying these constraints.

We first do a precomputation involving discriminants to find all exceptional places of $K(x)$, and reparametrize so as to move these places to the origin. (Even this step will be difficult without the assumptions above). We now want to resolve the point $(0, 0)$. In other words, we have $f_\tau(x_\tau, y_\tau)$ and $\Delta_\tau = \{x_\tau \in \mathbf{I}, y_\tau \in \mathbf{I}\}$, thus we want to find all places of R whose maximal ideal contains x_τ and y_τ, along with their local parameters and ramification indices over $K(x)$.

We cause the tree to branch at a vertex τ by adding new constraints that may partition the set of places associated with that vertex or give more ramification information. The new constraints are obtained from the Newton polygon of f_τ. We continue to expand the tree at the leaves, reparametrizing, adding more constraints, and refining the partition, until each leaf τ determines a unique place of R and x_τ is a local parameter.

When we are done, the ramification index of the unique place of R associated with a leaf τ is given by the product of the integers labeling the edges along the path from the root to τ. The composition of the birational transformations labeling the edges along the path gives x_τ and y_τ in terms of x and y.

3.2 Formal Description

We are given an irreducible polynomial $f(x, y) \in K[x, y]$.

1. Create the root ρ of the tree and label it $x_\rho = x$, $y_\rho = y$, $f_\rho = f$, $\Delta_\rho = \emptyset$.
2. Create a new vertex σ and an edge labeled 1 from ρ to σ. Reparametrize under the map $x_\sigma = 1/x_\rho$, $y_\sigma = y_\rho$ and label σ with the resulting irreducible polynomial $f_\sigma(x_\sigma, y_\sigma)$. Add the new constraint $x_\sigma \in \mathbf{I}$. Further resolution of this vertex will give all places of R over the place $x_\rho = \infty$ of $K(x_\rho)$.
3. Compute the discriminant $d(x) = \mathbf{disc}_y\ f(x, y)$ and factor it. For each distinct root a, create a new vertex τ and an edge with label 1 from ρ to τ. Reparametrize under the map $x_\tau = x_\rho - a$, $y_\tau = y_\rho$ and label τ with the resulting irreducible polynomial $f_\tau(x_\tau, y_\tau) = f_\rho(x_\tau + a, y_\tau)$, along with the new constraint $x_\tau \in \mathbf{I}$. Further resolution of this vertex will give all places of R over the place $x_\rho = a$ of $K(x_\rho)$.
4. For each child σ of ρ, if $f_\sigma(0, y_\sigma)$ is a nonzero constant, go on to step 5; there are no finite points $(0, b)$ on the curve $f_\sigma(x_\sigma, y_\sigma) = 0$, *i.e.*, no places or R

with $x_\sigma \in \mathbf{I}$ and $y_\sigma \in \mathbf{P}$. (Note $f_\sigma(0, y_\sigma)$ cannot be identically zero, since $f_\sigma(x_\sigma, y_\sigma)$ is irreducible.) Otherwise, $f_\sigma(0, y_\sigma)$ is of positive degree. Factor it, and for each distinct root b, create a new vertex τ and an edge with label 1 from σ to τ. Reparametrize under the map $y_\tau = y_\sigma - b$, $x_\tau = x_\sigma$ and label τ with the resulting irreducible polynomial $f_\tau(x_\tau, y_\tau) = f_\sigma(x_\tau, y_\tau + b)$. Add the constraints $x_\tau \in \mathbf{I}$, $y_\tau \in \mathbf{I}$.

5. Also for each child σ of ρ, determine whether the point $(0, \infty)$ lies on the curve $f_\sigma(x_\sigma, y_\sigma) = 0$ (*i.e.*, whether there is a place of R such that $x_\sigma \in \mathbf{I}$ and $1/y_\sigma \in \mathbf{I}$). This is done by checking whether the lead coefficient of f_σ, considered as a polynomial in y_σ with coefficients in $K[x_\sigma]$, is divisible by x_σ. If not, go on to step 6. Otherwise, create a new vertex τ and an edge with label 1 from σ to τ. Reparametrize under the map $y_\tau = 1/y_\sigma$, $x_\tau = x_\sigma$ to get the irreducible polynomial $f_\tau(x_\tau, y_\tau)$. We must have $f_\tau(0, 0) = 0$, otherwise the test above would have failed. Label τ with $f_\tau(x_\tau, y_\tau)$ and add the constraints $x_\tau \in \mathbf{I}$, $y_\tau \in \mathbf{I}$.

We now have at each leaf σ an irreducible polynomial $f_\sigma(x_\sigma, y_\sigma)$ and constraints $x_\sigma \in \mathbf{I}$, $y_\sigma \in \mathbf{I}$. By construction, f_σ has no constant term, *i.e.* $f_\sigma(0, 0) = 0$.

6. If $f_\sigma(x_\sigma, y_\sigma)$ has a linear term of the form cy_σ, stop expanding this branch. In this case x_σ is a local parameter, and there is a unique place satisfying Δ_σ.
7. Otherwise, assume $f_\sigma(x_\sigma, y_\sigma)$ has no constant term or term of the form cy_σ. For each edge of the Newton polygon of $f_\sigma(x_\sigma, y_\sigma)$ with slope $-m/n$, m and n relatively prime, create a new vertex τ and an edge with label m from σ to τ. Let s and t be a minimal pair of nonnegative integers such that $sm - tn = 1$. Reparametrize with respect to $x_\tau = x_\sigma^s / y_\sigma^t$ and $y_\tau = y_\sigma^m / x_\sigma^n$ as described in (1). Label τ with the resulting irreducible polynomial $f_\tau(x_\tau, y_\tau)$ and add the constraints $x_\tau \in \mathbf{I}$, $y_\tau \in \mathbf{P} - \mathbf{I}$.
8. For each vertex τ created in step 7 with label $f_\tau(x_\tau, y_\tau)$, factor $f_\tau(0, y_\tau)$. For each *nonzero* root a of $f_\tau(0, y_\tau)$ (at least one such root must exist), reparametrize with respect to the map $y_\upsilon = y_\tau - a$, $x_\upsilon = x_\tau$ to get $f_\upsilon(x_\upsilon, y_\upsilon) = f_\tau(x_\upsilon, y_\upsilon + a)$. Create a new vertex υ with that label and edge labeled 1 from τ to υ. Add the constraints $x_\upsilon \in \mathbf{I}$, $y_\upsilon \in \mathbf{I}$.

Continue expanding leaves according to steps 7 and 8 until the stopping condition of step 6 obtains. This must happen eventually. When done, there is exactly one place of R for each leaf σ, the ramification index over $K(x_\rho)$ of this place is the product of all the edge labels along the path from ρ to σ, and x_σ is a local parameter of the place.

Example 3. Consider the hyperbola $y^2 = x^2 - 1$ of Example 1. We first compute the discriminant

$$\mathbf{disc}_y \ f(x, y) = 4(x^2 - 1)$$

with roots ± 1. Substituting either of these values for x in the equation of the curve gives $y^2 = 0$, which has a double root 0. Thus $(\pm 1, 0)$ are the two finite exceptional points.

Reparametrizing by the map $x \mapsto 1/x$, we obtain

$$x^2y^2 = 1 - x^2$$

and substituting 0 for x gives $0 = 1$, therefore there are no points on the curve of the form (∞, b) for b finite. A similar calculation shows that there are no points on the curve of the form (a, ∞) with a finite.

Reparametrizing by the map $(x, y) \mapsto (1/x, 1/y)$ gives

$$x^2 = y^2 - x^2y^2 \tag{2}$$

and the point $(0, 0)$ lies on this curve. It is an exceptional point since 0 is a multiple root of the polynomial $y^2 = 0$ obtained by substituting 0 for x.

We reparametrize the original equation under the maps $x \mapsto x \pm 1$ to move the two finite exceptional points to $(0, 0)$. The third equation (2) is already of the desired form. We now have three irreducible equations

$$y^2 = x^2 - 2x$$
$$y^2 = x^2 + 2x$$
$$x^2 = y^2 - x^2y^2$$

labeling three children of the root, and wish to resolve the point $(0, 0)$ for each.

Now we do step 7 for the first equation. We plot the three exponent vectors $(0, 2)$, $(2, 0)$, $(1, 0)$, and observe that the Newton polygon has a single edge through $(0, 2)$ and $(1, 0)$ with slope -2. According to step 7, we take $m = 2$, $n = 1$, $s = 1$, $t = 1$, $u = x/y$, $v = y^2/x$, $x = u^2v$, $y = uv$, and reparametrize to obtain the irreducible polynomial

$$q(u, v) = v - u^2v + 2$$

We create a child with this label and an edge to it labeled 2. Now we do step 8. We compute $q(0, v) = v + 2$, which has one nonzero root -2. We create a child with label

$$q(x, y - 2) = y - x^2y + 2x^2$$

and an edge labeled 1. The polynomial has a linear term y, so we stop expanding this branch.

There is one leaf and the product of the edge labels on the path from the root to this leaf is 2, so we conclude that there is one place of R over the place $x = -1$ of $K(x)$ with ramification index 2.

A similar computation holds for the second equation.

For the third equation

$$x^2 = y^2 - x^2y^2$$

we again plot the exponent vectors $(2, 0)$, $(0, 2)$, and $(2, 2)$. The Newton polygon has a single edge through $(2, 0)$ and $(0, 2)$ with slope -1, giving $m = 1$, $n = 1$, $s = 1$, $t = 0$, $u = x$, $v = y/x$, $y = uv$. Reparametrizing, we obtain

$$q(u, v) = 1 - v^2 + u^2v^2 \ .$$

Then $q(0,v)$ is $1-v^2$ which has two nonzero roots ± 1. We create two new children labeled

$$q(x,y-1) = -y^2 + 2y + x^2y^2 - 2x^2y + x^2$$
$$q(x,y+1) = -y^2 - 2y + x^2y^2 + 2x^2y + x^2$$

with edges to these vertices each labeled 1. The stopping condition holds at each leaf, since each has a linear term $\pm 2y$. We conclude that there are two places of R over the place $x = \infty$ of $K(x)$, each with ramification index 1.

4 Second Algorithm

Suppose that the coefficients of f lie in a subfield k of K. The algorithm of the previous section is polynomial time, assuming unit-cost arithmetic in finite extensions of k and the ability to factor. If we count the cost of the representation, a naive approach may require exponential time.

In this section we describe an implicit representation of the algebraic numbers used in the algorithm of the previous section that allow them to be manipulated using only arithmetic in the field k. We will use gcd heavily, but factorization is unnecessary. This technique is called *passive factorization* and has been used by Dicrescenzo and Duval [5, 6] and Teitelbaum [13]. We use this in conjunction with the squarefree decomposition algorithm of [1] to avoid intermediate coefficient swell when $k = \mathbb{Q}$.

4.1 Overview

Each of steps 3, 4, and 8 of the algorithm of §3 factored a polynomial $q(y)$, producing roots a. For each such a, we created a new child τ of the current vertex σ and associated a with τ. The reparametrization of f_σ to obtain f_τ involved the algebraic number a.

Here, instead of factoring q to obtain the algebraic numbers a, we will just do a squarefree decomposition of q using the algorithm of [1] to give finitely many pairwise relatively prime factors q_τ of q. We will create one new child τ for each such factor q_τ and associate q_τ with τ. Subsequently, all computation involving roots a of q_τ will be done *symbolically* by introducing a new indeterminate a and working modulo $q_\tau(a)$. In effect, we will be dealing with all the roots of q_τ *simultaneously* in one subtree, instead of having a separate subtree for each root as in the previous section.

It may happen that some later computation causes $q_\tau(a)$ to split further. This will occur when a discrepancy in the behavior of the resolution over two different roots of $q_\tau(a)$ is discovered. This discovery takes the form of another polynomial in a that has a nontrivial common factor with $q_\tau(a)$. When that occurs, we can return to τ and split q_τ using the gcd algorithm, then split τ into a separate vertex for every new factor of $q_\tau(a)$, replicating the entire subtree below each new vertex. The tree never gets too big because of the absolute bound n on the number of places.

It may happen that q_τ never splits. In that case the resolution over all roots of q_τ looks exactly the same, and the number of the places determined by the subtree below τ can just be multiplied by the degree of q_τ. Thus we have a single node representing multiple places, but we know how many and their ramifications.

Formally, at any vertex in the tree at any point in time, we have sequences

$$\overline{a} = a_0,\ a_1,\ a_2,\ \ldots,\ a_k \tag{3}$$

$$\overline{q} = q_0(a_0),\ q_1(a_0, a_1),\ q_2(a_0, a_1, a_2),\ \ldots,\ q_k(a_0, a_1, \ldots, a_k) \tag{4}$$

of indeterminates and polynomials describing them. The polynomials $\overline{q}$ generate a zero-dimensional variety $V(\overline{q})$ in K^{k+1}. Arithmetic on coefficients will be done modulo these q_i. We maintain the invariants that each q_i is reduced modulo $q_1, \ldots, q_{i-1}$ and that the q_i are squarefree. We also maintain the invariant that the coefficients of q_i, considered as polynomials in the indeterminates $a_1, \ldots, a_{i-1}$, are relatively prime to q_{i-1}.

These conditions are determined by signs (zero or nonzero) of various subresultants, which are minors of the Sylvester matrix [8]. These are polynomials $p(a_0, a_1, \ldots, a_{i-1})$ in the coefficients of q_i. For each such p, either its sign is determined uniquely by $\overline{q}$, *i.e.* p either vanishes on all of $V(\overline{q})$ or does not vanish on all of $V(\overline{q})$, which allows the computation to be carried on; or p and q_{i-1} have a nontrivial common factor, *i.e.* p vanishes on some nonempty subset of $V(\overline{q})$ and does not vanish on some other nonempty subset, in which case q_{i-1} splits and we get a refinement. When all the splitting is done, we again have the property that $p(\overline{a})$ either vanishes on every point of $V(\overline{q})$ or on no point of $V(\overline{q})$.

The key idea being exploited here is the same used in the BKR algorithm for the theory of real closed fields [1], namely that we do not need to know the algebraic numbers a themselves to solve the resolution problem, but only the signs of certain polynomials in these algebraic numbers. This idea leads to an efficient implementation of the passive factorization method of [5, 6, 13].

Coefficients remain small under these symbolic operations. For example, in step 7, we reparametrize

$$f(x, y) = \sum_{ij} p_{ij}(\overline{a}) x^i y^j$$

to obtain a new polynomial $g(u, v)$. The terms of $g(u, v)$ and those of $f(x, y)$ are in one-to-one correspondence, and corresponding terms have the same coefficient. Thus $g(u, v)$ is also of the form

$$g(u, v) = \sum_{k\ell} p_{k\ell}(\overline{a}) u^k v^\ell \ .$$

The polynomial $g(0, y)$ has at least two nonzero terms. Let $g(0, y) = y^k h(y)$ where $h(y)$ is not divisible by y. In step 8, we factored $h(y)$ and extended the tree with a separate branch for each root, but here we will instead do a squarefree decomposition of $h(y)$ and extend the tree with a branch for each factor $q_{n+1}(y)$.

For each such branch, a new indeterminate a_{n+1} is created and $q_{n+1}(a_{n+1})$ is added to the list $q_0 \ldots, q_n$. Then we reparametrize symbolically with respect to the map $(u, v) = (x, y - a)$ to get

$$\begin{aligned} g(u, v + a) &= \sum_{ij} p_{ij}(\overline{a}) u^i (v + a)^j \\ &= \sum_{ij} p_{ij}(\overline{a}) u^i (\sum_{k=0}^{j} \binom{j}{k} v^k a^{j-k}) \\ &= \sum_{ij} p_{ij}(\overline{a}) (\sum_{k=0}^{j} \binom{j}{k} u^i v^k a^{j-k}) \\ &= \sum_{ik} (\sum_{j \geq k} \binom{j}{k} p_{ij}(\overline{a}) a^{j-k}) u^i v^k \ . \end{aligned}$$

The new coefficients are

$$p'_{ik}(\overline{a}, a) = \sum_{j \geq k} \binom{j}{k} p_{ij}(\overline{a}) a^{j-k}$$

which are small. We reduce these modulo the q_i and test the signs (zero or nonzero) of the new coefficients $p'_{ik}(\overline{a}, a)$ modulo the q_i. As above, this may cause further splitting of $\overline{q}$ in case the signs are not uniquely determined. When done, we again have the property that no coefficient of $g(u, v + a)$ vanishes at any point of the variety $V(\overline{q})$.

5 Analysis

The tree is of small depth, since at each vertex, either there are multiple children; or $m > 1$, in which case we get ramification (the maximum is n, and they are multiplicative along edges); or the degree of the discriminant strictly decreases [2, Theorem 15.1]. The number of branches is bounded by n, since there are at most n places in all.

There are at most as many q_i at any vertex as the depth of the vertex. The splitting of the tower (4) may cascade, but each such split uses only subresultant computations and takes polynomial time. There are at most n splits since each one creates a new branch of the tree. The product of the degrees of the q_i is at most n, since each sequence of roots (3) represented by the q_i determines a distinct place.

Acknowledgements

We thank Len Adelman, Ming-Deh Huang, Doug Ierardi, Susan Landau, John Little, Paul Pedersen, Moss Sweedler, Barry Trager, Emil Volcheck, Peter Walsh,

and especially Richard Zippel for their help. The support of the National Science Foundation under grant CCR-9317320 and the U.S. Army Research Office through ACSyAM, a branch of the Mathematical Sciences Institute of Cornell University under contract DAAL03-91-C-0027 is gratefully acknowledged.

References

1. M. Ben-Or, D. Kozen, and J. Reif, *The complexity of elementary algebra and geometry*, J. Comput. Syst. Sci., 32 (1985), pp. 251–264. Invited special issue.
2. G. A. Bliss, *Algebraic Functions*, Amer. Math. Soc., 1933.
3. C. Chevalley, *Introduction to the Theory of Algebraic Functions of One Variable*, American Mathematical Society, 1951.
4. A. L. Chistov, *Polynomial complexity of the Newton-Puiseux algorithm*, in Math. Found. Comput. Sci., vol. 233 of Lect. Notes. Comput. Sci., Springer, 1986, pp. 247–255.
5. C. Dicrescenzo and D. Duval, *Computations on curves*, vol. 174 of Lect. Notes in Comput. Sci., Springer, 1984, pp. 100–107.
6. ———, *Algebraic computations on algebraic numbers*, in Informatique et Calcul, Wiley-Masson, 1985, pp. 54–61.
7. W. Fulton, *Algebraic Curves*, Addison Wesley, 1989.
8. D. Ierardi and D. Kozen, *Parallel resultant computation*, in Synthesis of Parallel Algorithms, J. Reif, ed., Morgan Kaufmann, 1993, pp. 679–720.
9. S. Landau, *Factoring polynomials over algebraic number fields*, SIAM J. Comput., 14 (1985), pp. 184–195.
10. S. Lang, *Introduction to Algebraic and Abelian Functions*, Springer-Verlag, second ed., 1972.
11. A. K. Lenstra, *Factoring polynomials over algebraic number fields*, in Proc. EuroCal 1983, vol. 162 of Lect. Notes in Comput. Sci., Springer, 1983, pp. 245–254.
12. A. K. Lenstra, H. W. Lenstra, and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann., 261 (1982), pp. 515–534.
13. J. Teitelbaum, *The computational complexity of the resolution of plane curve singularities*, Math. Comp., 54 (1990), pp. 797–837.
14. B. M. Trager. Personal communication, 1994.
15. B. M. Trager, *Integration of Algebraic Functions*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, September 1984.
16. P. G. Walsh, *The Computation of Puiseux Expansions and a Quantitative Version of Runge's Theorem on Diophantine Equations*, PhD thesis, University of Waterloo, 1994.
17. R. Zippel. Personal communication, 1994.