J. Symbolic Computation (1995) **11**, 1–000

An algorithm for computing an integral basis in an algebraic function field

Mark van Hoeij

Department of mathematics

University of Nijmegen

6525 ED Nijmegen

The Netherlands

e-mail: hoeij@sci.kun.nl

(Received 6 September 1993)

Algorithms for computing integral bases of an algebraic function field are implemented in some computer algebra systems. They are used e.g. for the integration of algebraic functions. The method used by Maple 5.2 and AXIOM is given by Trager in [Trager,1984]. He adapted an algorithm of Ford and Zassenhaus [Ford,1978], that computes the ring of integers in an algebraic number field, to the case of a function field.

It turns out that using algebraic geometry one can write a faster algorithm. The method we will give is based on Puiseux expansions. One can see this as a variant on the Coates' algorithm as it is described in [Davenport,1981]. Some difficulties in computing with Puiseux expansions can be avoided using a sharp bound for the number of terms required which will be given in Section 3. In Section 5 we derive which denominator is needed in the integral basis. Using this result 'intermediate expression swell' can be avoided.

The Puiseux expansions generally introduce algebraic extensions. These extensions will not appear in the resulting integral basis.

1. Definitions and notations

The following conventions are used in the rest of the paper:

- L is an algebraically closed field of characteristic 0.
- x is transcendental over L.
- y is algebraic over $L(x)$ with minimal polynomial f. We assume that y is integral over $L[x]$. This means that f is not only a monic polynomial over $L(x)$ but over $L[x]$. The case where y is not integral over $L[x]$ can be reduced to this case after multiplying y by an element of $L[x]$.
- K is a subfield of L that contains all coefficients of f.
- $\overline{K}$ is the algebraic closure of K in L.
- n is the degree of f.
- $\overline{L[x]}$ is the integral closure of $L[x]$ in the algebraic function field $L(x, y)$.

0747–7171/90/000000 + 00 \$03.00/0

- An integral basis is a sequence $b_0, \ldots, b_{n-1}$ of elements of $L(x,y)$ such that $\overline{L[x]} = L[x]b_0 + \cdots + L[x]b_{n-1}$

We want to find an integral basis of $\overline{L[x]}$, which has all its coefficients in K. More precisely: we want an integral basis $b_0, \ldots, b_{n-1}$ where $b_i \in K(x,y)$. For example L can be $\overline{\mathbb{Q}}$ or $\mathbb{C}$ in the test examples of Section 6. K must contain the coefficients field of f, i.e. the smallest field that contains the coefficients of f. We can take $K = \mathbb{Q}$ in these examples. We want an integral basis for $L(x,y)$ which has coefficients in K. However, the Puiseux expansions that we will compute usually require algebraic extensions of K.

Because we assumed y to be integral over $L[x]$ we know that every element of $L[x,y]$ is integral:

$$L[x,y] \subset \overline{L[x]} \subset L(x,y).$$

From now on we will view elements of $L(x,y)$ as polynomials in y of degree $< n$. Then we can uniquely define the degree of an element of $L(x,y)$. Whenever we write $a(p)$ for an element $a \in L(x,y)$ we will mean a with p substituted for y.

There exists an integral basis $b_0, \ldots, b_{n-1}$ of the following form:

$$b_i \in K(x,y) \text{ and } \mathrm{degree}(b_i) = i, \qquad i = 0 \ldots n-1.$$

For the proof that b_i can be chosen in $K(x,y)$ we refer to [Trager,1984]. The condition $\mathrm{degree}(b_i) = i$ can be satisfied as follows. If $b_i = \sum_{j=0}^{n-1} b_{ij} y^j$ then (b_{ij}) is a matrix over $K(x)$. This matrix can be transformed to a triangular matrix (b'_{ij}) by elementary matrix operations over $K[x]$ because $K[x]$ is a principal ideal domain. Then we define for $i = 0 \ldots n-1$

$$b'_i \stackrel{\mathrm{def}}{=} \sum_{j=0}^{n-1} b'_{i,j} y^j.$$

The sequence $b'_0, \ldots, b'_{n-1}$ is an integral basis which satisfies $\mathrm{degree}(b'_i) = i$. Our algorithm will compute an integral basis in this form.

We start with $b_0 = 1$ in our algorithm. The problem we have to solve now is the following. Given a number $0 < d < n$ and a sequence $b_0, \ldots, b_{d-1}$ such that

$$L[x]b_0 + \cdots + L[x]b_{d-1} = \{a \in \overline{L[x]} \mid \mathrm{degree}(a) < d\}$$

and

$$b_i \in K(x,y) \text{ and } \mathrm{degree}(b_i) = i, \qquad i = 0 \ldots d-1,$$

compute an element $b_d \in K(x,y)$ such that

$$L[x]b_0 + \cdots + L[x]b_d = \{a \in \overline{L[x]} \mid \mathrm{degree}(a) \leq d\} \text{ and } \mathrm{degree}(b_d) = d.$$

Once this problem is solved we can find an integral basis $b_0, b_1, \ldots, b_{n-1}$ step by step.

2. Computing b_d

2.1. Outline of the algorithm

Our approach for finding an element b_d satisfying the above conditions is the following:

1 Our first guess is: $b_d = yb_{d-1}$. We could also have chosen y^d, but yb_{d-1} results in a faster algorithm. Define

$$V = \{a \in \overline{L[x]} |\ \mathrm{degree}(a) \leq d\} \setminus L[x]b_0 + \cdots + L[x]b_d.$$

2 As long as $V \neq \emptyset$ we repeat

(a) Choose $a \in V$ such that a can be written as

$$a = \frac{a_0 b_0 + \cdots + a_d b_d}{k}$$

with $a_0, \ldots, a_d, k$ in $K[x]$ and $a_d = 1$.

(b) Since $a_d = 1$ we have

$$L[x]b_0 + \cdots + L[x]b_d \subset L[x]b_0 + \cdots + L[x]b_{d-1} + L[x]a \subset \overline{L[x]}.$$

So we replace b_d by a. Now the $L[x]$-module $L[x]b_0 + \cdots + L[x]b_d$ is a bigger part of the set of integral elements with degree $\leq d$. So the set V is smaller now.

To complete this algorithm for computing b_d there remain three problems:

We must prove that this algorithm ends.
We must prove that a can be chosen in the desired form if $V \neq \emptyset$.
We need to find a way to decide if $V \neq \emptyset$ and if so to compute $a_0, \ldots, a_{d-1}$.

2.2. Proof that the algorithm terminates

The following proof can be found in [Trager,1984], but we will repeat it here because the discriminant plays an important role in our algorithm. We will use the discriminant only for the field extension $L(x) \subset L(x, y)$ in this paper.

Definition 2.1. *The* discriminant $\mathrm{disc}(e_1, \ldots, e_n)$ *of n elements $e_1, \ldots, e_n \in L(x, y)$ is defined as:*

$$\begin{vmatrix} \mathrm{Tr}(e_1 e_1) & \ldots & \mathrm{Tr}(e_1 e_n) \\ \vdots & & \vdots \\ \mathrm{Tr}(e_n e_1) & \ldots & \mathrm{Tr}(e_n e_n) \end{vmatrix}$$

where Tr *stands for the trace map of the extension $L(x) \subset L(x, y)$ of degree n.*

The discriminant is an element of $L(x)$. If $e_1, \ldots, e_n \in \overline{L[x]}$, then the discriminant is an element of $L[x]$. We can compute the discriminant of $y^0, \ldots, y^{n-1}$ using the following relation

$$D \stackrel{\mathrm{def}}{=} \mathrm{disc}(y^0, \ldots, y^{n-1}) = \mathrm{Res}_y(f, \frac{\partial f}{\partial y}) \in L[x]$$

In fact $D \in K[x]$ because all coefficients of f are elements of K. For more information about the discriminant see any book on algebraic number theory.

Each time b_d is replaced by

$$\frac{a_0 b_0 + \cdots + a_d b_d}{k}$$

with $a_d = 1$, we know that $\mathrm{disc}(b_0, \ldots, b_d, y^{d+1}, \ldots, y^{n-1}) \in K[x]$ changes by a factor $1/k^2$. So the degree of this discriminant decreases in every loop. Hence the algorithm must end.

In every loop $\mathrm{disc}(b_0, \ldots, b_d, y^{d+1}, \ldots, y^{n-1})$ divides $\mathrm{disc}(y^0, \ldots, y^{n-1})$ because $L[x]y^0 + \cdots + L[x]y^{n-1} \subset L[x]b_0 + \cdots + L[x]b_d + L[x]y^{d+1} + \cdots + L[x]y^{n-1}$. Therefore the only possible factors in the denominator of b_d are those factors k for which $k^2|D$.

Note that the discriminant D is useful for finding the singularities of an algebraic curve. If in a singularity $x = \alpha$ holds, then $(x - \alpha)^2$ divides the discriminant.

2.3. Proof that a can be chosen in the desired form

Suppose that $V \neq \emptyset$ in the algorithm of Section 2.1. Choose $a \in V$. By induction we may assume that $L[x]b_0 + \cdots + L[x]b_{d-1}$ contains all integral elements of degree $< d$. Hence $\mathrm{degree}(a) = d$. Furthermore

$$a = \frac{a_0 b_0 + \cdots + a_d b_d}{k'} \in L(x, y)$$

for some $a_0, \ldots, a_d$ and k' in $L[x]$, because $b_0, \ldots, b_d$ is a basis for the $L(x)$-vectorspace of elements of $L(x, y)$ of degree $\leq d$. Note that $k' \notin L$.

After multiplying a by an element of $L[x]$ we may assume that k' is irreducible in $L[x]$ and still $a \in V$. So we can write

$$k' = x - \alpha$$

for some $\alpha \in L$. Then $(x - \alpha)^2$ must divide D. Hence α is algebraic over K. If we add a multiple of $x - \alpha$ to one of the a_i, then a remains integral, because $b_0, \ldots, b_d$ are all integral. Hence we may assume that all a_i are elements of L. We know that $a_d \neq 0$ because the $\mathrm{degree}(a) = d$. After multiplying a by $1/a_d$ we get $a_d = 1$, and still $a \in V$.

Given $\alpha \in \overline{K}$ we want to prove that there is at most one sequence $a_0, \ldots, a_d \in L$ such that $(a_0 b_0 + \cdots + a_d b_d)/(x - \alpha) \in \overline{L[x]}$ and $a_d = 1$. Suppose we had two such sequences, then the difference of the two corresponding integral elements would be an integral element of degree $< d$ and not an element of $L[x]b_0 + \cdots + L[x]b_{d-1}$. This is a contradiction.

Suppose $(a_0 b_0 + \cdots + a_d b_d)/(x - \alpha)$ is integral and $a_d = 1$. Now we want to prove that $a_i \in K(\alpha)$. If a_i were trancendental over $K(\alpha, a_0, \ldots, a_{i-1}, a_{i+1}, \ldots, a_d)$ then we would get a sequence with the same property by substituting any number in K for a_i. This is a contradiction because there is at most one such sequence. Hence all a_i are algebraic over $K(\alpha)$. We see that a_i must be in $K(\alpha)$, otherwise the conjugates of a_i would form other sequences with the same property.

Let $k \in K[x]$ be the minimal polynomial of α over K. So $k'|k$. Then we have $K[x]/(k) \simeq K(\alpha)$. Elements of $K(\alpha)$ can be written as polynomials in α of degree $< \mathrm{degree}(k)$. Now we write each a_i as a polynomial in α, and then substitute x for α. This changes the a_i only by a multiple of $x - \alpha$. So a is still integral. Now a is an element of V, and the a_i in the numerator of a are elements of $K[x]$.

If P is a place (cf. next section) of the function field $L(x, y)$ we denote the discrete valuation in P by v_P. The fact that a is integral is equivalent to $v_P(a) \geq 0$ in all finite places P. So $v_P(a_0 b_0 + \cdots + a_d b_d) \geq v_P(x - \alpha)$ in all finite places P. Because $a_0 b_0 + \cdots + a_d b_d \in K(x, y)$ we have $v_P(a_0 b_0 + \cdots + a_d b_d) \geq v_P(x - \alpha_i)$ for all finite places where the α_i are the conjugates of α over K. We can conclude that $v_P(a_0 b_0 + \cdots + a_d b_d) \geq$

$v_P(k)$ in all finite places P. Hence

$$\frac{a_0b_0 + \cdots + a_db_d}{k} \in K(x,y)$$

is integral. Since it is not an element of $L[x]b_0 + \cdots + L[x]b_d$ it is an element of V. This finishes the proof.

2.4. Computing $a_0, \ldots, a_d$

In this Section we will, given an element $\alpha \in L$, describe a method to find elements $a_0, \ldots, a_d \in K(\alpha)$, $a_d = 1$ such that

$$\frac{a_0b_0 + \cdots + a_db_d}{x - \alpha}$$

is integral, assuming that such $a_0, \ldots, a_d$ exist. In Section 2.3 we showed how we can construct from this an integral element of the form

$$\frac{a_0b_0 + \cdots + a_db_d}{k}$$

with $a_0, \ldots, a_d, k \in K[x]$, k irreducible.

First we will summarize the facts about places and valuations that we will use. A place P of a field M is a discrete valuation ring of M. A place P has a valuation $v_P : M \to \mathbb{Z} \cup \{\infty\}$. We will consider only places of $L(x,y)$ in this paper. If $x \in P$ we will call P a finite place. An element $a \in L(x,y)$ is integral if and only if $v_P(a) \geq 0$ in all finite places P, in other words:

$$\bigcap_{P \text{ finite}} P = \overline{L[x]}.$$

A Puiseux series around $x = \alpha$ is an element of

$$\bigcup_{t=1}^{\infty} L((\ (x-\alpha)^{\frac{1}{t}}))$$

Since this is an algebraically closed field which contains $L(x)$ we know that f has n different zeros in this field. We will call these zeros the Puiseux expansions $p_1, \ldots, p_n$ of f around $x = \alpha$. We have a valuation

$$v : L((\ (x-\alpha)^{\frac{1}{t}})) \to \mathbb{Q} \cup \{\infty\}$$

given by $v((x-\alpha)^{1/t}) = 1/t$. Since p_i are zeros of f we have for each i an embedding

$$a \in L(x,y) \to a(p_i) \in L((\ (x-\alpha)^{\frac{1}{t}}))$$

for some integer t. This gives a valuation $v_i : L(x,y) \to \mathbb{Q} \cup \{\infty\}$, $v_i(a) = v(a(p_i))$.

Suppose P is a place above $x = \alpha$ (i.e. $v_P(x-\alpha) > 0$), and $t = v_P(x-\alpha)$. Then there are precisely t Puiseux expansions p_i that correspond to P, i.e. $v_i(a) \geq 0 \iff v_P(a) \geq 0$ for all $a \in L(x,y)$. For these t values of i we have $p_i \in L((\ (x-\alpha)^{\frac{1}{t}}))$ and $v_P(a) = tv_i(a)$.

Once we have computed the Puiseux expansions p_i around $x = \alpha$ we can check for given $a_0, \ldots, a_d \in K(\alpha)$ if $(a_0b_0 + \cdots + a_db_d)/(x-\alpha)$ is integral. To do that we need to check if $v_i((a_0b_0 + \cdots + a_db_d)/(x-\alpha) \geq 0$ for all i.

For undetermined $a_0, \ldots, a_d$ we can find a set of linear equations in $a_0, \ldots, a_d$ over $K(\alpha)$ equivalent to $(a_0b_0 + \cdots + a_db_d)/(x-\alpha)$ being integral by substituting for y all different Puiseux expansions p_i above $x = \alpha$ in $(a_0b_0 + \cdots + a_db_d)/(x-\alpha)$. Each Puiseux expansion p_i then gives a number of equations by taking the coefficients of the negative powers of $x - \alpha$ in $(a_0b_0 + \cdots + a_db_d)/(x-\alpha)$ with $y = p_i$ substituted. We know that $(a_0b_0 + \cdots + a_db_d)/(x-\alpha)$ is integral, if and only if all these coefficients of negative powers of $x - \alpha$ are zero, because that is exactly when $v_i((a_0b_0 + \cdots + a_db_d)/(x-\alpha)) \geq 0$.

If we add an extra equation $a_d = 1$ we know that there is at most one solution. To find it we need to solve a system of linear equations over $K(\alpha)$.

Now there remains one problem if we want a complete algorithm. The Puiseux expansions p_i usually have infinitely many terms. We can compute only a finite number of terms. Hence, we need to know how many terms of the Puiseux expansions must be computed in order to determine whether an element is integral.

3. Bounds for the Puiseux expansions

In this Section $p_1, \ldots, p_n$ are the n different Puiseux expansions around $x = \alpha$. Also $r_1, \ldots, r_n$ are n different elements of $\bigcup_{t=1}^{\infty} L((\,(x-\alpha)^{\frac{1}{t}}))$. In our algorithm $r_1, \ldots, r_n$ will be approximations of $p_1, \ldots, p_n$, i.e. r_i is a truncation of p_i. Instead of substituting $y = p_i$ we will substitute $y = r_i$ in our algorithm. Now we want to know how many of the lowest terms the r_i must have in order to be sure that our algorithm produces the right answer. More precisely: we want bounds $N_1, \ldots, N_n \in \mathbb{Q}$ and we determine r_i such that

$$v(r_i - p_i) \geq N_i.$$

So we determine p_i modulo $(x-\alpha)^{N_i}$. These numbers N_i must be chosen in such a way that the following holds for all $G \in L[x,y]$ with $\mathrm{degree}_y(G) < n$.

$$\forall_{M \in \mathbb{N}} \left[(\forall_i v(G(r_i)) \geq M) \iff (\forall_i v(G(p_i)) \geq M) \right].$$

Then we can use r_i instead of p_i as we will show now. Suppose a is an element of $L(x,y)$ and we want to check if a has nonnegative valuations in the finite places above $x = \alpha$. We can multiply out all factors in the denominator of a except $x - \alpha$. Hence we may assume that for some $G \in L[x,y]$ with degree in y less than n, we can write

$$a = \frac{G}{(x-\alpha)^M}.$$

This has nonnegative valuations if and only if $\forall_i v(a(p_i)) \geq 0$, if and only if $\forall_i v(G(p_i)) \geq M$, if and only if $\forall_i v(G(r_i)) \geq M$. So this can be checked using r_i instead of p_i.

To derive conditions for $N_1, \ldots, N_n$ we introduce an undetermined integer M and a polynomial $G \in L[x,y]$. We may assume that $\mathrm{degree}_y(G) < n$. We will assume

$$\forall_i v(G(r_i)) \geq M$$

and we want to choose the numbers $N_1, \ldots, N_n$ in such a way that we can conclude

$$\forall_i v(G(p_i)) \geq M.$$

Then we are finished, because the same argument can be used to conclude $\forall_i v(G(r_i)) \geq M$ from $\forall_i v(G(p_i)) \geq M$.

We have assumed that all r_i are different. Then we know that the polynomial G is completely determined by the values $G(r_i)$, because the degree of G is smaller than n.

One has

$$G = \sum_i G(r_i)\mathrm{Ind}_i$$

where Ind_i stands for indication polynomial, a function which equals 1 in r_i and 0 in all r_j, $j \neq i$.

$$\mathrm{Ind}_i = \prod_{j \neq i} \frac{y - r_j}{r_i - r_j}.$$

Since $v(G(r_i)) \geq M$ for all i, we can conclude

$$v(G(p_j)) = v\left(\sum_i G(r_i)\mathrm{Ind}_i(p_j)\right) \geq M$$

once we have $v(\mathrm{Ind}_i(p_j)) \geq 0$. So the following is sufficient

$$\forall_{i,j} v(\prod_{k \neq i} \frac{p_j - r_k}{r_i - r_k}) \geq 0.$$

Note that this condition is independent of the number M. This fact is useful in our algorithm. Now we can choose the rational numbers $N_1, \ldots, N_n$ in the following way.

$$N_i > \max_{k \neq i} v(p_i - p_k) + \max_k(\mathrm{Int}_k) - \mathrm{Int}_i$$

where

$$\mathrm{Int}_i := \sum_{k \neq i} v(p_i - p_k).$$

In particular we have $v(r_i - p_i) \geq N_i > \max_{k \neq i} v(p_i - p_k)$, which means that r_i is closer to p_i than all p_k, $k \neq i$ are. The valuation v gives a metric $(1/2)^v$, for which the strong triangle inequality holds. Using this, it is easy to see that from the last inequality follows: $v(r_i - p_i) > v(r_i - p_k)$ if $k \neq i$. So p_i is determined by r_i. It is the Puiseux expansion closest to r_i. If $i \neq j$ we also have $v(p_i - p_j) = v(r_i - r_j) = v(r_i - p_j)$. For $i \neq j$ we compute

$$\begin{aligned}
v(\textstyle\prod_{k \neq i} \frac{p_j - r_k}{r_i - r_k}) &= v(p_j - r_j) + \sum_{k \neq i, k \neq j} v(p_j - r_k) - \sum_{k \neq i} v(r_i - r_k) \\
&= v(p_j - r_j) + \sum_{k \neq i, k \neq j} v(p_j - p_k) - \sum_{k \neq i} v(p_i - p_k) \\
&= v(p_j - r_j) + \mathrm{Int}_j - v(p_j - p_i) - \mathrm{Int}_i \\
&\geq N_j + \mathrm{Int}_j - v(p_j - p_i) - \mathrm{Int}_i \\
&> \max_{k \neq j} v(p_j - p_k) + \max_k \mathrm{Int}_k - v(p_j - p_i) - \mathrm{Int}_i \\
&\geq 0.
\end{aligned}$$

For $i = j$ we have

$$v(\prod_{k \neq i} \frac{p_j - r_k}{r_i - r_k}) = \sum_{k \neq i} (v(p_i - r_k) - v(r_i - r_k)) = 0$$

So $v(\mathrm{Ind}_i(p_j)) \geq 0$ for all i and j, hence $v(G(p_j)) \geq M$ for all j.

4. A theoretical algorithm

We will summarize the results of the previous Sections in the following algorithm.

PROGRAM: integral basis

Input: a monic irreducible polynomial f over $L[x]$
Output: an integral basis $b_0, \ldots, b_{n-1}$

```
K := the coefficients field of f
n :=degree(f), the degree in y
df := the set of monic irreducible polynomials k ∈ K[x] for which k^2|Res_y(f, ∂f/∂y)
for k in df
    α_k := a zero of k in L
    r_{k,i} := p_i mod (x − α_k)^{N_i}, i = 1...n
    # where {p_1, ..., p_n} are the Puiseux expansions, and N_i are the rational
    # numbers satisfying the inequalities of the last Section
b_0 := 1
for d from 1 to n − 1
    b_d := y b_{d−1}
    for k in df
        found_something:=true
        while found_something do
            a := (a_0 b_0 + ... + a_d b_d)/(x − α_k)
            # where a_0, ..., a_{d−1} are indeterminates, and a_d = 1
            equations:= the set of coefficients of negative powers of x − α_k in
                a(r_{k,1}), ..., a(r_{k,n})
            solve the equations
            # We know there is at most one solution
            if there is a solution then
                assign the solution for a_i to a_i, i = 0..d − 1
                # We know that a_i ∈ K(α_k)
                a_i := a_i written as a polynomial in α_k
                a_i := a_i with x substituted for α_k
                b_d := (a_0 b_0 + ... + a_d b_d)/k
            else
                found_something:=false
```

RESULT: $b_0, \ldots, b_{n-1}$

5. A more efficient algorithm

The algorithm in Section 4 works, but we need to do more work to make it efficient. I will describe a few tricks that we used in our Maple implementation to make the algorithm faster. I did not include them in the algorithm in Section 4 in order to keep it readable. One trick is the following. When we compute $a(r_{k,1}), \ldots, a(r_{k,n})$ in the algorithm, we can use the previous computations, where $b_i(r_{k,1}), \ldots, b_i(r_{k,n})$, $i = 0..d-1$ was computed. The following is most important for the efficiency of our algorithm. So we will write in a theorem.

THEOREM 5.1. *Suppose $k \in K[x]$ is irreducible, $\alpha \in L$ is a zero of k and N is the rational number* $\max_i(\mathrm{Int}_i)$. *Then the largest multiplicity of k in the denominators in the integral basis is $\lfloor N \rfloor$. This bound is sharp.*

Using this, we can clear out the high powers of $x - \alpha$ during the computation. If we did not do so we would get a large intermediate expression swell, which greatly reduces the size of the input that the algorithm can handle. Using this bound we know precisely which terms we need and which ones we do not need.

PROOF. Suppose $p_1, \ldots, p_n$ are the Puiseux expansions above $x = \alpha$. The elements in our integral basis are of the following form: G/d where $d \in K[x]$ and $G \in K[x, y]$ is monic in y of degree $< n$. We may assume that the degree of G is $n-1$, if the degree is lower we can multiply G with a power of y. We can factor G if we compute the Puiseux expansions $q_1, \ldots, q_{n-1}$ of G

$$G = \prod_{i=1}^{n-1} (y - q_i).$$

For each q_i we choose a p_j which is closest to q_i, i.e. we choose j such that $v(p_i - q_j)$ is maximal. If we replace the q_i by the chosen p_j in the product $\prod(y - q_i)$, we see that the valuation of this product does not decrease in any of the finite places above $x = \alpha$. We may also assume that, without decreasing the valuation of the product in any of the finite places above $x = \alpha$, G does not have multiple factors. If $y - p_i$ appears more than once, we can replace a p_i by a p_j that does not yet appear in the product. p_j is chosen as close as possible to p_i. The valuation v_{p_i} of the product remains the same. It is equal to $v(0) = \infty$. One of the Puiseux expansions $p_1, \ldots, p_n$ will not appear in this product, since there are only $n-1$ factors. If p_i does not appear in the product, the valuation v_{p_i} of the product is Int_i. Hence the multiplicity of $x - \alpha$ in the denominator d is $\leq \lfloor \mathrm{Int}_i \rfloor$ if G/d is integral.

Now we want to show that this upperbound is sharp. Choose a number i such that Int_i is maximal. Let M be the number $\lfloor \mathrm{Int}_i \rfloor$. We want to construct an integral element of which the multiplicity of $x - \alpha$ in the denominator is M. First we compute

$$\frac{\prod_{j \neq i}(y - p_j)}{(x - \alpha)^M} \in L((\ (x - \alpha))^{1/t})[y]$$

for some integer t. This expression has valuation ≥ 0 for all v_{p_i}. But we can not use this expression since it is not an element of $L(x, y)$. Therefore we remove the powers of $x - \alpha$ higher than M in the numerator. So we get an element of $L((x - \alpha)^{1/t})[y]$. We remove the fractional powers of $x - \alpha$ by applying the trace map of $L((x - \alpha)^{1/t}) : L(x - \alpha)$. We know that the trace of an integral element is always integral. Suppose the result is G. Then the multiplicity of $x - \alpha$ in the denominator of G is M. Because G is integral, it must be an element of $L[x]b_0 + \cdots + L[x]b_{n-1}$. Hence the multiplicity of $x - \alpha$ is $\geq M$ in at least one of the denominators of b_i, and hence the multiplicity of k is $\geq M$. $\square$

6. Implementation

In my implementation I do not check whether f is irreducible over $L(x)$. However, f must be square-free, otherwise the discriminant is 0 and the algorithm will not work. If f is reducible and square-free the algorithm returns an answer. The result will be a basis for the integral closure of $L[x]$ in $L(x)[y]/(f)$. I will only sketch the proof. Let L_α denote the algebraic closure of $L((x-\alpha))$, $L_\alpha = \bigcup_k L((\ (x-\alpha)^{1/k}))$. The set of integral elements of L_α is $\bigcup_k L[[\ (x-\alpha)^{1/k}]]$. There are precisely n $L(x)$ homomorphisms from $\sigma_{i,\alpha} : L(x)[y]/(f) \to L_\alpha$, which correspond to the zeros of f in L_α. Now an element $a \in L(x)[y]/(f)$ is integral precisely when the images $\sigma_{i,\alpha}(a) \in L_\alpha$ are integral for all i and for all $\alpha \in L$. The algorithm searches elements with this property. Hence the algorithm still works for square-free reducible f. Recall that the algorithm only needs to check those α for which $(x-\alpha)^2$ divides the discriminant. The $\sigma_{i,\alpha}(a)$ are determined by substituting the n Puiseux expansions for y in a.

My implementation is part of the Maple 5.2 share library. It is available by anonymous ftp at neptune.inf.ethz.ch under the name IntBasis in the directory /maple/5.2/share.

6.1. A few tests

I have tried the following tests

$$\begin{gathered}
f_1 = y^{10} + (-2494x^2 + 474)y^8 + (84366 + 2042158x^4 - 660492x^2)y^6 \\
+(128361096x^4 - 47970216x^2 + 6697080 - 761328152x^6)y^4 + \\
(-12024807786x^4 - 506101284x^2 + 15052058268x^6 + 202172841 + 134266087241x^8)y^2 \\
+34263110700x^4 - 228715574724x^6 + 5431439286x^2 + 201803238 \\
-9127158539954x^{10} - 3212722859346x^8,
\end{gathered}$$

$$\begin{aligned}
f_2 &= y^{20} + y^{13}x + x^4y^5 + x^3(x+1)^2, \\
f_3 &= y^{30} + y^{13}x + x^4y^5 + x^3(x+1)^2, \\
f_4 &= y^{40} + y^{13}x + x^4y^5 + x^3(x+1)^2.
\end{aligned}$$

The curve f_1 was found by Rob Koelman. It is an example of a curve of degree 10 with 26 cusps. For a curve of degree 10 this is the largest number of cusps according to the Plücker formulas. It has no other singularities, so the genus is 10.

On a 64Mb workstation none of these examples can be run by the integral basis algorithm included in Maple 5.2 and AXIOM. Using my implementation, however, the integral basis computation for these examples can be done on a 386 40 MHz PC with 4Mb of memory, using Maple 5.2 under DOS. The computation on the PC takes 306, 60, 95 and 228 seconds, resp.

6.2. Other integral bases

I also tried to implement this method in the following two cases. However, the implementation is incomplete. For both 1 and 2 the algorithm does not work in the case where wild ramifications occur. Implementations will become available once we have a solution for this case.

1 A function field $K(x,y)$ where y is algebraic over K, and K is a field with prime characteristic.
2 An algebraic number field $Q(y)$ where y is algebraic over the rational numbers.

For algebraic number fields we can use p-adic numbers instead of Puiseux expansions. I thank Michiel Frankenhuysen for this idea.

Computations with p-adic numbers can be performed fast. Especially because we can use the bound in Section 3. So for algebraic number fields without wild ramifications the algorithm is very efficient, The bottleneck for this case is the factorization of the discriminant. The bound on the denominator in Section 5 also holds for function fields with prime characteristic and number fields. However, this bound becomes an upper bound in these cases, not an exact bound.

6.3. Efficient computation of Puiseux expansions

For an introduction in Puiseux expansions see [Bliss,1966]. A Puiseux expansion is a zero of f in the algebraic closure of $L((x))$. So there are precisely n different Puiseux expansions. We will briefly describe how to find the next term of a Puiseux expansion (or better: an approximation of a Puiseux expansion)

$$r = r_0 x^{e_0} + \ldots + r_k x^{e_k}$$

when $r_0, \ldots, r_{k-1}$ are already known. The exponent e_k can be found using the Newton polygon (cf. [Bliss,1966]). An equation for the coefficient r_k is given by the lowest coefficient, as a polynomial in x, of $f(r)$. To find $f(r)$ one should not substitute r for y in f, because r may be a large expression. Instead one substitutes $r_k x^{e_k - e_{k-1}} + r_{k-1}$ for r_{k-1} in the previous expression $f(r)$, in which r_{k-1} was still undetermined.

In order to avoid intermediate expression swell we must remove the unnecessary terms in $f(r)$. Once r is computed far enough to distinguish it from other Puiseux expansions we know precisely which terms are needed in $f(r)$. Because in that case we know the number Int, which we defined as the sum of the valuations of the differences of r with the other Puiseux expansions. If p is the exact Puiseux expansion which is the one closest to r than we have

$$v(r-p) + \mathrm{Int} = v(f(r)).$$

So we can clear out all powers of x in $f(r)$ higher than the sum of the desired accuracy and the number Int.

References

Bliss, G.A., (1966). *Algebraic Functions*, Dover

Davenport, James Harold, (1981). *On the integration of algebraic functions*, Springer-Verlag, p32-48

Ford, David James, (1978). *On the Computation of the Maximal Order in a Dedekind Domain*, Ph.D. thesis, Ohio State University, Dept. of Mathematics

Trager, Barry Marshall, (1984). *Integration of algebraic functions*, Ph.D. thesis, Dept. of EECS, Massachusetts Institute of Technology