

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/37600680>

Integration of algebraic functions

Article · January 1984

Source: OAI

CITATIONS

103

READS

2,551

1 author:



Barry M. Trager

IBM

58 PUBLICATIONS 1,741 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



IBM z196 [View project](#)



Computer Algebra [View project](#)

INTEGRATION OF ALGEBRAIC FUNCTIONS

by

BARRY MARSHALL TRAGER

B.S. Massachusetts Institute of Technology
(1973)

M.S. Massachusetts Institute of Technology
(1976)

Submitted to the Department of
Electrical Engineering and Computer Science
in Partial Fulfillment of the
Requirements of the
Degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1984

© Barry M. Trager 1984

The author hereby grants to M.I.T. permission to reproduce and to distribute copies of this thesis document in whole or in part.

Signature of Author Barry Marshall Trager
Department of Electrical Engineering and Computer Science, August 31, 1984

Certified by Joel Moss Thesis Supervisor

Certified by Richard E Zipp Thesis Supervisor

Accepted by [Signature]
Chairman, Departmental Committee on Graduate Students
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

ARCHIVES

OCT 04 1984

LIBRARIES

INTEGRATION OF ALGEBRAIC FUNCTIONS

by

BARRY MARSHALL TRAGER

Submitted to the Department of Electrical Engineering and Computer Science
on August 31, 1984 in partial fulfillment of the requirements of the
Degree of Doctor of Philosophy in Computer Science

ABSTRACT

This thesis will provide a practical decision procedure for the indefinite integration of algebraic functions in terms of elementary functions. An elementary function is one that can be expressed using combinations of algebraic functions, exponentials and logarithms. We can either express the integral as an elementary function or guarantee that this cannot be done. The algorithms presented here are "rational" in the sense that no algebraic extensions are made which are not necessary to express the answer. The basic approach involves reducing the singularities of the integrand iteratively. This has the advantage that even when the original problem is not integrable, it can be partially integrated and reduced to a simpler form. We are able to prove that any integral can be reduced to one whose finite poles are all simple. Elimination of simple poles requires the construction of logarithmic terms. Whether or not this can be done reduces to the question of whether some multiple of a given divisor is principal. We present a novel algorithm for solving this problem which is able to compute the exact order of the given divisor by examining its behavior under "reduction modulo p ." We also present algorithms for finding absolutely irreducible factors of multivariate polynomials and computing the genus of curves. The latter is a simple consequence of our use of an integral basis for our function field to provide a global non-singular model of the curve. This basis allows us to determine the nature of the singularities of the integrand and provides an ideal-theoretic algorithm to compute a function with a given principal divisor.

Thesis Supervisor: Dr. Joel Moses
Title: Professor of Computer Science and Engineering

Thesis Supervisor: Dr. Richard Zippel
Title: Assistant Professor of Computer Science and Engineering

ACKNOWLEDGEMENTS

I would like to thank Professor Joel Moses for introducing me to the field of Computer Algebra in general and symbolic integration in particular, Professor Steven Kleiman for several illuminating discussions, Dr. James Davenport for constructive criticism and conversations about topics of common interest and Dr. Richard Zippel for his helpful comments.

I would also like to thank the many friends who encouraged, cajoled, and browbeat me into finally finishing this dissertation. Especially Patrizia who actually got me to finish, David and Joel who tried everything they could think of, Anni who supplied me with a continuing source of interesting integrals to solve, and Dick and the others at IBM research who provided the necessary time, tools, and financial support.

CONTENTS

Chapter 1	Introduction	5
Section 1.	Review of previous work.	5
Section 2.	Outline of thesis.	12
Chapter 2	Integral Bases.	15
Section 1.	Radical of the Discriminant.	21
Section 2.	Computing the Idealizer	25
Section 3.	Normalize at infinity.	26
Section 4.	Genus computation.	29
Section 5.	Simple radical extensions.	29
Chapter 3	Absolute Irreducibility	32
Section 1.	Exact coefficient fields and Regular extensions	34
Section 2.	Algorithmic considerations.	36
Section 3.	Binomial polynomials.	38
Chapter 4	The Rational part of Integral	40
Section 1.	Rational function integration.	40
Section 2.	Algebraic functions.	43
Section 3.	Poles at infinity	48
Chapter 5	Log terms and Divisors	51
Section 1.	Properties of logarithmic differentials.	51
Section 2.	Computing the residues.	55
Section 3.	Constructing Divisors	59
Section 4.	Dealing with Branch places	62
Chapter 6	Principal Divisors and Points of Finite Order	64
Section 1.	Principal Divisors	64
Section 2.	Good Reduction.	65
Section 3.	Algorithms for Divisor reduction	67
Chapter 7	Conclusions and Future Research.	71
Appendix A.	Integration of simple radical extensions	73
Section 1.	Structure theorems	73
Section 2.	A Generalized Risch algorithm	74
Section 3.	Summary and Conclusions.	79
References	81

CHAPTER 1 INTRODUCTION

This thesis will provide a practical decision procedure for the indefinite integration of algebraic functions. Unless explicitly noted, we will always assume x to be our distinguished variable of integration. An algebraic function y of x is defined as a solution of a monic polynomial equation in y with coefficients that are rational functions in x . Each of these rational functions in x can be written as a quotient of polynomials in x whose coefficients are constants (i.e. not dependent on x). If we adjoin each of these constant coefficients to our base field Q (rational numbers), then we have constructed a finitely generated extension of Q that we will call our coefficient field, K . We can assume the integral has the form $\int y dx$ where $f(x,y)$ is the unique monic polynomial of least degree that y satisfies. An elementary function is one that can be expressed using combinations of algebraic functions, exponentials and logarithms. We propose either to express the integral as an elementary function or guarantee that this cannot be done.

SECTION 1. REVIEW OF PREVIOUS WORK

During the 18th and 19th centuries this problem attracted much interest [1]. Euler (1748) and others studied elliptic functions and discovered they were not integrable in terms of elementary functions. Abel (1826) first studied the integrals of general algebraic functions, which later became known as abelian integrals. Liouville [41] proved a key theorem that forms the basis for the decision procedure. He showed that $\int y dx$, if integrable in terms of elementary functions, could be expressed as

$$\int y dx = v_0 + \sum c_i \log v_i \quad (1)$$

where the c_i are constants and the v_i are rational combinations of x and y . The basic idea in his proof is that integration cannot introduce any new functions other than constant multiples of logs since differentiation would not remove them. Liouville (1833) also gave an integration algorithm [31] for the special case of algebraic integrals that could be expressed without logarithms, but he found no method for solving the general problem. This problem was so difficult that Hardy (1916) pessimistically stated [21] "there is reason to suppose that no such method can be given."

The next major steps in this area were taken by Risch (1968). First he gave a complete decision procedure for integrating elementary functions that were free of algebraic functions [38]. Then he turned his attention to the algebraic problem, and in [39] he sketched a procedure that reduced algebraic integration to a problem in algebraic geometry, determining a bound for the torsion subgroup of the divisor class group of an algebraic function field. Finally in [40] by referring to some recent results in algebraic geometry he outlined a theoretical solution to the problem. While this indeed disproved Hardy's undecidability conjecture, it did not really present a practical algorithm that could be used to actually solve integration problems.

Risch refined equation (1) by showing one could assume $c_i \in \bar{K}$ and $v_i \in \bar{K}(x,y)$ where \bar{K} is the algebraic closure of K (see also [42]). He also showed that $v_0 \in K(x,y)$. The integration problem is now reduced to finding the c_i and the v_i . The basic approach that Risch used and that we will use is to construct these functions by analyzing their singularities. By considering algebraic functions on their Riemann surface, they are no longer multi-valued and have only a finite number of poles as singularities. At each point of the Riemann surface a function can be expressed locally as a Laurent series in terms of some uniformizing parameter. When the parameter is expressed as $(x-a)^{1/n}$ for some $a \in K$ and $n \in \mathbb{N}$ the series is called a Puiseux expansion. The finite initial segment composed of the terms with negative exponents is called the principal part of the series expansion. If we compute the principal parts of the integrand

at all its poles, then we can integrate these sums termwise. If we discard the log terms that arise by integrating terms of the form $r_j (x-a)^{-1}$, then the remaining terms form the principal parts of the expansions of v_0 at all of its poles. Thus we make the key observation that the singular places of v_0 are a subset of the singular places of the integrand. Now we are in the position of trying to find a function with a given set of principal parts. Since an algebraic function with no poles is constant, this function is only determined up to an additive constant, as expected from an integration problem. The least degree term of each principal part specifies the order n_j of the pole at that place p_j on the Riemann surface. This allows us to associate an integer with each place if we agree to associate the integer 0 with places where the integrand was non-singular. The formal sum over all places $\sum n_j p_j$ is called a *divisor*, and by the Riemann-Roch theorem the set of algebraic functions with orders not less than those specified by the divisor form a finite dimensional vector space over \bar{K} . Bliss [6] gives a technique for finding a basis for this vector space. A general member of this vector space can be expressed as a linear combination of basis elements with indeterminate coefficients. By equating the principal parts of a general member of the vector space with the principal parts of the integral we get a system of linear equations that give a necessary and sufficient condition for the existence of the algebraic part of the integral.

To find the logarithmic part we now examine the terms of the form $r_j (x - a)^{-1}$ that we ignored earlier. The coefficients r_j are precisely the residues of the integrand at the singular places p_j of the integrand. If the integral exists then it can be presented in a form where the c_j 's form a basis for the \mathcal{Z} -module generated by the residues. In fact as we shall show in chapter 6 the c_j generate the minimal algebraic extension of the coefficient field required to express the integral. The algorithms we present will avoid Puiseux expansions and never introduce algebraic quantities not required for the final answer.

To simplify our presentation of Risch's approach, we will now assume all the residues are rational numbers. The general case will be treated in chapter 6. Let m be a least common

denominator of all the residues. We will now construct a divisor D from our set of residues, in general we would have k divisors where k is the rank of the \mathbb{Z} -module generated by the residues. We define the order of D at each place p to be m times the residue of the integrand at p . Thus $D = \sum (mr_i)p_i$. Since there are only a finite number of places where the integrand has non-zero residue D is a well defined divisor. Since the sum of the residues of the integrand is zero, $\deg D = \sum (mr_i) = 0$. For any function $f \in K(x,y)$ we have a divisor $\text{div}(f) = \sum (\text{ord}_p f)p$ where $\text{ord}_p f$ is the order of f at p . A divisor is called principal if it is the divisor of some algebraic function. If the divisor associated with the residues is principal describing a function f then we have determined the log term $\frac{1}{m} \log f$. However, if the divisor is not principal it is possible for some integer multiple of it to become principal. If k times a divisor D is principal with associated function v , then we have a log term $\frac{1}{mk} \log v$. This is the fundamental theoretical obstruction to the integration of algebraic functions. The integrand contains complete information about the location and orders of the poles of the algebraic part of the integral, but for the log parts the residues enable us to find reduced divisors only (the coefficients of the places are relatively prime), not necessarily the actual divisors of the logands. The "points of finite order problem" is to find an integer bound B such that if kD is not principal for some divisor D and all integers $1 \leq k \leq B$ then no integer multiple of D is principal. Under componentwise addition and subtraction of the order coefficients the set of divisors becomes an abelian group. The quotient of the group of divisors of degree zero by the subgroup of principal divisors is called the divisor class group. The divisor classes for which some multiple is principal form the torsion subgroup of the divisor class group. Thus we can reformulate our question as finding a bound for the torsion subgroup. In order for such a bound to exist it is critical that our constant field be a finitely generated extension of the rationals and not an algebraically closed field.

The approach that Risch outlined for determining the bound uses a technique that has lately come into vogue in many areas of algebraic manipulation. We take a difficult problem and homomorphically map it into a simpler domain, hoping to find some technique for lifting

the solution back to the original domain. Although the original *points of finite order* problem seems quite difficult, Weil (1948) showed that the problem is easily solvable for function fields in one variable over finite fields [52]. For these fields the divisor class group is finite and Weil's rationality formula for the zeta function shows that the order of the class group can be computed by counting points on a non-singular model. Using Weil's proof of the extended Riemann hypothesis for such fields we can explicitly give bounds for the divisor class group over a finite field. An upper bound is $(\sqrt{q} + 1)^{2g}$ where q is the order of the finite field and g is the genus, and we obtain a lower bound of $(\sqrt{q} - 1)^{2g}$.

The natural approach to reducing our problem is "reduction mod p ," where p is the prime ideal for some discrete valuation of our constant field, e.g. if the constant field is \mathcal{Q} then p will be a prime ideal in \mathcal{Z} . But we must verify when such a reduction is "good", i.e. gives us useful information for determining our bound. Risch observed [40] that if we get a projective non-singular model for our function field and if its reduction mod p is still non-singular then the reduction is good. This means that the induced homomorphism of divisor class groups is injective for all divisors whose order is relatively prime to the characteristic of the finite field ([44] or [47]). Since the torsion subgroup is a finite abelian group, it can be decomposed into a product of the group of elements whose orders are relatively prime to p and the p -Sylow subgroup. Thus if we find two distinct rational primes that give us good reduction, we can multiply the bounds for the two divisor class groups and get our desired bounds. Risch claimed that all of these steps were known to be effective, but he provided no explicit algorithms. Dwork and Baldassarri [4] independently duplicated Risch's approach to this problem in the process of finding algebraic solutions for second order linear differential equations.

James Davenport has independently investigated this problem. His algorithms are constructed along the lines suggested by Risch. He uses Puiseux expansions and Coate's [12] algorithm to construct a basis for the multiples of a divisor. This construction is used for both

the algebraic and transcendental part of the answer. To bound the divisor torsion he uses special purpose algorithms depending on whether the curve is genus 1 and whether there are parameters present. In the case there are parameters present, he produces an explicit test for torsion without computing the bound. In the case of genus 1, he uses arithmetic on the curve to compute the torsion. While these special purpose tests can be reasonably efficient, in the general case, he reverts to the Weil bound given previously. His algorithms are general, but his implementation is currently limited to algebraic functions that can be expressed with nested square roots.

Risch and Davenport depend on Puiseux expansions to unravel the singularities of the function field defined by the integrand. Some such mechanism is necessary to distinguish apparent poles from actual ones. For example the function y/x has an apparent pole at the origin, but if $y^2 = x^2(x + 1)$ then the function is actually holomorphic there. We will use integral bases to determine the nature of the poles of an algebraic function. An integral basis for an algebraic function field of degree n is a set of n functions such that an element of the function field can be expressed as a linear combination of basis elements with polynomial coefficients if and only if that element has no singularities in the finite plane, i.e. the element is an integral algebraic function. Good algorithms for computing integral bases are a subject of ongoing research [56]. Bliss [6] shows that finding integral bases is no harder than computing Puiseux expansions, and in the very important case of function fields defined by a single radical, they are immediate.

In this thesis we present a new algorithm for integration of algebraic functions. We rely on the construction of an integral basis to provide an affine non-singular model for the curve. This will allow us to construct the algebraic part of the answer by a generalization of Hermite's algorithm for integrating rational functions. This integral basis can also be used to test whether divisors are principal and to test for good reduction.

A very serious problem in symbolic calculations is that intermediate expressions are frequently much larger than the final answer and can thus be very time and space consuming. One of the counterparts to intermediate expression swell that we are forced to deal with in this problem is intermediate constant field extensions. Risch and Davenport make the simplifying assumption of an algebraically closed constant field. Making unnecessary extensions of the constant field greatly increases the cost of arithmetic. We will present an algorithm that will perform all its operations in the minimum extension field in which the answer can be expressed. Risch's more extensive use of Puiseux expansions forced him to operate in an extension field of much higher degree for his intermediate computations.

We will present a new algorithm for algebraic function integration that is strongly analogous to recent efficient algorithms for rational function integration [49]. Using integral bases to normalize the problem, we will be able to reduce the finding of the algebraic part of the integral to solving a set of linear equations. Finding the logarithmic part is indeed more difficult and does involve determining whether a given divisor is of finite order to guarantee termination of the algorithm. In addition to obtaining bounds that guarantee termination, we will present a novel algorithm for actually obtaining the logand associated with a divisor. Unlike earlier approaches that constructed this function from the divisor alone, we will use the integrand to create the ideal of functions that are multiples of the divisor at all finite places. Generators for this ideal can be derived almost by inspection, and there remains only to determine whether there is a principal generator.

Algebraic function integration is significantly more complicated than rational function integration since one can no longer depend on unique factorization. Ideals were created by Dedekind to restore unique factorization to algebraic number fields. Since that time they have been studied in increasingly abstract settings, to the point that their origins are almost forgotten. We intend to actually use ideals, as Dedekind intended 100 years ago, to combat the lack of unique factorization in algebraic function fields. In addition, our explicit genera-

tors for the ideal associated with a divisor will permit us to reduce the generators and compute the exact order of this divisor mod p . Thus instead of working with bounds for the torsion, we can compute the exact order of torsion divisors. We then need only to test whether or not this particular power is principal, instead of testing all powers up to some calculated bound. This will be shown to lead to a very practical decision procedure for algebraic function integration.

Finally we will investigate the possibility of extending this procedure to include exponentials and logarithms enabling one to determine whether the integral of any elementary function can be expressed as an elementary function.

SECTION 2. OUTLINE OF THESIS

Chapter Two will present an algorithm for computing an integral basis for our function field. This is essentially the same algorithm presented by Ford [20] for algebraic number fields, with proofs of validity in this more general situation and extended to normalize the basis elements at infinity. This fundamental construction will be used throughout the thesis and effectively provides us with an affine non-singular model for our function field. It enables us to determine the poles of our integrand and characterize the form of the answer. We will also use this integral basis to help find principal generators for divisors and test for “good reduction”.

Since the running time of many of our algorithms depend critically on the degree of the defining relationship for our function field, it is very useful to guarantee that our defining polynomial is irreducible over the algebraic closure of our coefficient domain, i.e. is absolutely irreducible. In Chapter Three we present a new algorithm for finding an absolutely irreducible factor of a multivariate polynomial that seems to be significantly better than other known approaches. As pointed out by Duval [18], the number of absolutely irreducible factors of the defining polynomial is the same as the dimension of the vector space of functions that have no

poles. We can compute this simply using our normalized integral basis, and then we only need to perform our factorization algorithm when it is known to yield a lower degree factor. Since an irreducible polynomial will usually be absolutely irreducible, this can prevent a lot of wasted effort.

Armed with a minimal defining polynomial and an integral basis we are ready to find the purely algebraic part of the integral. Chapter Four proceeds by analogy with the standard approaches to rational function integration. It shows how Hermite's algorithm can be generalized to deal with algebraic functions. This approach will always succeed in reducing the integral to one with only simple finite poles and perhaps poles at infinity. In fact we are able to show that if the original problem had no poles at infinity, and after removing the algebraic part we introduce poles at infinity, then the original problem was not integrable. This approach has the advantage of allowing one to obtain canonical reduced form even for problems that are not integrable. In this stage only linear congruence equations are solved and no new algebraic numbers are generated. It is difficult to remove poles at infinity by a Hermite-like method, so the original integral is transformed by a simple change of variables so that there are no poles at infinity. This simplifying transformation is one of the problems to be overcome in trying to generalize this algorithm to handle mixed transcendental as well as algebraic extensions.

If the original problem was integrable, the remaining simple finite poles must be canceled by the derivative of a linear combination of logarithmic terms. In the rational function case these log terms can be found by factorization or by computing gcd's of polynomials. Unfortunately algebraic function fields are not unique factorization domains so this approach cannot be used. As described earlier we will use the residues of the integrand to construct divisors associated with each log term. In Chapter Five we present an algorithm that computes a polynomial whose roots are all the residues using resultants. This is an extension of the idea we used in [48] for rational functions and does not use power series expansions or extensions

of the coefficient domain. Given this polynomial we finally need to extend our coefficient domain to include its splitting field. We show that this is the minimal extension required for expressing the integral. After computing a Z -linear basis for the residues, we construct a model for the divisor associated with each basis element. Our construction of the model appears new, and provides us with a simple construction to find a principal generator if one exists.

Finally in Chapter Six we are ready to address the "points of finite order problem". We need to know for each of the divisors we have constructed, whether there is some power of it that is principal. As discussed in Risch [40] and Davenport [14], we will use the technique of "gcd reduction" to solve this problem. However our explicit representation of a divisor will allow us to compute the order of individual divisors exactly instead of merely computing a bound on the orders of all divisors. While we do perform a sequence of tests for principality on powers of divisors, these tests are all performed over finite constant fields and thus much less expensive than testing successive powers of a divisor over our original coefficient field. After computing what should be the order of our original divisor if it were finite, we merely perform a single test for that power of our divisor over the original function field. Both of these improvements should make a substantial difference in running time.

In Chapter Seven we summarize our contributions and suggest ways to extend the work done here. In an appendix we present one step toward a complete algorithm for handling both transcendental and algebraic extensions.

CHAPTER 2 INTEGRAL BASES

In later chapters we will be very concerned with the problem of creating functions with prescribed singularities. It will be very useful to be able to recognize and generate functions whose only poles lie at places over infinity. Such functions are called *integral algebraic functions*. In the field $K(x)$ these are simply polynomials in x , i.e. a rational function has a finite pole if and only if it has a nontrivial denominator. Any function that is algebraic over $K(x)$ satisfies a unique monic irreducible polynomial with coefficients in $K(x)$.

$$Z^m + a_1 Z^{m-1} + \dots + a_m \tag{1}$$

Such a function is integral over $K[x]$ if and only if the coefficients are in fact in $K[x]$, i.e. polynomials in x . In the rest of this thesis we will abbreviate "integral over $K[x]$ " to "integral".

Let $K(x,y)$ be a finite algebraic extension of degree n over $K(x)$, then the integral functions form a free module of rank n over $K[x]$, i.e. any such function can be written as linear combination of n basis functions with coefficients that are polynomials in x . Such a basis is called an *integral basis*. If we allow the coefficients to be rational functions in x then these same n functions comprise a vector space basis for $K(x,y)$ over $K(x)$. Thus each element of $K(x,y)$ has a unique representation in terms of a particular integral basis, and has no finite poles if and only if each coefficient is a polynomial, i.e. no denominators. In this chapter we present an algorithm for computing an integral

allow the coefficients to be rational functions in x then these same n functions comprise a vector space basis for $K(x,y)$ over $K(x)$.

One technique for finding such a basis is given in [Bliss]. What we call an integral basis, he would term *multiples except at infinity of the divisor 1*. His basic technique involves Puiseux expansions. We wish to avoid performing such expansions for two reasons: (1) A large

amount of code is required (2) Many algebraic numbers need to be introduced to compute Puiseux expansions even though none of them are actually required to express the final basis elements. While Puiseux expansions are very useful in their own right, we don't really need them and choose to avoid the time cost of doing unnecessary algebraic number computations and the space cost of all that additional code.

The algorithm we present here is based on work by Zassenhaus and Ford [20]. They were primarily interested in the case of algebraic number fields, but their algorithm also applies to function fields in one variable. In fact since we will always assume the characteristic of K is zero or greater than n , the algorithm can be somewhat simplified. The algorithm presented here is a generalization to function fields of the first of the two algorithms Ford presents. We chose to use this one since it is much simpler and it avoids fully factoring the discriminant.

We are given $K(x,y)$ where K is a computable field, x is a distinguished transcendental element, and $f(x,y)$ an irreducible separable polynomial of degree n over $K[x]$. Without loss of generality we can also assume f monic. If not then let $\hat{y} = ay$ where a is the leading coefficient, then \hat{y} satisfies a monic polynomial and generates the same function field. The elements of $K(x,y)$ that are integral over $K[x]$ form a ring called the integral closure of $K[x]$ in $K(x,y)$. As noted above this ring is also a free module of rank n . Since y is integral over $K[x]$, and the sum or product of integral elements are integral, $[1, y, \dots, y^{n-1}]$ constitutes a basis for an integral $K[x]$ module. This is our first "approximation" to an integral basis. Each iteration of the algorithm will produce a basis for a strictly larger integral $K[x]$ module until the integral closure is reached.

One important measure of the relative sizes of full (i.e. rank n) sub-modules of the integral closure is given by the discriminant. Let $[w_1, \dots, w_n]$ be n elements of $K(x,y)$. Since $K(x,y)$ is a separable algebraic extension of $K(x)$ of degree n , there are n distinct embeddings

σ_i into a given algebraic closure. The images of a element under these mappings are called the conjugates of that element. The conjugate matrix of $\bar{w} = [w_1, \dots, w_n]$ is defined by

$$M_{\bar{w}} = \begin{bmatrix} \sigma_1(w_1) & \cdots & \sigma_n(w_1) \\ \vdots & & \vdots \\ \sigma_1(w_n) & \cdots & \sigma_n(w_n) \end{bmatrix} \quad (2)$$

The *discriminant* of \bar{w} is the square of the determinant of the conjugate matrix. The discriminant is non-zero if and only if the w_i generate a full module (i.e. they are linearly independent). We can define the trace (*sp*) of an element $w \in K(x, y)$ as $sp(w) = \sum \sigma_i(w)$. Since this is a symmetric function of the conjugates, it is an element of $K(x)$. If we re-express the discriminant as the determinant of the product of the conjugate matrix and its transpose, we see that the product entries are traces of products of the original matrix entries. Thus the discriminant could be defined as $determinant(sp(w_i w_j))$.

$$SP_{\bar{w}} = \begin{bmatrix} \sigma_1(w_1) & \cdots & \sigma_n(w_1) \\ \vdots & & \vdots \\ \sigma_1(w_n) & \cdots & \sigma_n(w_n) \end{bmatrix} \begin{bmatrix} \sigma_1(w_1) & \cdots & \sigma_1(w_n) \\ \vdots & & \vdots \\ \sigma_n(w_1) & \cdots & \sigma_n(w_n) \end{bmatrix} = \begin{bmatrix} sp(w_1^2) & \cdots & sp(w_1 w_n) \\ \vdots & & \vdots \\ sp(w_n w_1) & \cdots & sp(w_n^2) \end{bmatrix} \quad (3)$$

If the w_i 's are integral functions then their traces are polynomials, and thus the discriminant is a polynomial. If $\bar{v} = [v_1, \dots, v_n]$ is a basis for a full module that contains \bar{w} then each w_i can be written as a polynomial combination of the v_i , i.e. $\bar{w} = A\bar{v}$ where the change of basis matrix A is an $n \times n$ matrix of polynomials. Thus the conjugate matrix $M_{\bar{w}} = A.M_{\bar{v}}$ and $Disc(\bar{w}) = \det(A)^2 Disc(\bar{v})$. \bar{w} and \bar{v} generate the same module if and only if A is invertible as a matrix over $K[x]$, i.e. $\det(A) \in K$. If \bar{v} strictly contains \bar{w} then $\det(A)$ is a polynomial p of nonzero degree and $Disc(\bar{v}) = Disc(\bar{w})/p^2$ thus each time we are able to produce a strictly larger $K[x]$ module, we eliminate a squared factor from the discriminant and the process can only continue for finitely many steps.

We will now state and prove the key algebraic result on which the algorithm is based. Let R be a *principal ideal domain*, in our case $K[x]$ while Ford and Zassenhaus assume $R = \mathbb{Z}$. Let V be a domain that is a finite integral extension of R . Then V is also a free module of

rank equal to the degree of $QF(V)$ (the quotient field of V) over $QF(R)$. Let $\bar{v} = [v_1, \dots, v_n]$ be a basis for V over R . The discriminant of \bar{v} generates an ideal in R that we will call the discriminant of V over R . The discriminant of any other basis of V over R differs by the square of a unit and thus generates the same ideal. If m is an ideal in a ring S , we define the idealizer $Id(m)$ to be the set of all $u \in QF(S)$ such that $um \subseteq m$. $Id(m)$ clearly contains S and is the largest ring in which m is still an ideal.

Theorem: [1] Under the above conditions V is integrally closed if and only if the idealizer of every prime ideal containing the discriminant equals V .

We will break this into a succession of small lemmas.

Lemma: [1] If m is a nonzero ideal in V the idealizer of m is integral over R .

Proof: Since V is a finite R -algebra m is finitely generated over R . Let m_1, \dots, m_k span m over R and let $u \in QF(V)$ such that $um \subseteq m$. Then $um_i = \sum_j r_{ij}m_j$ with $r_{ij} \in R$. Let M be the matrix $r_{ij} - \delta_{ij}u$ where δ_{ij} is the Kronecker index. Then M annihilates the vector $[m_1, \dots, m_k]$, and if we multiply by the adjoint of M we see that $\det M$ kills each of m_i and thus annihilates m . Since V is an integral domain, $\det(M)$ must be zero, but this gives a monic polynomial over R that u satisfies, and hence u is integral over R . \square

Next we define the *inverse* of an ideal. If m is an ideal in V then m^{-1} is the set of all $u \in QF(V)$ such that $um \subseteq V$. This notion is similar to the idealizer. The idealizer of an ideal is the subset of the quotient field that sends an ideal into itself, while the inverse of an ideal sends it into the ring V . Having defined the inverse of an ideal we will say that an ideal m is *invertible* if $mm^{-1} = V$. By definition we have $mm^{-1} \subseteq V$. If V is integrally closed then it is a Dedekind domain and has the property that every non-zero ideal is invertible. By proposition 9.7 and 9.8 in [2] p.97 we have

Lemma: [2] If all the non-zero prime ideals of V are invertible then V is integrally closed.

Lemma: [3] Any prime ideal not containing the discriminant of V over R is invertible.

Proof: Let m be a prime ideal not containing the discriminant of V over R . Then if we localize at m the discriminant becomes a unit and thus the local ring is integrally closed and locally m is invertible, in fact principal. If we localize at any other maximal ideal of V m contains units and is its own inverse. Since the localization of m at each prime ideal is invertible then by proposition 9.6 [2] m is invertible. \square

Corollary: [1] If V is not integrally closed there is some prime ideal containing the discriminant that is not invertible.

Proof: By lemma 2 there is some prime ideal that is not invertible and by lemma 3 it must contain the discriminant. \square

The following lemma is proved in [27] p. 607.

Lemma: [4] If V properly contains an ideal m then m^{-1} properly contains V .

Now we are ready to finish the proof of theorem 1.

Proof: If V is integrally closed the idealizer of any non-zero ideal equals V by lemma 1. If V is not integrally closed there is some prime ideal m that contains the discriminant but is not invertible. Thus $m^{-1}m$ is an ideal containing m but properly contained in V . Since m is maximal we must have $m^{-1}m = m$. Thus in this case $m^{-1} = Id(m)$. By lemma 4 we have the idealizer of m properly contains V . \square

Using theorem 1 we could compute the integral closure by computing the idealizer of the finitely many ideals that contain the discriminant. Either the result in each case will be V in which case V must be integrally closed, or we will find a ring strictly larger than V that is integral over V . This can happen only a finite number of times since each such ring will

remove a squared non-unit factor from the discriminant. We will improve this idea by dealing with all the ideals dividing the discriminant at the same time. First we observe the following property of idealizers:

Lemma: [5] If m and n are ideals then the idealizer of the product contains the idealizer of either ideal.

Proof: Any element of mn is of the form $\sum m_i n_i$ where $m_i \in m$ and $n_i \in n$. If $u \in Id(m)$ then $u m_i \in m$ and thus $u m_i n_i \in mn$. \square

Next we must introduce the notion of the radical of an ideal. The radical of an ideal $m \in V$ is the set of all $u \in V$ such that some power of u is in m and can also be characterized as the intersection of all prime ideals containing m . Our algorithm is based on the following corollary to theorem 1.

Corollary: [2] V is integrally closed if and only if the idealizer of the radical of the discriminant equals V .

Proof: Since all the non-zero prime ideals of V are maximal, the radical of the discriminant is also the product of all prime ideals containing the discriminant. By theorem 1 if V is not integrally closed the idealizer of one of these primes must be strictly larger than V . By lemma 5 the idealizer of the radical contains the idealizer of that prime and thus must also strictly contain V . If V is integrally closed again the idealizer of any ideal must equal V . \square

Thus our algorithm for computing the integral closure of V is:

1. find the radical of the discriminant of V over R
2. compute the idealizer \hat{V} of that radical.
3. If \hat{V} is strictly larger than V then set V to \hat{V} and go to step (1)

4. return V

Although this algorithm works, there are a few optimizations possible. The first time through $V = R[\alpha]$ where α satisfies equation (1). To compute the discriminant in this case, one simply computes the resultant of equation (1) and its derivative. Prime factors of the discriminant that appear only to the first power can be ignored. When returning from step 3 to step 1 one only needs to concentrate on the factors of the discriminant that have actually been reduced in the previous iteration. Since if there is some p whose p -radical is invertible, it will stay that way throughout the rest of the computation. This leads us to the following improved version of the algorithm for computing the integral closure of $V = R[\alpha]$ over R assuming $f(X)$ is the minimal equation of integral dependence for α .

0. Let $d = \text{Resultant}(f, f')$ and $k = d$

1. Let $q = \prod p_i$ such that p_i is prime, $p_i | k$, and $p_i^2 | d$. If q is a unit then return V .

2. Find $J_q(V)$, the radical of (q) in V

3. Find \hat{V} , the idealizer of $J_q(V)$ along with M the change of basis matrix from \hat{V} to V .

4. Let k be the determinant of M . If k is a unit then return V

5. Set $d = d/k^2$ and $V = \hat{V}$ and go to 1.

SECTION 1. RADICAL OF THE DISCRIMINANT

The discriminant is a principal ideal generated by some element d of R . We wish to compute the radical of the ideal d generates in V . Since R is a principal ideal domain it is also a unique factorization domain. Let (p_1, \dots, p_k) be the distinct prime factors of d in R . Since the radical of (d) is intersection of the prime ideals containing d , it is also the intersection of the radicals of the p_i . Let us therefore consider how to compute the radical in V of a principal

ideal generated by a prime element p of R . Following Ford we call such an ideal the p -radical of V .

u is in the p -radical if and only if the coefficients a_i in equation (1) of the monic minimal polynomial for u over R are divisible by p . ([2] Proposition 5.14 and 5.15) This gives us a membership test, but we are looking for ideal generators. Zassenhaus and Ford observed that under appropriate conditions the trace map from V to R provides us with linear constraints on the members of the p -radical. For any u in V the degree of u over R must divide the rank of V over R that is also the degree of $QF(V)$ over $QF(R)$. If m is the degree of u over R then $sp(u) = -(n/m)a_1$. Thus if $u \in p$ -radical then p divides $sp(u)$. Again following Ford we define the p -trace-radical as the set of $u \in V$ such that for all $w \in V$, $p \mid sp(uw)$. This leads us to the following lemma:

Lemma: [6] p -radical $\subseteq p$ -trace-radical

Proof: If u is in the p -radical then for any $w \in V$, uw is in the p -radical. But then the previous argument shows that $sp(uw)$ is divisible by p . \square

Now we find conditions under which the two sets of lemma 6 are the same. If w is in the p -trace-radical then $sp(w^k) \equiv 0 \pmod{p}$ for all $k > 0$. Let m be the degree of w over R . Then $R[w]$ is a free R -module of rank m dividing n , the rank of V over R . There is a reduced trace map from $R[w]$ to R denoted by sp_w satisfying $(n/m)sp_w(u) = sp(u)$ for any $u \in R[w]$. If $n/m \notin (p)$ then p dividing $sp(u)$ implies p divides $sp_w(u)$ and thus the p -trace-radical of $R[w]$ equals the intersection of the p -trace-radical of V with $R[w]$. If n/m is zero in $R/(p)$ then the characteristic of $R/(p)$ must divide n . To avoid this problem, we now make the assumption that the characteristic of $R/(p)$ is greater than n the rank of V over R . In our application where $R = K[x]$ the characteristic of $R/(p)$ is the same as the characteristic of K , and we shall see that the restriction will not cause any problems. In Ford and Zassenhaus' case where $R = \mathbb{Z}$, $p \in \mathbb{Z}$ and the characteristic of $R/(p)$ equals p . Thus when the discriminant has small

prime factors, our assumption would be invalid, so they compute the p -radical for small values of p using the kernel of powers of the Frobenius automorphism instead of the trace map.

Assume w is in the p -trace-radical of $R[w]$. Then $sp_w(w^k)$ is divisible by p for all $k > 0$. We wish to relate the traces of powers of w to the coefficients of the minimal polynomial for w in equation (1). This is provided by Newton's identities ([25] p. 203). If we let $s_k = sp_w(w^k)$ then for $1 \leq k \leq n$ we have:

$$s_k + a_1 s_{k-1} + \dots + a_{k-1} s_1 = -ka_k \quad (4)$$

Thus $a_1 = -s_1$ is divisible by p and by induction assume a_i is divisible by p for all $i < k \leq n$. Then the left hand side of equation (4) is divisible by p and as long as the characteristic of $R/(p)$ is greater than n we can divide by k and a_k must also be divisible by p . Since the coefficients of its minimal polynomial are divisible by p , w must be in the p -radical of V , which proves the following partial converse to lemma 6.

Theorem: [2] If the characteristic of $R/(p)$ is greater than the rank of V over R , then the p -trace-radical of V equals the p -radical of V .

Subsection 1.1. Computing the p -trace-radical

Let $[w_1, \dots, w_n]$ be a basis for V over R . The p -trace-radical $J_p(V)$ was defined as the set of $u \in V$ such that $p \mid sp(uw)$ for all $w \in V$.

$$sp(uw) \equiv 0 \pmod{p} \text{ for all } w \in V \iff$$

$$sp(uw_i) \equiv 0 \pmod{p} \text{ for } 1 \leq i \leq n \iff$$

$$\sum_{j=1}^n u_j sp(w_j w_i) \equiv 0 \pmod{p} \text{ for } 1 \leq i \leq n$$

Using the trace matrix $SP_{\overline{w}}$ defined by equation (3) we can write this as:

$$SP_{\bar{w}} \cdot \bar{u} = \begin{bmatrix} sp(w_1^2) & \cdots & sp(w_1 w_n) \\ \vdots & & \vdots \\ sp(w_n w_1) & \cdots & sp(w_n^2) \end{bmatrix} \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} \in pR^n \quad (5)$$

where pR^n represents the set of vectors of length n whose elements are divisible by p . $J_p(V)$ is determined by the solutions to (5) with $u_i \in R$. We actually wish to compute $J_q(V)$ where q is a certain product of distinct primes dividing the discriminant.

$$J_q(V) = \bigcap_{p_i | q} J_{p_i}(V)$$

Thus we want to find all $\bar{u} \in R^n$ such that the left side of equation (5) is in fact in qR^n . We need to add some equations to (5) to guarantee that each element u_i of our solution vectors must lie in R as opposed to $QF(R)$. Let I_n be the $n \times n$ identity matrix. Then $\bar{u} \in R^n$ if and only if $qI_n \cdot \bar{u} \in qR^n$. Thus if we let M_q be the vertical concatenation of $SP_{\bar{w}}$ and qI_n , $J_q(V)$ is the set of all $u \in QF(V)$ such that $M_q \cdot \bar{u} \in qR^{2n}$. If we left multiply M_q by an invertible R -matrix, the R -module of solutions remains unchanged. Invertible R -matrices are called unimodular, and are characterized by having determinants that are units in R . Since R is a principal ideal domain, there is some unimodular matrix that converts M_q into an upper triangular matrix. ([37] Theorem II.2 gives a constructive proof of this). This process allows us to reduce the $2n$ relations imposed by M_q and equation (5) to an equivalent set of n independent relations. Once this is done, we invert the square matrix determined by the first n rows of the reduced M_q . The columns of this inverse matrix provide a basis for the solutions to equation (5).

In our application we have the even stronger restriction that R is a euclidean domain. This allows us to triangularize M_q by elementary row operations only. This process is called Hermitian row reduction and is somewhat analogous to gaussian elimination that is used for matrices over fields. With gaussian elimination any nonzero element can be used to zero out its entire column. With hermitian row reduction one can only multiply by elements of R and a nonzero element can reduce the other members of its column to be smaller than it. Since R

is a euclidean domain this process can only continue for finitely many steps before we find an element that divides everyone else in its column. Finally this element can be used to clear the column. Let d be the size function associated with R , for $R = K[x]$ we use the degree, and for $R = \mathbb{Z}$ we use absolute value. To simplify the presentation of the following algorithm we define $d(0) = \infty$. Let $nrows$ and $ncols$ be respectively the numbers of rows and columns of a given matrix M . We assume that $nrows \geq ncols$ and M has rank $ncols$ in the following algorithm for hermitian row reduction.

1. loop for $j = 1$ thru $ncols$
 2. choose k such that $d(M_{kj})$ is minimal for $j \leq k \leq nrows$.
 3. exchange rows j and k
 4. loop for $i = j + 1$ thru $nrows$
 5. let q be the polynomial part of M_{ij}/M_{jj}
 6. replace M_i with $M_i - qM_j$
 7. if there is some $M_{ij} \neq 0$ for $j < i \leq nrows$ then go to 2.
8. return M

SECTION 2. COMPUTING THE IDEALIZER

Given (m_1, \dots, m_n) that form an R -basis for an ideal m in V , we wish to compute an R -basis for the idealizer of m . The idealizer of m was defined as the set of $u \in QF(V)$ such that $um \subseteq m$. This concept is very similar to the inverse ideal of m , m^{-1} , which was defined as the set of $u \in QF(V)$ such that $um \subseteq V$. Although we don't need to compute inverses to find an integral basis we will need them in chapter V, and we present both algorithms here to display their similarities. The inverse procedure will be given first since it is slightly simpler.

We assume $\bar{v} = (v_1, \dots, v_n)$ forms a basis for V over R that we hold fixed throughout this section. $u \in m^{-1}$ if and only if $um_i = \sum r_{ij}v_j$ with $r_{ij} \in R$ for $1 \leq i \leq n$. Multiplication by m_i is a

linear transformation on V . Let M_i represent multiplication by m_i with respect to our fixed choice of basis \bar{v} . (for details on constructing M_i see [48]) Since \bar{v} also forms a basis for $QF(V)$ over $QF(R)$, u can be represented as $\sum u_i v_i$ where $u_i \in QF(R)$. Then $M_i [u_1, \dots, u_n]^t$ yields the vector of coefficients of $u m_i$. $u \in m^{-1}$ if and only if the product coefficients lie in R for all i . Thus we are left with the linear algebra problem of finding a basis for R -module of vectors $[u_1, \dots, u_n]$ such that $M_i([u_1, \dots, u_n]^t)$ is a vector of elements of R for all i . Let M be the $n^2 \times n$ matrix that is the vertical concatenation of the M_i . Then we are looking for all vectors \bar{u} over $QF(R)$ such that $M\bar{u}$ is an n^2 vector of elements of R . By Hermitian row reduction we can zero out the last $n^2 - n$ rows of M , so we have reduced to an $n \times n$ matrix \hat{M} such that $M\bar{u} \in R^{n^2}$ if and only if $\hat{M}\bar{u} \in R^n$. The columns of \hat{M}^{-1} form a basis for m^{-1} .

Essentially the same approach will allow us to find the idealizer of m . Now we require that $u m_i = \sum r_{ij} m_j$ with $r_{ij} \in R$. The M_i still represent multiplication by m_i but now the input and output bases are different. The M_i necessary to compute the idealizer again take inputs expressed in terms of \bar{v} but give output vectors expressed in terms of the basis for m . Other than this one change the algorithm for idealizers is identical to the previous one for inverses. A summary of both algorithms follows:

1. Let M_i be the matrix representing multiplication by m_i with input base \bar{v} and output basis \bar{v} for computing inverses or \bar{m} for calculating the idealizer.
2. Let \hat{M} be the first n rows of the Hermitian row reduction of the vertical concatenation of the M_i .
3. Return the columns of \hat{M}^{-1} as the result expressed with respect to \bar{v}

Note the transpose of \hat{M}^{-1} is the change of basis matrix required in step 3 of the integral basis algorithm.

SECTION 3. NORMALIZE AT INFINITY

In this section we return to the special case $R = K[x]$. Having constructed an integral basis we can recognize finite poles of functions, but we also need to deal with singularities at "infinity". Given an arbitrary basis for a $K[x]$ module (an integral basis is a special case), we wish to minimize the sum of the orders of the basis elements at infinity. A characteristic property of an integral basis $[w_1, \dots, w_n]$ is that $\sum a_i(x)w_i$ is integral over $K[x]$ if and only if each $a_i(x)w_i$ is. In other words there is no cancellation of singularities from different summands. We would like our basis to have the same property with respect to the local ring of $K(x)$ at infinity and we will characterize this by saying that such a basis is *Normal at infinity*.

We will need the concept of a *local ring* at a place p of the function field $K(x)$. The local ring at p is defined as the set of functions in $K(x)$ that have no pole at p . If p is a finite place centered at $x = a$ then the local ring at p consists of rational functions whose denominators are not divisible by $x - a$. The order of a rational function at *infinity* is the degree of its denominator less the degree of its numerator. Thus the local ring of $K(x)$ at ∞ consists of those rational functions whose numerator degree does not exceed their denominator degree. A function in $K(x, y)$ is said to be integral over the local ring at p (for brevity integral at p) if it satisfies a monic polynomial with coefficients in the local ring at p . Analogous to the global integral basis, there exists a local integral basis at each place p of $K(x)$ such that all functions in $K(x, y)$ that are integral at p can be written as a linear combination of basis elements with coefficients in the local ring at p of $K(x)$. We will find it convenient to introduce the slightly weaker concept of a normal basis. $[w_1, \dots, w_n]$ is a normal basis at p if there exist rational functions $r_i \in K(x)$ such that $r_i w_i$ form a local integral basis at p . In other words there exists rational scaling factors that convert a normal basis into an integral basis.

Armed with this terminology we see that a basis is normal at infinity if and only if some rational multiple of the basis elements is a local integral basis at infinity. Let $[w_1, \dots, w_n]$ be a

basis for a $K[x]$ module. Assume we are given a local integral basis at infinity $[v_1, \dots, v_n]$. We wish to modify the original basis to make it normal at infinity without disturbing its basis properties everywhere else. We first represent the w_i in terms of the v_j :

$$w_i = \sum_{j=1}^n M_{ij} v_j \text{ where } M_{ij} \in K(x)$$

If w_i were a normal basis at infinity there would be rational functions r_i such that the change of basis matrix $r_i m_{ij}$ would have a determinant that was a unit in the local ring of $K(x)$ at infinity, i.e. a rational function whose numerator has the same degree as its denominator.

Define the representation order, $k(w_i)$ of w_i at infinity as the $\min_j \text{ord}_\infty M_{ij}$ for $1 \leq j \leq n$. We will initially choose $r_i = x^{-k(w_i)}$. This guarantees that $r_i w_i$ is integral at infinity and its representation order is 0. Note that the order at infinity of the determinant of M is always $\geq \sum k(w_i)$. We will show that our basis becomes normal when these numbers are equal. Let \hat{M} be the change of basis matrix for $r_i w_i$. Each row of \hat{M} is just r_i times the corresponding row of M . $r_i w_i$ is an integral basis if and only if the determinant of \hat{M} has order zero at infinity. Since this determinant is integral at infinity, this is equivalent to it having a non-zero value at infinity. Let N be a matrix where N_{ij} is the value of \hat{M}_{ij} at infinity. Since taking determinants commutes with evaluation, the determinant of N equals the value of the determinant of \hat{M} at infinity. Thus $r_i w_i$ is a local integral basis at infinity if and only if N has nonzero determinant. If the determinant of N is zero then there are a set of constants $c_i \in K$ such that $\sum_{i=1}^n c_i N_{ij} = 0$ for $1 \leq j \leq n$. Let $i_0 = i$ such that $c_i \neq 0$ and $k(w_i)$ is minimal. Define

$$\hat{w}_{i_0} = \sum_{i=1}^n c_i x^{k(w_{i_0}) - k(w_i)} w_i$$

Then replacing w_{i_0} by \hat{w}_{i_0} still yields a global integral basis. Similarly replacing $\hat{M}_{i_0 j}$ by $\sum c_i \hat{M}_{ij}$ yields a row whose orders are strictly positive. Thus the representation order $k(\hat{w}_{i_0})$ is strictly greater than $k(w_{i_0})$. The order of the determinant of the change of basis matrix is preserved by our new basis. After a finite number of such steps this order will be equal to the

sum of the representation orders of our basis elements and the basis will be normal.

We have shown how to make an arbitrary basis normal at infinity given a local integral basis at infinity. We next show how to compute the latter. If we replace $x = 1/z$ then infinity gets transformed to zero in the z - space. In order to compute a local integral basis at zero, we can use the algorithm of the previous section after making one optimization. In the local ring at 0 any polynomial not divisible by z is a unit. Thus we can replace the discriminant by the maximal power of z dividing it. After computing this local integral basis, we merely substitute $1/x$ for z and obtain a local integral basis at infinity.

SECTION 4. GENUS COMPUTATION

Our integral basis algorithm will also afford us an easy way to compute the genus of a function field. The global discriminant divisor of $K(x,y)$ over $K(x)$ is the product of the local discriminant divisors over each place of $K(x)$. As a by-product of our integral basis computation, we have computed two discriminants, a polynomial $disc_{finite}(x)$ that is the product of the discriminants over all finite places of $K(x)$, and a monomial in $z = 1/x$, $disc_{\infty}(1/x)$ that is the local discriminant at infinity. The degree of the discriminant divisor of our function field $K(x,y)$ over $K(x)$ is the sum of degree of $disc_{finite}(x)$ and the order at infinity of $disc_{\infty}(1/x)$. If we call this total discriminant degree d and we assume that K is the exact constant field of $K(x,y)$, then the genus of our function field $K(x,y)$ can be computed by the following formula given in [19] p. 134:

$$g = d/2 - [K(x,y):K(x)] + 1$$

SECTION 5. SIMPLE RADICAL EXTENSIONS

If F is a field with y algebraic over F of degree n such that $y^n \in F$, then we will call $F(y)$ a simple radical extension of F . If the characteristic of F is relatively prime to n , then extending F if necessary we can assume that F contains ω a primitive n^{th} root of unity. There is a unique differential automorphism of $F(y)$ over F such that $\sigma(y) = \omega y$. Define the operator

$$T_i = \frac{1}{n} \sum_{j=0}^{n-1} \frac{\sigma^j}{\omega^{ij}}$$

Note that $T_i(y^j) = y^j$ if $i = j$ else 0. Thus letting $g = \sum g^i y^i$ with $g^i \in F$, we have $T_i(g) = g^i y^i$. Since σ sends integral functions to integral functions, and sums and products of integral functions are integral, we have that the operators T_i map integral functions to integral functions. If g is an integral function this shows that each $g^i y^i$ must also be integral, which means that the basis y^i is normal everywhere proving the following:

Proposition: If $K(x,y)$ is a simple radical extension of $K(x)$ of degree n relatively prime to the characteristic, then the natural basis, $1, y, \dots, y^{n-1}$ is normal everywhere.

Without loss of generality we can assume that y satisfies the following equation:

$$y^n = \prod_{i=0}^{n-1} p_i$$

where $p_i \in K[x]$ and has no repeated factors. Thus to convert our natural basis into an integral basis we have to find polynomials $d_i(x)$ of maximal degree such that $y^i/d_i(x)$ is integral. Raising this expression to the n th power this implies that $\prod p_j^{ij}/d_i^n \in K[x]$. It is easy to show that the maximal $d_i(x)$ is the following:

$$d_i = \prod_{j=0}^{n-1} p_j^{[ij/n]}$$

Thus the following functions provide an integral basis for our simple radical extension:

$$\frac{y^i}{a_i(x)}$$

CHAPTER 3 ABSOLUTE IRREDUCIBILITY

We have assumed that the defining polynomial $f(x,y)$ for the integrand y is irreducible over $K(x)$, but during the integration process we may need to extend our coefficient field K , and over the extended field f may no longer be irreducible. In fact, in order to compute our bound for the "points of finite order", we need to be able to guarantee that f will remain irreducible after any algebraic extension of our coefficient field and is thus said to be absolutely irreducible. Another difficulty is the precise determination of our coefficient field. Initially we defined our coefficient field K to be the minimal extension of \mathcal{Q} necessary to express the defining polynomial. Any element of our function field that is algebraic over k could also have been considered part of the coefficient field. For example, if $f(x,y) = y^4 - 2x^2$ then y^2/x is a $\sqrt{2}$ and thus algebraic over \mathcal{Q} . Note that once we adjoin $\sqrt{2}$ to K , f is no longer irreducible and y satisfies a polynomial of degree 2 over this extended coefficient field. It will be advantageous to make our coefficient field as large as possible since that will decrease the degree of our function field and thus speed up our computation time that is strongly dependent on this degree. We now define the *exact coefficient field* K^o of $K(x,y)$ to be the set of all elements of $K(x,y)$ that are algebraic over K , also called the relative algebraic closure of K in $K(x,y)$. From the previous example the existence of elements of $K(x,y)$ that are algebraic over K seems connected with the question of the absolute irreducibility of $f(x,y)$. In the next section we will prove that this is indeed the case and in fact the process for finding an absolutely irreducible polynomial for y will lead us to discover the true coefficient field of $K(x,y)$.

As the previous example seems rather contrived, one might be led to suspect that defining polynomials that are irreducible but not absolutely irreducible are quite rare in practice. This is indeed the case, however the integral basis computation from the previous chapter can be used to perform a quick test for absolute irreducibility. The integral basis for that example is $1, y,$

$y^2/x, y^3/x$. Note that the first and third of these functions have no poles anywhere and are thus have divisor (1). As observed by Duval in [18] the dimension of the multiples of the divisor (1) is the number of absolutely irreducible components of the function field. Since we have already computed an integral basis that is normal at infinity, we merely need to test how many of our basis elements have no poles at infinity. If the answer is one, we can skip the algorithm derived in this chapter, since we are guaranteed that our defining polynomial is absolutely irreducible.

In this chapter we will permit K to be a perfect field of arbitrary characteristic. As a consequence if F is any finite algebraic extension of K , it can be generated by a single element and is thus a simple extension of K , $F = K(\alpha)$. Assuming that $f(x,y)$ is irreducible over K , we will present a new algorithm for performing absolutely irreducible factorizations, i.e. factorization over the algebraic closure of K , \bar{K} . The initial difficulty is that all current algebraic factoring algorithms only operate over a finitely generated field and the algebraic closure of K is not finitely generated. Thus we must find some subfield of \bar{K} that is finitely generated and is sufficient for performing the factorization. Risch showed the problem was decidable in [38] p. 178. His approach was to convert a multivariate polynomial to a univariate one by the Kronecker substitution ([50] p. 135).

$$f(x_1, x_2, \dots, x_v) \mapsto f(t, t^d, \dots, t^{d^{v-1}})$$

The key property of this substitution is that different power products of the x_i go into different powers of t assuming d is chosen larger than the degree of any variable appearing in f . Risch argued that the splitting field of this univariate polynomial suffices for the factorization. Assuming the original polynomial had v variables each with maximum degree d then his splitting field can be an algebraic extension of degree $d^v!$, whereas the algorithm presented below can test for absolute irreducibility or find an absolutely irreducible factor by operating over an extension of degree d_{\min} , the minimum of the degrees of all the variables. By

examining the algebraic structure of the fields involved, we will also discover some quick tests for absolute irreducibility.

SECTION 1. EXACT COEFFICIENT FIELDS AND REGULAR EXTENSIONS

First we will need some purely algebraic results (see also [45] pp 194-198). All fields are assumed to be contained in some universal field. If E and F are fields then EF is the compositum that is the field generated by $E \cup F$. If K is a finite algebraic extension of k then $[K:k]$ denotes the degree of this extension. From the previous section we see that determining the exact coefficient field requires us to find the relative algebraic closure of one field in another. The next lemmas give some properties of such fields.

Lemma: [1] If x is transcendental over k then k is algebraically closed in $k(x)$.

Proof: Let y be an element of $k(x)$ that is not in k . y can be written as $u(x)/v(x)$ with $u, v \in k[x]$. Then x satisfies the polynomial $P(X) = u(X) - v(X)y$. If P is not identically zero, this implies x is algebraic over $k(y)$. Let $u(X) = \sum u_i X^i$ and $v(X) = \sum v_i X^i$ and choose j such that $v_j \neq 0$. If P were identically zero then $u_j - yv_j = 0$, but since $u_j, v_j \in k$ this implies $y \in k$ contrary to the assumptions. Thus x is algebraic over $k(y)$ and y cannot be algebraic over k . \square

Lemma: [2] Let $K \supseteq k$ be fields with k algebraically closed in K and $k(\alpha)$ a simple algebraic extension of k . Then $[K(\alpha):K] = [k(\alpha):k]$.

Proof: Any factor of the monic minimal polynomial for α over k has coefficients that are polynomials in the conjugates of α and thus algebraic over k . If these coefficients were in K they would also be in k since the latter is algebraically closed in the former. \square

Corollary: [1] Let x be transcendental over k and F be a simple algebraic extension of k , then $[F:k] = [F(x):k(x)]$.

We are looking for a defining polynomial for y that is irreducible over \bar{K} . Factoring the bivariate polynomial f over \bar{K} is exactly the same as factoring it over $\bar{K}(x)$ by Gauss's lemma. If a perfect field k is algebraically closed in K , then K is said to be a *regular* extension of k . Thus the search for an absolutely irreducible defining polynomial is equivalent to finding a presentation of our function field as a regular extension of the coefficient field. We have the following key theorem that connects our twin problems of absolute irreducibility and exact coefficient fields.

Theorem: [1] If $f(x,y)$ is irreducible over a perfect field k then it is absolutely irreducible if and only if k is algebraically closed in $k(x,y)$, i.e. $k = k^o$

Proof: f is absolutely irreducible if and only if it is irreducible over any finite algebraic extension F of k . We wish to prove that f is irreducible over any such F if and only if k is algebraically closed in $k(x,y)$.

$$[F(x,y):k(x)] = [F(x,y):F(x)][F(x):k(x)] = [F(x,y):k(x,y)][k(x,y):k(x)] \quad (1)$$

By corollary 1 $[F(x):k(x)] = [F:k]$ and using equation (1) we have:

$$[F(x,y):F(x)] = [k(x,y):k(x)] \iff [F:k] = [F(x,y):k(x,y)] \quad (2)$$

$f(x,y)$ is irreducible over F if and only if $[F(x,y):F(x)] = [k(x,y):k(x)]$, and using equation (2) we have f is absolutely irreducible if and only if $[F:k] = [F(x,y):k(x,y)]$ for any finite algebraic extension F of k .

If k is algebraically closed in $k(x,y)$ then by lemma 1 $[F(x,y):k(x,y)] = [F:k]$ for any such F . Conversely if f is absolutely irreducible then choose F to be the algebraic closure of k in $k(x,y)$. Thus $[F(x,y):k(x,y)] = 1$ since $F \subseteq k(x,y)$. By equation (2) we have $[F:k] = 1$ showing that k is in fact algebraically closed in $k(x,y)$. \square

Corollary: [2] If $f(x,y)$ is irreducible over k^o , the algebraic closure of k in $k(x,y)$, then f is absolutely irreducible.

Proof: k^o is algebraically closed in $k^o(x,y) = k(x,y)$. Thus f is absolutely irreducible by the theorem. \square

Corollary 2 gives us a finitely generated field to factor over.

Note that we have also demonstrated that $[K^o:K]$ divides $[K(x,y):K(x)]$. Our choice of x as the independent variable and y as algebraic over x was somewhat arbitrary. If we reverse the roles we see that $[K^o:K]$ must divide both $\deg_y f$ and $\deg_x f$. So if these two numbers are relatively prime then f must be absolutely irreducible. More generally for an irreducible multivariate polynomial, if the gcd of all the degrees appearing in the polynomial is 1, then it must be absolutely irreducible.

SECTION 2. ALGORITHMIC CONSIDERATIONS

Armed with these algebraic results, we return to the question of finding an absolutely irreducible equation for y . We now realize that K^o is the field we want to factor over, but we have no explicit presentation of K^o . We know that $K^o \subseteq K(x,y)$ with each element of K^o algebraic over K . Since each element of K^o is independent of x , we can also view K^o as a subfield of the field $K(u,w)$ where w satisfies $f(u,w) = 0$. Again as in chapter II without loss of generality we will assume that $f(X,Y) \in K[X,Y]$ is monic as a polynomial in Y .

The minimal polynomial of each element of K^o is monic with coefficients in K . Thus $K^o[u]$ is certainly an integral algebraic extension of $K[u]$. Let A be the integral closure of $K[u]$ in $K(u,w)$. Any factorization of $f(x,y)$ with coefficients in K^o yields a factorization with coefficients in A . By the results of the last chapter any element of A can be written as a polynomial in u and w divided by the discriminant of $K[u,w]$ over $K[u]$. This discriminant is a polynomial in u and is the same as the discriminant of $f(u,w)$ viewed as a polynomial in w . The discriminant of a monic polynomial is non-zero if and only if the polynomial is square-

free, i.e. has no repeated factors. Since we have assumed that K is perfect, any irreducible polynomial over K must be square-free and hence the discriminant of f is non-zero.

If K is sufficiently large then we can find a $u_0 \in K$ such that $\text{disc}(f)$ doesn't vanish at $u = u_0$. Then the map sending u to u_0 is a well defined homomorphism from A onto a ring B . B can be presented as $K[w]$ modulo $f(u_0, w)$. Let $f_1(w)$ be an irreducible factor of $f(u_0, w)$ over K . There is a natural homomorphism from B onto the field $B_1 = K[w]/(f_1(w))$. If we have a non-trivial factorization of f with coefficients in A , then applying both of the above homomorphisms we arrive at a non-trivial factorization with coefficients in B_1 . Our choice of an irreducible factor of $f(u_0, w)$ was arbitrary, so we may as well choose the factor of least degree. In particular, if there are any linear factors, then $f(x, y)$ must already be absolutely irreducible. In any case, we have found a presentation for an algebraic extension of K that contains K° . Note that it was unnecessary to compute $\text{disc}(f)$ to verify our choice of u_0 . $\text{Disc}(f)$ is zero if and only if f has a repeated factor. Thus we pick successive values of u_0 until we find one such that $f(u_0, y)$ remains square-free. If m is the degree of f in u , then we must test at worst $2mn$ values until we find one that works.

We have justified the following algorithm for finding an absolutely irreducible factor of a polynomial $f(x, y)$ whose discriminant with respect to y is nonzero and irreducible over a sufficiently large field K .

0. If $\text{gcd}(\deg_x f, \deg_y f) = 1$ then return $f(x, y)$
1. Find an x_0 in K such that $f(x_0, y)$ is square free. (may fail if K is finite)
2. Factor $f(x_0, y)$ over K and let $f_1(y)$ be a factor of least degree.
3. If $f_1(y)$ is linear return $f(x, y)$.
4. Else factor $f(x, y)$ over $K[w]/(f_1(w))$ and return a factor of minimal degree.

For the more general problem of finding absolutely irreducible factors of multivariate polynomials, the same approach works. Step 0 is changed to take the gcd of the degrees of all variables present, and in step 1 we must find values for all variables but one such that the resulting univariate polynomial is square free and the leading coefficient doesn't vanish. We have assumed that K is large enough so that we can find substitution points that leave f square free. This may fail if K is a finite field. In that case if we let m be the gcd of the degrees of all variables present, then we do know that $[K^o:K]$ divides m . Thus K^o is a subfield of the unique extension of K of degree m . It therefore suffices to factor over that extension.

Once we have found an absolutely irreducible defining polynomial for y , then we adjoin the coefficients of all the monomials to K and this generates K^o the exact coefficient field.

SECTION 3. BINOMIAL POLYNOMIALS

If $f(x,y)$ is of the form $y^n - g(x)$ then a much simpler algorithm exists for obtaining an absolutely irreducible factorization. This is based on the following theorem proven in [Lang] p.221.

Theorem: [2] Let k be a field and n an integer ≥ 2 . Let $a \in k$, $a \neq 0$. Assume that for all prime numbers p such that $p | n$ we have $a \notin k^p$, and if $4 | n$ then $a \notin -4k^4$. Then X^{n-a} is irreducible in $k[X]$.

We define square-free factorization of a polynomial $g(x) \in K[x]$ to be

$$g(x) = c \prod g_i^{e_i} \tag{3}$$

where $c \in K$, $\gcd(g_i, g_j) = 1$ for $i \neq j$, and each g_i has no repeated factors. If we assume that K is perfect, then $g_i(x) \in K[x]$ for all i . This leads us to the following simple algorithm for finding an absolutely irreducible factor of $y^n - g(x)$.

1. Compute a square-free factorization of $g(x)$ as in equation (3).
2. Let d be the gcd of all the e_i and n .
3. If $d = 1$ then f is absolutely irreducible and return it.
4. An absolutely irreducible factor of $y^n - g(x)$ is

$$y^{\frac{n}{d}} - c \prod g_i^{\frac{e_i}{d}}$$

CHAPTER 4 THE RATIONAL PART OF INTEGRAL

By Liouville's theorem if the integral of an algebraic function is expressible in terms of elementary functions, it can also be written as the sum of an algebraic function and constant multiples of logs of algebraic functions. We will label the sum of logs the *transcendental part* of the integral and the rest the *rational part*. This terminology is carried over from rational function integration. We use it since we want to draw strong parallels between our algorithms for algebraic function integration and the well known ones for rational function integration.

SECTION 1. RATIONAL FUNCTION INTEGRATION

There are primarily three basic algorithms for finding the rational part of the integral of a rational function, however they all share a common first stage. Polynomial division is employed to convert the integrand into the sum of a polynomial and a proper rational function. (Proper means the degree of the numerator is less than that of the denominator). The polynomial part can be trivially integrated termwise. At this point the three algorithms diverge.

Subsection 1.1. Full factorization

The simplest algorithm conceptually completely factors the denominator of the reduced integrand over the algebraic closure of the constant field. If only approximations to the roots were used this algorithm would be acceptable, but we want exact solutions and this requires constructing the splitting field of the denominator. This construction involves algebraic factoring and can be very expensive. Then a complete partial fraction decomposition is performed. Each term in the result is easy to integrate, but the result contains many algebraic

quantities. They are all unnecessary since the rational part of the integral can always be expressed without extending the constant field, but the task of trying to eliminate the algebraics can be costly. This algorithm is close in spirit to the proposed by Risch in [39] and partially implemented by Davenport in [14]. The similarity becomes more evident if we reexpress the algorithm in terms of power series. The portion of the partial fraction expansion involving a particular root of the denominator is identical to the principal part (negative degree terms) of the power series expansion of the integrand at that root. Thus the algorithm can be expressed as power series computation and termwise integration. The additional step in Risch's algorithm involves reconstruction of the rational part of the integral from the principal parts of its expansions at all its poles. This is a fairly complex process for an algebraic function, however a rational function is simply determined up to an additive constant as the sum of its principal parts.

Subsection 1.2. Linear equations

Another algorithm for rational function integration was proposed by Horowitz [26]. This approach is global as contrasted with the local techniques used above. First a square-free factorization of the denominator is performed.

$$D = \prod D_i^i, \text{ gcd}(D_i, D_j) = 1 \text{ for } i \neq j \tag{1}$$

and each D_i is square-free (has no multiple factors). By observing that the integral of $(x - c)^{-k-1}$ is $-(x - c)^{-k}/k$, we see that integration reduces the order of a pole by one.

Thus

$$\int \frac{A}{\prod D_i^i} = \frac{B}{\prod D_i^{i-1}} + \int \frac{C}{\prod D_i}$$

where the last integral produces the transcendental part. The degrees of B and C are constrained since both are numerators of proper rational functions. By letting B and C be

polynomials with undetermined coefficients, differentiating both sides of the above equation and equating coefficients of the same powers of x we get a system of linear equations for the coefficients of B and C . Since the rational part of the integral is uniquely determined up to an additive constant and we have constrained the constant by requiring a proper rational function, the above system has a unique solution. This method is relatively insensitive to sparseness in the input or output.

Subsection 1.3. Hermite's algorithm

The third technique is due to Hermite [24]. It attempts to reduce the "complexity" of the problem using a succession of linear first degree congruence equations. Again we start by performing a square-free factorization of the denominator. But now instead of trying to find the entire rational part in a single step, we will repeatedly find pieces of the rational part that can be used to reduce the multiplicity of the denominator of the integrand. The algorithm we present here treats the factors of the denominator one at a time. Mack [Mack.] presents a variant that treats all the factors at once, but his algorithm would only make our formulas more complicated, and the underlying theory is essentially the same.

Again we assume a square-free factorization of the denominator of the integrand as in (1). Let $V = D_{j+1}$ for some $j > 0$, i.e. V is a multiple factor of the denominator. Let U be the cofactor of V , $U = D/V$. By (1) U and V are relatively prime and we can write the integral as

$$\int \frac{A}{UV^{j+1}} dz$$

We will attempt to repeatedly reduce the multiplicity of the denominator while constructing portions of the final answer. We claim that there exist polynomials B and C such that

$$\int \frac{A}{UV^{j+1}} = \frac{B}{V^j} + \int \frac{C}{UV^j}$$

After differentiating both sides and multiplying by UV^{j+1} we get

$$A = UVB' - jBUV' + VC$$

This is a differential equation with two unknowns so it would seem that we have made the problem more complicated. We will additionally claim, however, that there is a unique solution B such that $\deg(B) < \deg(V)$. We then reduce the equation modulo V . This eliminates the B' and C terms from the above equation and we are left with

$$A \equiv -jBUV' \pmod{V} \quad (2)$$

This equation will indeed have a unique solution as long as $j \neq 0$ and $\gcd(V, V') = 1$. But V is square-free by construction so the latter requirement is satisfied. As long as $j > 0$ we can find a unique B solving (2) and then subtracting $(B/V^j)'$ from the integrand will reduce the multiplicity of the denominator. By repeating this process with all multiple factors of the denominator, we see that the integral of any rational function can always be expressed as the sum of a rational function and an integral whose denominator has multiplicity one. The latter integral has no rational part, i.e. it is expressible exclusively as a sum of logs.

SECTION 2. ALGEBRAIC FUNCTIONS

We choose to base our algorithm for finding the rational part of the integral of an algebraic function on Hermite's method for rational function integration. In fact all three of the algorithms presented in the previous section can be generalized to handle algebraic functions. The generalization of the first approach requires Puiseux expansions and algebraic number computations that we wish to avoid. The advantage of the third approach over the second is that we are provided with insight into how an integral may fail to be elementary and its step by step reductive nature will often allow us to return partial results instead of merely returning "not integrable".

To simplify matters we will transform the integral so that there are no poles or branch points at *infinity*. We assume an integral of the form $\int \sum R_i(x)y^i/Q(x) dx$ where R_i and Q are polynomials and y satisfies $f(y,x) = 0$. Let a be an integer that is neither a root of Q nor a root of the discriminant of f . We can "move" the point a to ∞ by defining $z = 1/(x - a)$ or $x = a + 1/z$. Our new integral is

$$\int \sum \frac{R_i(a + 1/z)y^i}{Q(a + 1/z)} (-z^{-2}) dz$$

After performing the integration we can apply the inverse transformation to express the answer in terms of x instead of z . If we let m be the degree of f in x , then the transformed minimal polynomial for y is $g(y,z) = z^m f(y, a + 1/z)$. Using the results of last chapter, we can find a basis, $[w_1, \dots, w_n]$ for the integral closure of $K[z]$ in $K(z,y)$. In terms of this basis the integrand can be expressed as $\sum A_i(z)w_i/D(z)$ where D and A_i are polynomials. Since this integrand has no poles at ∞ , $\deg(A_i(z)) < \deg(D(z))$ for all i .

We now attempt to imitate Hermite's algorithm for rational functions. Again we start by performing a square-free factorization of the denominator $D(z)$ yielding $D = \prod D_i^i$. Let $V = D_{k+1}$ for some $k > 0$ and $U = D/V$. Following Hermite's algorithm we might now look for polynomials B_i and C_i such that

$$\int \sum A_i \frac{w_i}{UV^{k+1}} dz = \sum B_i \frac{w_i}{V^k} + \int \sum C_i \frac{w_i}{UV^k} dz \quad (3)$$

Unfortunately we won't be able to find them in general without additional restrictions on U . The difficulty is caused by the fact that y' has a non-trivial denominator. Hermite depended on the fact that the derivative of a polynomial is a polynomial, but this is not necessarily true for algebraic functions. If y satisfies $f(y,z) = 0$ then we can find $y' = dy/dz$ by taking the total derivative of the defining polynomial.

$$\frac{\partial f}{\partial y} dy + \frac{\partial f}{\partial z} dz = 0$$

$$\frac{dy}{dz} = -\frac{\partial f/\partial z}{\partial f/\partial y}$$

Thus y' is a rational function in y and z . Similarly w_i' in general has a nontrivial denominator. Let E be the least common denominator of w_i' for all i . Then the derivative of the second term in (3) will in general have E as a factor of its denominator. This will force us to permit E to be a factor in the denominator of the third term. We want to use (3) iteratively; this means the third term on one iteration will become the first term for the next iteration. Therefore we may as well assume the denominator of the first term is also divisible by E . In the last iteration $k = 1$ making the denominator of the third term is UV , so our additional restriction on equation (4) is that $E \mid UV$. We can always guarantee this by multiplying the A_i and U by a suitable factor of E . Note that there is then an uncanceled gcd between the numerator and denominator of the first term in (3). We need to be sure that the new U is still relatively prime to V . This can only be guaranteed if we know E to be square-free. Fortunately that is always the case.

Lemma: [1] If $v_p w \geq 0$ and $v_p(z - a) > 0$ then $v_p(z - a)w' > 0$.

Proof: $v_p w \geq 0$ implies $v_p dw \geq 0$, and $v_p(z - a) > 0$ implies $v_p(z - a) = v_p dz + 1$. ([Chev51] IV.8

Lemma 1) But $dw = w' dz$ and thus $0 \leq v_p w' dz < v_p(z - a)w'$. \square

We therefore begin with equation (3) with the following conditions:

$$E \mid UV, \gcd(U, V) = 1, \gcd(V, V') = 1 \tag{4}$$

Now as before we perform the differentiation and multiply through by UV^{k+1} yielding

$$\sum A_i w_i = U \sum (VB_i' + B_i V^{k+1} \left(\frac{w_i}{V^k}\right)') + V \sum C_i w_i \tag{5}$$

We then reduce the equation modulo V .

$$\sum A_i w_i \equiv \sum B_i UV^{k+1} \left(\frac{w_i}{V^k} \right)' \pmod{V} \quad (6)$$

We must show that this equation always has a unique solution. This is equivalent to showing that the $S_i = UV^{k+1} \left(\frac{w_i}{V^k} \right)'$ is a local integral basis, i.e. any integral function can be expressed as a linear combination of them with rational function coefficients and denominator relatively prime to V . There are two ways this can fail, either the S_i are not linearly independent over $K(z)$ or there exists an integral function whose representation requires a factor of V as a denominator. Since the former case implies that zero is a nontrivial linear combination of the S_i , both cases may be summarized by saying there exists an integral function that can be represented as

$$\frac{1}{V} \sum T_i UV^{k+1} \left(\frac{w_i}{V^k} \right)' \quad \text{where } V \text{ does not divide } UT_i \text{ for some } i \quad (7)$$

For purposes of generating a contradiction we assume that the S_i do not form a local integral basis and thus there exists an integral function (7). We can add $\sum (UT_i)' w_i$ to equation (7) yielding a new integral function G such that

$$G = \sum V^k \left((UT_i)' \frac{w_i}{V^k} + UT_i \left(\frac{w_i}{V^k} \right)' \right) = V^k \sum \left(UT_i \frac{w_i}{V^k} \right)' \quad (8)$$

Assuming that such an integral function G exists implies there exists a function for whom differentiation doesn't increase the order of its poles; this is the source of our contradiction. Let $F = \sum UT_i w_i / V^k$. The restriction on the T_i in (7) says that a smaller value of k would be insufficient, i.e. there is some place p such that $v_p F < v_p (1/V^{k-1})$ and $v_p V > 0$ where v_p is the order function at p .

Lemma: [2] If u is a nonzero function such that $v_p u \neq 0$ and p is a finite place (not over ∞) then $v_p u' = v_p u - r$ where r is the ramification index of p with respect to $K(z)$.

Proof: We embed our function field in its p -adic completion and let t be a uniformizing

parameter at p . We can then write

$$u = \sum_{i=j}^{\infty} c_i t^i$$

The coefficients in the series are algebraic over our constant field and we can extend our derivation d/dz uniquely to the completion yielding:

$$u' = \sum_{i=j}^{\infty} i c_i t^{i-1} t'$$

Since $j \neq 0$, $v_p u' = j - 1 + v_p t'$. Since p is a finite place, the p -adic expansion of z begins $z = a_0 + a_1 t + \dots$ where $a_1 \neq 0$ so $v_p(dz/dt) = r - 1$. Thus $v_p t' = v_p(dt/dz) = 1 - r$ and $v_p u' = (j - 1) + (1 - r) = j - r$. \square

Since V is a square-free polynomial in z and $v_p V > 0$ we have $v_p V = r$ and thus $v_p(1/V^k) = v_p(1/V^{k-1}) - r$. If $k > 0$ then $v_p F < v_p(1/V^{k-1})$ implies $v_p F < 0$. Using lemma 2 we see that $v_p F' < v_p(1/V^{k-1}) - r = v_p(1/V^k)$. But according to equation (8) F' can be written as G/V^k where G is an integral function. This contradicts the fact that $v_p F' < v_p(1/V^k)$. Thus the S_i do indeed form a local integral basis, and equation (6) will have a unique solution modulo V as long as $k > 0$.

By our choice of E there must exist polynomials M_{ij} such that

$$E w_i' = \sum_j M_{ij} w_j$$

Since $E \mid UV$ let $TE = UV$ for some polynomial T .

$$UV w_i' = TE w_i' = T \sum_j M_{ij} w_j \tag{9}$$

Substituting (9) into equation (6) yields:

$$\sum_i A_i w_i \equiv \sum_i (-kUV' B_i) w_i + \sum_i B_i T \sum_j M_{ij} w_j \pmod{V} \tag{10}$$

If we now equate the coefficients of w_i on both sides of (10) we get a set of linear congruence

equations with B_i as unknowns.

$$A_i \equiv -kUV'B_i + T \sum_j B_j M_{ji} \pmod{V} \quad (11)$$

We have shown that this system will always have a unique solution for $k > 0$. Thus the determinant of the coefficient matrix must be relatively prime to V and the system can always be solved, e.g. using Cramer's rule.

For efficiency reasons it is important to recognize when the system (11) decouples, i.e. each equation involves only one unknown. When $V|T$ the summation in (11) vanishes and we are reduced to solving a succession of equations of the form

$$A_i \equiv -kUV'B_i \pmod{V}$$

This case of algebraic integration is almost exactly the same as in rational function integration. $V|T$ if and only if $\gcd(V,E) = 1$. Thus it can become worthwhile to split V into two factors, one that divides E and one relatively prime to E and treat each case separately.

The other situation in which the system decouples is when the matrix M_{ij} in (9) is diagonal. This means that $w_i' = R_i w_i$ where R_i is a rational function in z . We can solve this differential equation, and the solution will be an algebraic function if and only if $w_i^{m_i} \in K(z)$ [42]. Thus the matrix can be diagonal if and only if $K(z,y)$ is a compositum of single radical extensions.

SECTION 3. POLES AT INFINITY

Repeated applications of the reductions presented in the previous section will leave us with an integral whose denominator is square-free. One might hope that we have removed all singularities from the integrand except those that should be cancelled by log terms as in the

rational function case. Unfortunately while we have been reducing the order of the finite poles, we may have introduced poles at ∞ . In this section we will prove that if our basis is normal at infinity the presence of such poles will prove that the original problem was non-integrable.

The problem is caused by the second term in equation (3), $E = \sum B_i w_i / V^k$. By construction we have the restriction that $\deg(B_i) < \deg(V)$. Thus at each iteration of the previous algorithm, the coefficients of the answer produced will be proper rational functions, i.e. degree of numerator less than degree of denominator. This guarantees that the rational function coefficients of the portion of the answer we have produced will have positive order at ∞ , but the w_i themselves will in general have poles at ∞ . We will see that if this algebraic portion of the answer has poles at infinity, then the original problem was not integrable.

Lemma: [3] If u is a nonzero function such that $v_p u \neq 0$ with p a place over ∞ , then $v_p u' = v_p u + r$ where r is the ramification index of p .

Proof: The proof is identical to Lemma 2 except that $v_p t'$ is different. The p -adic expansion of z is $z = a_{-r} t^{-r} + \dots$ so $v_p(dz/dt) = -r - 1$. Thus $v_p t' = v_p(dt/dz) = r + 1$ and $v_p u' = (v_p u - 1) + (r + 1) = v_p u + r$. \square

Lemma: [4] If u is a nonzero function such that $v_p u \neq 0$ at some place p , then $v_p du = v_p u - 1$.

Let us assume that the third term in equation (3) has poles at infinity. Since the original integrand had zero residue at infinity, and the derivative of any algebraic function has zero residue everywhere this third term must also have zero residue at infinity. Thus if it has poles there, they must be at least double poles. If this term is integrable it must be expressible as an algebraic function and a sum of constant multiples of logs. This algebraic function can not have any finite poles, since the integrand has only simple poles in the finite plane. Thus this algebraic function must be expressible as a polynomial multiple of our basis elements w_i . Since

this function has a pole at infinity, either one of the coefficients is a polynomial of positive degree, or some non-constant basis element has a non-zero coefficient. Although the second term in equation (3) could also have poles at infinity, when we add these functions those poles do not cancel. The coefficients of the basis elements in the second term are all proper rational functions, i.e. degree of numerator less than degree of denominator. This is a consequence of the fact that $\deg(B_i) < \deg(V)$ since each B_i is computed modulo V . When we add this polynomial multiple of our basis elements to the second term, we arrive at a function with at least one coefficient that is an improper rational function. If this coefficient has a pole at infinity then the function itself has a pole at infinity since the basis is normal at infinity and no basis elements have zeros at all places over infinity. If none of the coefficients have poles at infinity, then some non-constant basis function must have a coefficient of order zero at infinity. But since any non-constant basis function must have a pole at infinity, both possibilities imply the complete algebraic part of the integral must have a pole at infinity. But by Lemma 4 this would imply that the integrand would have a pole at infinity that contradicts our initial assumption. Thus we see that if in the process of finding the algebraic part, we introduce poles at infinity then the original problem was not integrable and there is no need to try and remove these poles. We have demonstrated the following theorem and corollary:

Theorem: [1] If gdx is a differential with zero residues and no poles at infinity, then the algorithm of the previous section computes a function h such that $(g - h')dx$ has no finite poles and is zero if gdx was integrable.

Corollary: If gdx also has non-zero residues then $(g - h')dx$ has only simple poles in the finite plane and has no poles at infinity if gdx was integrable.

CHAPTER 5 LOG TERMS AND DIVISORS

In this chapter we will present algorithms for finding the logarithmic or transcendental part of the integral. Using the results of the previous chapter we can assume that the rational part has been removed and the integrand has been reduced to a differential with only simple poles. Using Liouville's theorem we know that if the integral is expressible as an elementary function, then it can be written as

$$\int R(x,y)dx = \sum c_i \log v_i(x,y) \quad (1)$$

where $c_i \in K'$ and $v_i \in K'(x,y)$ with K' a finite algebraic extension of K . To find the rational part of the integral no extension of the ground field was necessary, but the same is not true for the logarithmic part. We will show that the residues of the integrand generate the unique K' of minimal degree over K sufficient to express the answer. Additionally these residues provide us with clues about the orders of the poles and zeros of the v_i .

SECTION 1. PROPERTIES OF LOGARITHMIC DIFFERENTIALS

We will start by examining the relationship between $f \in K(x,y)$ and the differential $\frac{df}{f}$. We are interested in the local behavior at a place p with t as a uniformizing parameter. If a function f has order k at p then $f = t^k g$ where g is a function whose value at p is finite and nonzero. Any differential $f dg$ can be written as $h dt$ where $h = f \frac{dg}{dt}$. The order at p of the differential $f dg$ is defined to be the order of the function h above. The residue at p of the differential $h dt$ is the coefficient of the t^{-1} term in the series expansion of h in powers of t . Both the residue and the order of a differential are independent of the choice of uniformizing parameter.

Next we examine the properties of differentials of the form $\frac{df}{f}$ at a place p with local uniformizing parameter t . If f has order k at p it can be written as $f = t^k g$ for some function g of order 0 at p and similarly the differential can be rewritten as:

$$\frac{df}{f} = k \frac{dt}{t} + \frac{dg}{g}$$

As shown in chapter four, since g has no pole at p , dg must have non-negative order at p also. By construction g has order 0 at p and thus $\frac{dg}{g}$ has non-negative order at p . Since t is a local uniformizing parameter at p the order of $k \frac{dt}{t}$ is the same as the order of $\frac{k}{t}$ that is precisely -1 as long as k is nonzero. Similarly the residue of $\frac{df}{f}$ is the same as the residue of $k \frac{dt}{t}$ that is precisely k . Thus we have shown that the order of $\frac{df}{f}$ is always greater than or equal to -1 , and its residue at any place p is the same as the order of f .

This leads to a solution to a special case of our original problem. When can a differential be expressed in the form $\frac{df}{f}$ for some function f in our function field. The following two necessary properties give quick failure tests:

1. The order of the differential must be greater than or equal to -1 everywhere.
2. The residues must all be integers since they correspond to the orders of the desired function f .

If a differential passes those tests, then we try to determine if there exists a function whose order at every place is equal to the residue of the differential at that place. Thus the residues of the differential provide us with a formal specification of the location and orders of the poles and zeros of the desired function. Since the differential only has a finite number of poles, it can only have nonzero residue at a finite number of places. A divisor is a formal integer linear combination of places that has a finite number of non-zero coefficients. Divisors that correspond to the orders of the poles and zeros of actual functions are called principal divisors. A differential with integer residues everywhere immediately provides us with a

divisor and we are left with the problem of determining whether or not it is principal. An algorithm for answering this question will be presented later in this chapter.

The solution to the problem discussed in the previous paragraph leads to some computational difficulties, but no great theoretical problems. But if we generalize things slightly we arrive at a problem that many mathematicians around the turn of the century worked on but were unable to solve. In fact Hardy even went so far as to state "there is reason to believe that no solution exists". The generalization involves introducing one additional constant coefficient. Instead of asking whether a given differential can be expressed as $\frac{df}{f}$, we allow one more degree of freedom and try to express it as $1/m \frac{df}{f}$ for some integer m . By the same reasoning as above we see that if we multiply all the residues of the original differential by m we arrive at the desired divisor for f . But we don't know the value of m . This is the source of the theoretical difficulty. We derive a divisor from the residues of the integrand and we can test whether or not it is principal. If it is then we are done. If it is not, however, perhaps by scaling all the orders specified by two, we arrive at a principal divisor. If not try scaling by three, ... etc. The real difficulty is knowing when to stop. More formally given a divisor D we need to be able to determine a bound M such that if for all positive integers $j < M$, jD is not principal, then we are guaranteed that there exists no multiple of D that is principal. It is only in the last thirty years that a solution to this problem has been discovered using the technique of good reduction from algebraic geometry. Basically the original coefficient field is reduced to a finite field. There each divisor has finite order, i.e. there is some finite value of j such that jD is principal. This information is used to limit the potential set of j 's that needs to be examined. We will present this construction in the next chapter.

Armed with some intuition from the previous special cases, we will now investigate the fully general problem indicated by equation (1). We wish to write the given differential as

$$R(x,y)dx = \sum c_i \frac{dv_i}{v_i} \quad (2)$$

for some constants c_i and some rational functions $v_i(x,y)$. As indicated earlier the constants c_i and the coefficients of the v_i will in general lie in a finite algebraic extension K' of the field of constants of the original function field. Our first step involves the construction of the minimal extension of K sufficient to express the answer.

The decomposition of the $R(x,y)dx$ indicated by equation (2) is certainly not unique. In fact let b_j be another set of constants such that each c_i can be expressed as integer multiples of the b_j , $c_i = \sum n_{ij}b_j$. Define new functions $w_j = \prod_i v_i^{n_{ij}}$ then:

$$\sum_i c_i \frac{dv_i}{v_i} = \sum_j b_j \frac{dw_j}{w_j}$$

Since a linear dependence among the coefficients implies we can express the sum with fewer terms, the answer with the smallest number of summands will have coefficients that are linearly independent over the rationals.

We next investigate the relationship between the coefficients of the log terms and the residues of the integrand. Since the residue of $\frac{dv}{v}$ is always an integer for any function v , we see that the residues of the integrand are always integer linear combinations of the c_i 's. Thus the coefficients of the log terms generate a \mathbb{Z} -module that contains all the residues of the integrand. We will show that the c_i 's can be chosen so that they generate the same vector space over the rationals that the residues of the integrand do. Let a_j form a basis for the vector space generated by the residues. Let b_k extend this basis to include the coefficients of the log terms. Thus each c_i can be written uniquely as $\sum r_{ij}a_j + \sum s_{ik}b_k$ where the r_{ij} and s_{ik} are rational numbers. We will in fact assume the r_{ij} and s_{ik} are integers, which can be accomplished by suitably scaling the basis elements. Then as shown in the previous paragraph, we can construct functions w_j and u_k such that:

$$\sum_i c_i \frac{dv_i}{v_i} = \sum_j a_j \frac{dw_j}{w_j} + \sum_k b_k \frac{du_k}{u_k} \quad (3)$$

Since however the a_j 's form a basis for the vector space containing all the residues of the entire sum and the b_k 's and the a_j 's form a linearly independent set by construction, the residues of each $\frac{du_k}{u_k}$ must be zero everywhere. This implies that each u_k has order zero everywhere, is thus a constant. Since the differential of any constant is zero, the second sum in equation (3) is identically zero.

We have shown that if a differential has a decomposition as in equation (2), it has one where the coefficients form a basis for the vector space spanned by the residues of $R(x,y)dx$. If we compute such a basis r_i , then by Lemma 1, we are guaranteed that there exist integers n_i such that $\frac{r_i}{n_i}$ can be chosen as the coefficients of the log terms. Let K' be K extended by all the residues of the integrand. Assume we have found a representation for the integrand fdx as a sum of logarithmic differentials, $\sum c_i \frac{dv_i}{v_i}$. We have shown that we can assume the coefficients of the log terms c_i lie in K' . We wish to show that we can also assume that the $v_i \in K'(x,y)$. If instead $v_i \in E(x,y)$ where E is a finite algebraic extension of K' of degree j then by applying a trace from E to K' , $tr_{\frac{E}{K'}}$ we arrive at a solution whose constant field is exactly K' .

$$fdx = \sum c_i \frac{dv_i}{v_i}$$

$$tr(fdx) = \sum tr(c_i \frac{dv_i}{v_i})$$

Since $f \in K(x,y)$, $tr(fdx) = jfdx$ and $tr(dv/v)$ is the same as $d(Nv)/(Nv)$ where N is the norm from E to K' .

$$jfdx = \sum c_i \frac{dNv_i}{Nv_i}$$

But Nv is a rational function with coefficients in K' , so we have shown if the integral can be expressed over some algebraic extension of K then in fact it can be expressed over the extension of K generated by the residues of the integrand, and no smaller extension will suffice.

SECTION 2. COMPUTING THE RESIDUES

Now that we have demonstrated the importance of the residues of the integrand, we need an efficient way to compute these residues. If the differential is expressed in the form $f dt$ where t is the local uniformizing parameter for some place p , then the residue was defined to be the coefficient of t^{-1} in the series expansion of f at p . If f is known to have order greater than or equal to -1 at p , then we can just compute the value of tf at p . Thus we need to be able to find local uniformizing parameters for places, and be able to compute the value of functions at places. We can view the Riemann surface associated with the differential as a multisheeted covering of the complex X -plane. Thus each finite place p_0 can be associated with some x -value x_0 by projection. The order of the line of projection at a place defines the branch index of that place. Branch places are those where the line of projection is tangent to the Riemann surface and thus have branch index greater than one. Since $x = x_0$ is the equation of the line of projection from p to the X -plane, the order of the function $x - x_0$ at p is equal to the branch index of p . If p is not a branch place, $x - x_0$ has order 1 and can thus be used as a local uniformizing parameter.

Theorem: Let $f dx$ be a differential with order greater than or equal to -1 at some place p with branching index r centered at x_0 . The residue of $f dx$ at p is equal to the value of the function $r(x - x_0)f$ at p .

Proof: Let t be a uniformizing parameter at p . Since $x - x_0$ has order r at p , it can be written as

$$x - x_0 = t^r g \tag{5}$$

where g has order zero at p

$$dx = (rt^{r-1}g + t^r \frac{dg}{dt})dt$$

Since dg/dt has non-negative order at p , dx has order $r - 1$ at p and f must have order greater

than or equal to $-r$ at p

$$f dx = r t^{r-1} f g dt + t^r f \left(\frac{dg}{dt} \right) dt \quad (6)$$

the second term on the right side of equation (6) is holomorphic at p so the residue of $f dx$ at p is the same as the residue of the first term on the right side of (6). By the argument of the preceding paragraph, this residue can be computed by evaluating the function $r t^r f g$ at p . Using equation (5) this function can be written as $r(x-x_0)f$. \square

Theorem: Let $f dx$ be a differential with at most simple poles in the finite plane. Let $D(x)$ be a polynomial whose roots include the x -projections of the poles of f . Let $g = fD$, then $g dx$ is holomorphic in the finite plane and the residue of $f dx$ at any place p with branch index r centered over a root of $D(x)$ is equal to the value of rg/D' at that place.

Proof: Since $f dx$ has only simple poles in the finite plane, $g dx$ has no finite poles. Let p be a place over a root x_0 of $D(x)$. From the previous theorem we only need to show that at p $(x-x_0)f = \frac{g}{D'}$. Let $D(x) = (x-x_0)C(x)$. Then since $f = \frac{g}{D}$, $(x-x_0)f = \frac{g}{C}$. But $D' = C + (x-x_0)C'$ and thus $D'(x_0) = C(x_0)$ at p . \square

We could use the previous theorem to compute separately each residue of the integrand, but it will be more convenient to find them all at once. We will construct a polynomial whose roots are rational multiples of the residues of the differential $f dx$. f is in general a rational function of x and y . After rationalizing the denominator we can assume $f(x,y) = g(x,y)/D(x)$ where g and D are polynomials. Let Z be a new indeterminate and define $R(Z) = \prod Z D' - g$ where the product is taken over all places centered above roots of D . The roots of R are the residues of $f dx$ divided by the branch orders. Since the branch orders are always positive integers, the splitting field of R is precisely the minimum extension of K containing all the residues of the integrand, and thus the smallest extension of the coefficient field in which the integration can be performed. Similarly a Q -basis for the roots of R provides us with a Q -basis for the residues. A key observation is that R can be computed without extending the coeffi-

cient field. Let $F(x,y) = 0$ be the defining polynomial for our function field. For a particular value of x , x_0 , $\prod ZD'(x_0) - g(x_0,y)$ taken over all places sitting over x_0 , is just the resultant ${}_Y(ZD'(x_0) - g(x_0,Y), F(x_0,Y))$ both viewed as polynomials in Y . Extending that product over all roots x_0 of D is just another resultant of the previous result with $D(X)$, both viewed as polynomials in X . Thus we can compute

$$R(Z) = \text{resultant}_X(\text{resultant}_Y(ZD'(X) - g(X,Y), F(X,Y)), D(X)) \quad (7)$$

Since the roots of R are nonzero rational multiples of the residues, the splitting field of R is the minimal extension containing all the residues. We have found R using only rational operations over K but to actually compute the log terms we will have to work in a coefficient field that contains all the roots of R . [48] gives an algorithm for computing the splitting field of R using algebraic factoring. This step can be very expensive, since if R has degree n the splitting field may be of degree n factorial. However there is no escaping this expense since this is the smallest extension in which the answer can be expressed.

We next need to compute a basis for the vector space spanned by the roots of R over the rationals. If the coefficients of R are all rational numbers, then we can view the splitting field of R as a vector space over the rationals. Each root is thus a vector with rational coefficients, and we are interested in finding a basis for the space spanned by these vectors. This can be done using standard techniques from linear algebra. In general however the coefficients of R will come from some finitely generated extension of \mathcal{Q} . This occurs when the integrand contains additional parameters or algebraic numbers. Let K be the field generated by the coefficients of R . The splitting field of R is a vector space over K and thus each root of R is representable as a vector of elements of K with respect to some chosen basis of the splitting field over K . We need to view these roots as elements of a finite dimensional vector space over \mathcal{Q} . We will do this by replacing each coefficient from K by a finite dimensional vector over \mathcal{Q} . Let b_i be a basis for the splitting field of R over K . Then the roots of R can be

represented as $r_j = \sum c_{ij}b_i$ where $c_{ij} \in K$. The c_{ij} will in general be rational quantities but by a different choice of basis elements we can guarantee they are polynomials in the generators of K' over K . Let d_i be a common denominator for the i th components of all the vectors, i.e. for c_{ij} with fixed i . If we choose as a new basis for K' over K , b_i/d_i then with respect to this basis all the coefficients will be polynomials. Now we have represented the roots as vectors of multivariate polynomials with rational coefficients. If we choose a basis for the set of monomials appearing in these polynomials, then with respect to the tensor product of this monomial basis and the original basis of K' over K , our roots are expressible as vectors of rational numbers and can thus be viewed as elements of a finite dimensional vector space over \mathcal{Q} . At this point we apply gaussian elimination to this collection of vectors adjoined to an identity matrix, to find a minimal basis over \mathcal{Q} , and the rational linear equations expressing all the roots in terms of this basis set.

SECTION 3. CONSTRUCTING DIVISORS

The basis for the residues over \mathcal{Q} become our candidates for coefficients of the log terms. They are only candidates since they may be integer multiples of the correct coefficients. For each candidate we next proceed to compute an associated divisor. The minimum multiple of this divisor that is principal will provide us with the appropriate scaling of the candidate coefficient. We first need to construct a set of building blocks from which we will construct our divisors. For each root of R we will construct a divisor that has order one at each place where the integrand has that root as residue, and order zero elsewhere.

We assume the integrand is of the form $(G(x,y)/D(x))dx$ as above. In order to simplify matters we will also assume that all places centered above roots of D are not branch places. In the next section we will show how to convert the general problem to satisfy this restriction. Under this assumption the roots of R are precisely the residues of the integrand, and at any place above the root x_0 of D , $x-x_0$ is a local uniformizing parameter. By theorem 2 G/D' is

a function whose value over all places centered above roots of D is the residue of the integrand. Let r be a root of R , then $G/D' - r$ is a function that vanishes wherever the integrand has residue equal to r and at the other places above roots of D , has a nonzero value. Define $B(x) = D(x)/\gcd(D, D')$. Then B contains all factors of D taken with multiplicity one.

Proposition: The minimum of the orders of $B(x)$ and $G - rD'$ at any finite place p is equal to one if the integrand has residue r at p and zero otherwise.

Proof: If the integrand has residue r at p then $G - rD'$ vanishes at p and $B(x)$ vanishes to order one at p . If p is any place over a root of D where the integrand has residue different from r then $G - rD'$ has order zero at p . If p is a finite place not centered over a root of D then $B(x)$ has order zero at p . \square

If our function field had unique factorization we would proceed by computing the greatest common divisor of $B(x)$ and $G - rD'$ and use these to build our divisors. But unless the function field happens to have genus 0, we are not guaranteed that the notion of gcd is well defined. We will work instead with the ideal generated by these two functions over the ring of integral algebraic functions. In a gcd-domain this ideal would have a single generator but not in general.

Now that we have created these ideal building blocks, we need algorithms for multiplication and division. In general a set of generators for the product of two ideals can be computed simply as the set cross products of generators from one times generators from the other. This means that the product of two ideals with m and n generators respectively will require mn generators. Our primitive building blocks have only two generators, and we will first show how this property can be maintained under multiplication and division.

First we will modify our building blocks slightly. We define the support of a divisor as the set of places in the divisor with nonzero order. Our building blocks are of the form $\{h(x, y, r), B(x)\}$. We know that the zeros of h at places over roots of B coincide with the zeros

of the desired divisor, but h can have zeros of multiplicity greater than one, which is compensated in the ideal by the fact that B has only zeros of order one. Our first step is to modify h so that it has only simple zeros. This can be done by adding a random integer multiple of B to h . For all but a small finite number of choices, this will produce an h with only simple zeros at places over the roots of B . To complete this step we only need a test to guarantee our random integer was not "unlucky." The norm of h is a polynomial in X whose degree is the sum of the order of the finite zeros of h . For any chosen integer j we can compute $N(x) = \text{Norm}(h + jB)$. We can split write $N = N_1 N_2$ where N_2 is the part of N that is relatively prime to B . We wish to choose j so that N_1 has as small a degree as possible. When this is done, the degree of N_1 will be the same as the number of places where the integrand had residue r that is the same as the multiplicity of r in $R(x)$. Thus j is "lucky" as long as degree N_1 is the same as the multiplicity of r in $R(x)$.

Now we have divisor descriptions of the form $(h(x,y,r), A(x))$ for divisor D and satisfying the following properties:

1. $\text{order}(h) = \text{order}(D)$ at all places over the roots of the $A(x)$
2. $\text{order}(D) = 0$ at all other places
3. both h and A are multiples of D except at infinity

(h,A) can be viewed as a locally principal model of the divisor D since at any place either the $\text{order}(D) = \text{order}(h)$ or $\text{order}(D) = \text{order}(A)$. Given two such descriptions $(h_1(x,y), A_1(x))$ for divisor D_1 and $(h_2(x,y), A_2(x))$ for divisor D_2 , where the support of A_1 equals the support of A_2 , we claim the description $(h_1 h_2, A_1 A_2)$ is of the same form for divisor $D_1 D_2$.

Quotients of divisors are slightly more complicated. h_1/h_2 satisfies property 1 but not necessarily property 3. It may have extraneous poles outside the support of D_1 and D_2 . These can be removed by rationalizing its denominator and removing any factors that are relatively

prime to A_1 or A_2 . If we let h_3 be this normalization of h_1/h_2 whose finite poles are all above A_1 or A_2 , then $(h_3, A_1 A_2)$ satisfies all three properties for the divisor D_1/D_2 .

Now that we have the basic building blocks, and simple algorithms for multiplying and dividing them, we are ready to construct the divisors associated with each of the residue basis elements. Let b_i be the i th basis element. Each root r_j of $R(x)$ can be represented uniquely as an integer linear combination of the basis elements $r_j = \sum c_{ij} b_i$. Also each r_j can be associated with the set of places P_j where the integrand had residue equal to r_j . The divisor D_j associated with basis element b_i is defined as $\prod P_j^{c_{ij}}$. We have seen how to represent P_j as the basic building block associated with r_j , and we now construct D_i as appropriate products and quotients of the P_j . This gives us a computable representation for the divisors associated with our candidate logands. In the next chapter we will see how to determine if there is some multiple of these divisors that are principal and thus whose generators will furnish us with the desired log terms.

SECTION 4. DEALING WITH BRANCH PLACES

In the previous section we assumed that the branch places of the function field did not lie above any of the poles of the integrand. We now will show how to bring the integrand into this form. If the integrand involves only a single unnnested radical, then we claim this assumption is guaranteed. In this case the defining polynomial for our function field is of the form $Y^n = F(x)$. The integrand can be written as $(\sum G_i(x) Y^i) dx$. As shown in the appendix, this is integrable if and only if each summand is. Thus we are reduced to integrands of the form $G(x) Y dx$ with possibly a different choice of Y . We can also assume without loss of generality that $F(x)$ is a polynomial whose roots all have multiplicity less than n . In this case the finite branch places occur precisely at the places above the roots of F . We claim that the integrand cannot have a nonzero residue at any branch place. If p is a branch place centered above a root x_0 of F with branch index r , then the order of $G(x)$ must be a multiple of r and the order

of dx is precisely $r - 1$. But the order of Y is some positive integer j less than r . In order for the integrand to have a simple pole at p , we must have $k*r + r - 1 + j = -1$. But this equation implies j is divisible by r contrary to our assumption. Thus with simple radicals we never have nonzero residues at branch places.

In the more general situation we have to change our model of the function field to achieve this condition. Branch places occur when the line of projection from our defining curve down to the x -axis is tangent to the curve. By a different choice of independent variable we change tangency points and thus arrive at a different set of branch points. For each pole of our integrand, there are only a finite number of projection directions that are tangent to the curve. If we replace X by $X + mY$ for some random integer m then for almost all choices of m , the resulting function field will not have branch places above any of the poles of our integrand. This can be checked by computing the discriminant of an integral basis for the new presentation of the function field. The branch places all lie above roots of this discriminant so we wish to choose m such the integrand has zero residue at all places above zeros of the discriminant. If we let $D(x)$ be this new discriminant and let $g = fD(x)$ where $f dx$ is the integrand, then equation (7) computes the residues of the integrand over all roots of $D(x)$. We then check that $R(Z)$ has no non-zero roots, i.e. is a pure monomial in Z .

CHAPTER 6 PRINCIPAL DIVISORS AND POINTS OF FINITE ORDER

In this chapter we will present a decision procedure for computing the log term associated with each element of a \mathcal{Q} -basis for the residues. In chapter five we showed how to construct a reduced divisor that described the pole zero ratios for the desired log term. We need to be able to test whether there exists some multiple of this divisor that corresponds to an actual function in $K'(x,y)$.

We will first present an algorithm for determining whether the given divisor is principal, i.e. is the divisor of a function in $K'(x,y)$. If this algorithm succeeds then we are done, but if not we need to try multiples of this divisor. If each of these tests fails then we need to know when we can stop, i.e. when we are guaranteed that there is no multiple of this divisor that is principal.

SECTION 1. PRINCIPAL DIVISORS

We start with a divisor description of the form $(h(x,y),g(x))$. This describes a divisor whose order at places over ∞ is 0, and whose order at all other places is the minimum of the orders of $h(x,y)$ and $g(x)$. We wish to determine if the divisor is principal, i.e. if there is a single function that has exactly the same orders as this divisor at all places in our function field. In particular such a function is a multiple of our divisor, and we will base our construction on the following proposition and corollary:

Proposition: Let h_1, \dots, h_k be functions and D be a divisor such that $\min_i \text{ord}_p h_i = \text{ord}_p D$ at all finite places. Then the ideal generated by the h_i over the ring of integral functions coincides with the multiples of D except at infinity.

Proof: This is a restatement of the isomorphism between fractional ideals in an algebraic function field and multiples of divisors ignoring places at infinity. [19]

Using the proposition we see that our desired principal generator will be an integral linear combination of $h(x,y)$ and $g(x)$. Every such linear combination will be a multiple of D except at infinity but most of them will have poles at infinity. If we can find one that has no poles at infinity then it will be a multiple of D everywhere. Since the degree of D is zero, any such function must have orders exactly equal to those specified by D and is thus uniquely determined up to scaling by a constant. This proves the following:

Corollary: If D is a divisor of degree 0 with no places at infinity, and h_1, \dots, h_k are as in the proposition, then D is principal if and only if the ideal generated by the h_i has an element with no poles at infinity.

In chapter 2 we showed how to compute an integral basis $[w_1, \dots, w_n]$ for our function field. Any integral multiple of $g(x,y)$ can thus be written as $\sum a_i w_i g$ where $a_i \in K[x]$. Thus we can rewrite our ideal as the $K[x]$ -module generated by $(w_1 g, \dots, w_n g, w_1 h, \dots, w_n h)$. We need to determine whether this module contains a function that has no poles at infinity. If we have a $K[x]$ -module basis for our ideal that is normal at infinity then a linear combination of basis elements is integral at infinity if and only if each summand is. Since each summand is a polynomial times a basis element and a polynomial can never have a zero at infinity, if the ideal contains a function with no poles at infinity, one of the basis elements must have no poles at infinity. Thus after computing a normal basis for our ideal, we only have to check whether any of the basis elements have no poles at infinity.

Theorem: If D is a divisor of degree zero with no places at infinity, then D is principal if and only if a normal basis for the ideal of multiples of D except at infinity has an element that is regular at infinity.

SECTION 2. GOOD REDUCTION

The algorithm in the previous section will enable us to determine whether any given divisor is principal, but we need to know whether there is some power of the divisor that is principal. This problem that was thought to be insoluble in the early 1900's will be solved by the technique of "good reduction". Function fields in one variable whose constant fields are finite fields have the property that any divisor of degree zero has some power that is principal. Thus if we start testing the successive powers of a divisor, we are guaranteed that this process will terminate. Repeated applications of the algorithm in the previous section will enable us to determine the minimum power of any divisor that is principal. Thus we need to be able to reduce the constant field of our function field to a finite field in such a way that we can calculate the order of a divisor from the order of its image.

If we assume the defining polynomial for our function field $f(x,y)$ is monic in y and has integer coefficients, then its coefficients can be reduced modulo p , a prime integer. As long as the reduced polynomial remains irreducible, it defines a function field over the finite field \mathbb{Z}/p . If the genus of the reduced field is the same as the original then we will say that our original field has "good reduction" modulo p . It can be shown [19] for any particular defining polynomial, $f(x,y)$, $Q(x,y)$ has good reduction at all but finitely many primes.

More generally if we have a discrete valuation of our coefficient field, we can choose a defining polynomial for our function field that is monic and such that all of its coefficients are contained in the valuation ring. Then we can apply the natural homomorphism to the residue class field of the valuation. If the resulting polynomial remains absolutely irreducible and the genus of the reduced function field is unchanged, then we say we have "good reduction" at that valuation. Again we will have good reduction at "almost all" valuations. In the rest of this chapter when we say reduction modulo p , we mean reduction by a discrete valuation of the constant field, which extends the natural p -adic valuation of the rationals.

If we have good reduction then there is a homomorphism between the group of divisors of $K(x,y)$ and their images. Since principal divisors map to principal divisors, we in fact have a homomorphism of the quotient group of divisors modulo principal divisors. Let D be a divisor of $K(x,y)$ that has finite order n , i.e. such that D^n is principal. If n is relatively prime to the characteristic of the reduced field then a consequence of good reduction is that the order of the image of D under the reduction is also n . This is a consequence of the key theorem observed by Risch [40] and Dwork and Baldassarri [4] that makes "good reduction" useful.

Theorem: The homomorphism between divisor class groups under good reduction is an isomorphism when restricted to divisors whose orders are relatively prime to the characteristic of the reduced function field.

Let p be the characteristic of the reduced field,

Corollary: if the divisor D has order $p^k n$ where $\gcd(n,p)=1$, then the reduction of D must have order $p^j n$ for some $j \leq k$.

Proof: Let the order of the reduction be $p^j m$. Since reduction is a group homomorphism, we must have $m | n$ and $j \leq k$. Since D^{p^k} has order exactly n , its reduction must have order exactly n . But the order of its reduction is a divisor of m and thus $n | m$ and so finally we have $n = m$.

We have shown that good reduction preserves the "prime to p " part of the orders of divisors, but we need a way to recover the entire order of the divisor. The solution to this problem is to use two different reductions whose residue fields have different characteristics. Let p and q be the characteristics of the residue class fields of two different valuations each of which gives good reduction. If the order of a divisor D is $np^j q^k$ where n is relatively prime to both p and q , reducing D by the first valuation enables one to determine n and k while the second valuation provides a determination of j , thus completing the computation of the order of D . Note that if the two values of n obtained from the two reductions do not match, then D could not have finite order, and we have a simple test that will frequently yield an early termination to our computation.

SECTION 3. ALGORITHMS FOR DIVISOR REDUCTION

The previous section discussed the properties of good reduction stemming from any discrete valuation of our constant field. In general our constant field will be a finitely generated extension of \mathcal{Q} . We can assume this is presented as a sequence of transcendental extensions followed by a single algebraic extension. If n is the transcendence degree of our constant field over \mathcal{Q} , then in order to reduce our constant field to a finite field we need to perform $n + 2$ reductions as outlined in the previous section. The first n reductions will replace parameters by integer values. After these substitutions the polynomial defining our constant field may no longer be irreducible. The factors correspond to different choices of discrete valuations we can make. So we might as well choose the least degree factor. At this point we have reduced our constant field to a finite algebraic extension of \mathcal{Q} , and if all the reductions were good all divisors of finite order will map to divisors of exactly the same order (since the characteristic is still zero). Finally we need to choose two prime integers, p and q , and perform reductions modulo p and then modulo q . Instead of testing whether each reduction step is “good” it is enough that the reduction from the original coefficient field to the finite fields are “good”. This guarantees that all the intermediate reductions were safe.

We now need algorithms for reducing various objects of interest and guaranteeing that our reductions are “good”. For simplicity we will choose all our reductions so that the denominators and leading coefficients of all polynomials of interest do not vanish, i.e. the defining polynomial for our function field and the components of our divisor models and integral basis. Again this will be true for “almost all” reductions. We can use the integral basis algorithm in Chapter 2 to test both that the defining polynomial remains absolutely irreducible and that the genus is unchanged. We merely need verify that our original integral basis for our function field is still an integral basis, and that there are no non-constant functions with no poles. The second condition guarantees absolute irreducibility and the first implies that the genus remains unchanged.

Given that our function field reduces well, we now ask whether our divisor descriptors reduce to locally principal models for the reduction of the desired divisor. Eichler [19] shows that the image of a principal divisor is the divisor associated with the reduction of the generator as long as the generator has a well-defined non-zero reduction. He also shows that this can be accomplished by a proper scaling of the generator. In our situation we have a divisor descriptor of the form $(g(x,y),p(x))$ where g provides a locally principal model at places over the roots of p and the divisor has order zero everywhere else. Thus p defines the projection of the support of the divisor. We assume that g and p have been appropriately scaled so they have well-defined non-zero reductions. We first require that our reduced divisor also have zero order at places over infinity. Since the reduction of p defines the projection of the support of the reduced divisor, this is equivalent to the non-vanishing of the leading coefficient of p . The remaining problem comes from the zeros and poles of g that did not lie above roots of p . It is possible that they may reduce to a place that is above a root of p . In order to ensure that our the reduction of our divisor model does in fact reduce to a model of the desired divisor, we must guarantee this doesn't happen. We compute another polynomial $q(x)$ which is zero precisely at the finite zeros and poles of g that do not lie above roots of p . By definition p and q have no common zeros. Our divisor model reduces well if this is still true for the reductions of p and q . If this condition fails to hold, we can either choose a different reduction or choose a different g as a model for the desired divisor. Assuming we have chosen g to have all of its extraneous poles at infinity, we can obtain another model for the divisor by taking $g + c \cdot p^k$ where c is a constant and k bounds the order of of the zeros of our divisor. With a g of this form, q represents the locations of the finite zeros of g that are not above roots of p . The zeros of g are the poles of $1/g$ and these can be found by representing $1/g$ in terms of our integral basis and computing a common denominator $d(x)$ of the rational function coefficients. $d1 = d/\gcd(d,d')$ has the same roots with multiplicity one. We can now compute $q = d1/\gcd(d1,p)$. The test that our divisor description reduces well is then simply that the leading coefficient of p doesn't vanish and that p and q remain relatively prime.

The preceding discussion can be re-expressed using the standard terminology from algebraic geometry [22]. A locally principal or Cartier divisor is usually specified by giving a collection of open sets U_i and functions f_i such that the open sets cover the Riemann surface for our function field, and f_i/f_j has no zeros or poles on $U_i \cap U_j$. In our case it is convenient to define two closed sets V_1 and V_2 which are the complements of U_1 and U_2 respectively. V_1 is the set of places over the roots of q and at infinity. V_2 is the set of places over the roots of p . We define $f_1 = g$ and $f_2 = 1$. The condition that U_1 and U_2 are an open cover is equivalent to the statement that closed sets V_1 and V_2 have no points in common. By construction g has no zeros or poles on $U_1 \cap U_2$, i.e. the complement of $V_1 \cup V_2$. Thus our initial divisor description yields a well-defined Cartier divisor. We now check whether this remains true after "reduction modulo p ." Again by construction $g = f_1/f_2$ has no zeros or poles except at places above roots of p or q or infinity. The condition that V_1 and V_2 have no points in common translates to p and q remaining relatively prime and no roots of p moving to infinity, i.e. the leading coefficient of p remaining non-zero. Thus we again arrive at the same conditions for good reduction of our divisor description.

CHAPTER 7 CONCLUSIONS AND FUTURE RESEARCH

This thesis has presented an algorithm for integrating algebraic functions that is “rational” in the sense that no algebraic extensions are made beyond those required to express the answer. This yields a more efficient and direct solution to the problem than the previous approach presented by Davenport and Risch. We have stressed the analogies between algebraic function integration and the well-known techniques for rational function integration. The main source of complications comes from the lack of unique factorization in an algebraic function field. We have restored unique factorization by working with ideals. The gcd of two elements is represented by the ideal they generate. After performing the necessary arithmetic on these ideals, we determine whether some power of the resulting ideal is principal giving rise to a logarithmic term in the answer. This ability to give an explicit presentation for an ideal or the associated divisor allows us to compute its order exactly instead of merely finding an upper bound. We perform more work in algebraic function fields over finite fields in order to perform less work over our original constant field where operations are much more costly.

Another central theme in our approach is the iterative reduction of singularities of the integrand. By developing a Hermite-like reduction we are able to reduce non-integrable problems to a simpler form. In fact we are able to prove as Hermite does that one can always reduce the integral of an algebraic function to one which has only simple poles at finite places.

The fundamental construction that we use is the integral basis. This is used to determine the actual singularities of the integrand, to find principal generators for ideals, and to test our function field for “good reduction.” Two other useful applications of an integral basis are computing the genus of our function field and verifying whether our defining polynomial is absolutely irreducible. The existence of this basis allows us to reduce many problems to elementary row operations on matrices of rational functions. If a more efficient algorithm can

be found for computing integral bases, then our entire algorithm will be similarly improved. The integral basis provides us with a global affine non-singular model for our function field and replaces the Puiseux series and Coates' algorithm used by Davenport. A "rational" algorithm to find the multiples of a divisor would yield a similar improvement in Davenport's implementation.

The natural extension of this work would involve handling mixed towers of algebraic, exponential and logarithmic extensions, i.e. general elementary function fields. In the appendix I have given a partial algorithm for the case where you have a purely transcendental tower followed by an unnested radical extension. As Davenport observes in [15] appendix 4, Risch's original transcendental algorithm holds over any differential function field in which one can compute integrals and solve first order linear differential equations. Once you have computed bounds for the orders of the poles in the differential equations, both the algorithms presented here and those of Davenport would suffice for finding the solutions. Thus we can now also perform integrations over any elementary transcendental extension of an algebraic function field.

Extending this work for general elementary function fields may be simpler if first performed using local power series expansions as advocated by Risch in [39]. Once this local approach is well understood, then perhaps a global "rational" approach can be developed. The difficulties seem to involve dealing with places at infinity since one can no longer simply transform infinity to a finite place without creating a non-elementary function field tower. Also in attempting to solve a first order linear differential equation using "rational" techniques, one seems to generate a system of coupled differential equations over the function field one level down. This "reduction" seems to lead to an apparently more difficult problem. The development of algorithms for finding solutions to systems of first order linear differential equations over a given function field could yield a very elegant solution to the general problem of integration in "finite terms." [16]

APPENDIX A. INTEGRATION OF SIMPLE RADICAL EXTENSIONS

Although Risch has presented an outline of an algorithm for integrating mixed towers of algebraic and transcendental elementary functions in [39] and [40]; unfortunately his algorithms require considerably more complex machinery than his earlier ones for purely transcendental functions [38]. Moses' implementation of the transcendental case [35] demonstrated its practicality, whereas there are as yet no implementations for Risch's more general algorithm [39].

This appendix will show how a combination of Risch's earlier techniques and the algorithms presented in this thesis can be generalized to begin to handle mixtures of transcendentals and unnested radicals. While this may seem a severe restriction, perusing an integral table such as [7] will show that fewer than 1 % of the problems are excluded.

We will assume the reader is familiar with the terminology and results of [38]. We will use the term field to mean a differential field of characteristic 0. If F is a field with y algebraic over F of degree n such that $y^n \in F$, then we will call $F(y)$ a simple radical extension of F . Any element of $F(y)$ can be written as a polynomial in y of degree $n - 1$ with coefficients in F .

SECTION 1. STRUCTURE THEOREMS

Our first result will be a refining of Liouville's structure theorem for simple radical extensions. Let F be any differential field with K its field of constants and y radical over F . Risch's Strong Liouville Theorem states that $g \in F(y)$ is integrable if and only if there is a

$v \in F(y)$, $c_i \in K(d)$, and w_i in $F(y, d)$ with d algebraic over K such that:

$$f = v' + \sum c_i \frac{w_i'}{w_i} \quad (1)$$

We will call v the rational part of the integral and the rest the transcendental part of the integral. If we assume that F contains ω a primitive n^{th} root of unity then there is a unique differential automorphism of $F(y)$ over F such that $\sigma(y) = \omega y$. Define the operator

$$T_i = \frac{1}{n} \sum_{j=0}^{n-1} \frac{\sigma^j}{\omega^{ij}}$$

Note that $T_i(y^j) = y^j$ if $i = j$ else 0. Thus letting $g = \sum g_i y^i$ and $v = \sum v_i y^i$, we have $T_i(g) = g_i y^i$. Also noting that T_i commutes with the derivation, $T_i(v') = (v_i y^i)'$. Since T_i sends the transcendental part of the integral to a sum of the same form, by successively applying T_j to equation (1) for $0 \leq j \leq n-1$ we deduce that (1) g is integrable if and only if each $g_i y^i$ is and (2) the rational part of $\int g_i y^i$ is $v_i y^i$.

Now let G be a compositum of simple radical extensions, i.e. $G = F(y_1, \dots, y_k)$ where $y_i^{e_i} \in F$ and $[G:F] = \prod e_i$. Any $g \in G$ can be written as a polynomial in the y_i 's with coefficients in F where the degree of y_i in g is less than e_i . Then by repeating the previous argument for each y_i , one can show g is integrable if and only if each term is integrable. The subfield of G generated by a single such term over F is differentially isomorphic to a simple radical extension of F of degree at most the least common multiple of the e_i 's. Thus integrals over compositums of simple radical extensions can be reduced to integrals over simple extensions frequently of much lower degree.

SECTION 2. A GENERALIZED RISCH ALGORITHM

Let F be arbitrary differential field and $E = F(\theta)$ where θ is transcendental over F and F and E have the same constant subfield. We will additionally assume exactly one of the

following is true:

(1) $\theta' = 1$

(2) $\theta' = v'/v$ for some $v \in F$, i.e. $\theta = \log(v)$

(3) $\theta'/\theta = v'$ for some $v \in F$, i.e. $\theta = \exp(v)$

We will be making use of the fact that $F(\theta)$ is a constructive Euclidean domain. Thus we can compute gcd's and hence square-free decompositions. We are interested in the case where $G = E(y)$ is a simple radical extension of degree n . Additionally we will require that y must depend on θ , i.e. y^n is not in F . By the previous section we are reduced to considering integrands of the form Sy^i with $S \in E$. We will find it convenient to rewrite this as R/y^{n-i} where $R = Sy^n$. By changing our choice of y we can assume an integrand of the form R/y . (Note this may involve changing our value of n). Without loss of generality we can finally assume $y^n = P(\theta) \in F[\theta]$ where P has no factors of multiplicity $\geq n$.

Let $R = A/B$ with $A, B \in F[\theta]$ and B monic. After finding a square-free basis for P and B and performing a partial-fraction decomposition on A/B we can split our integrands into three cases: $C/(V^k y)$, $C/(W^k y)$, and C/y where V is relatively prime to P but W is a square-free factor of P . Unlike the previous section, integrability of R/y does not guarantee integrability of each term in the partial fraction decomposition. However after the splitting we will be able to apply a variety of reduction formulae to these cases. There will be a strong similarity between our algorithms for reducing the integrands and Hermite's algorithm for rational function integration.

Since the case $\theta = \exp(v)$ has additional complications, we will treat it later. For the remaining cases a polynomial Q is square-free if and only if $\gcd(Q, Q') = 1$.

The first problem we encounter is that y' introduces new denominators, so we will choose a polynomial f such that $(f/y)' = g/y$ for some $g \in F[\theta]$. In fact we will need an f of least possible degree. If $P = d \prod P_i^{e_i}$ is our square-free decomposition of P into monic factors with $d \in F$, then define $f = \prod P_i$.

$$(f/y)' = f/y \left(\sum (1 - e_i/n) P_i' / P_i - d' / nd \right) = g/y$$

Clearly $g \in F[\theta]$.

Subsection 2.1. Case 1 $C/(V^k y)$ where $\gcd(V, f) = 1$

We want to find a function whose derivative when subtracted from the integrand will decrease k and not introduce any new denominators. We want a polynomial B such that

$$\frac{(Bf)'}{V^{k-1}y} - \frac{C}{V^k y} = \frac{D}{V^{k-1}y}$$

for some polynomial D . Since

$$\frac{(Bf)'}{V^{k-1}y} = \frac{(1-k)V' Bf}{V^k y} + \frac{B'f + Bg}{V^{k-1}y}$$

Thus we must choose B such that $(1-k)V'fB \equiv C \pmod{V}$. Since $\gcd(fV', V) = 1$ we can find B as long as $k > 1$. Thus we can continue this reduction process until $k = 1$.

Subsection 2.2. Case 2 $C/(W^k y)$ where $W = P_j$

Since W divides f we must start with an apparent denominator of the form $W^k y$. Letting $h = f/W$ we have

$$\left(B \frac{f}{W^k y} \right)' = \frac{Bg - kW'h}{W^k y} + \frac{B'h}{W^{k-1}y}$$

Thus we want to choose B such that $Bg - kW'h \equiv C \pmod{W}$. $W = P_j$ implies $g \equiv (1 - e_j/n)W'h \pmod{W}$. Thus we have $B(1 - k - e_j/n)W'h \equiv C \pmod{W}$. Since $W'h$ is relatively prime to W and $e_j < n$, this equation is solvable for any k . Thus by repeated applications of this reduction step we can eliminate all factors of f from the denominators of our integrands.

Subsection 2.3. Case 3 C/y

Here we will try to find a B such that the reduced integrand has lower degree. Let $B = b_j \theta^j$, with $b_j \in F$.

$$(b_j \theta^j f/y)' = b_j' \theta^j f/y + j b_j \theta^{j-1} \theta' f/y + b_j \theta^j g/y \quad (2)$$

Letting $m = \deg f$ we have two subcases depending on whether d is constant or not. If $d' = 0$ then $\deg(g) = m - 1$ else $\deg(g) = m$. Let $C = \sum c_i \theta^i$.

We will first assume that $d' = 0$. If $b_j' = 0$ then equation (2) has degree $j + m - 1$ else degree $j + m$. Thus equating formally highest degree terms we obtain:

$$c_{j+m} = (j + 1)b_{j+1} \theta' + b_{j+1} \text{lcf}(g) + b_j' \quad (3)$$

where b_{j+1} is a constant. $\text{Lcf}(g)$ is the leading coefficient of g .

$$\text{lcf}(g) = \sum (1 - e_i/n) \text{lcf}(P_i')$$

If $P_i = \theta^k + a_i \theta^{k-1} + \dots$ then $\text{lcf}(P_i') = k\theta' + a_i'$. If $\theta' = 1$ then equation (3) reduces to :

$$c_{j+m} = (j + 1 + \sum \deg(P_i)(1 - e_i/n)) b_{j+1} \quad (4)$$

If $j + 1 \geq 0$ then the coefficient of b_{j+1} will always be nonzero. Thus we can always reduce C until $\deg(C) < m - 1$.

If $\theta = \log(v)$ then equation (3) takes the form

$$c_{j+m} = b_{j+1} \left((j + 1) \frac{v'}{v} + \sum (1 - e_i/n) (\deg(P_i) \frac{v'}{v} + a_i') \right) + b_j' \quad (5)$$

The coefficient of $b_{j+1} v'/v$ in equation (5) is precisely the coefficient of b_{j+1} in equation (4) and is therefore nonzero if $j + 1 \geq 0$. If the original problem is integrable then c_{j+m} must be integrable. Just as in [38] b_{j+1} is uniquely determined since θ is a monomial over F while b_j is determined up to an additive constant. In this case we can reduce C until

$\deg(C) < m - 1$ only if either the original problem is integrable or at least the necessary set of coefficients are integrable.

We will now treat the subcase where d' is nonzero. Note that we must have $\theta = \log(v)$ for this to hold. Here we must first assume that there is no constant s such that sd has an n th root in F . If there is such an s then we can pick a new generator $y(sd)^{-1/n}$ for G that puts us back in the previous subcase. Otherwise we have $\deg(g) = \deg(f)$ and after equating leading terms we have $c_{k+m} = b_k' - \frac{d'}{nd}b_k$. Our assumption about d guarantees that this equation can have at most one solution in F . Thus we need to be able to solve first order linear differential equations for a solution in F . If F is a tower of monomial extensions then [38] shows how to do this. If all the equations we set up have solutions we will reduce C so that $\deg(C) < m$.

Subsection 2.4. $\theta = \exp(v)$

The distinguishing characteristic of the exponential function is that it is a factor of its derivative. Thus we can no longer claim that a square-free polynomial must be relatively prime to its derivative. It will only be necessary to treat factors of the form θ^k specially. We begin by rewriting the square-free decomposition of P . $P = d\theta^j \prod P_i^{e_i}$ where no P_i is divisible by θ . We will again define $f = \prod P_i$, noting that now f is not divisible by θ . We next verify that $(f/y)'$ still is of the form g/y for some $g \in F[\theta]$.

$$\left(\frac{f}{y}\right)' = \frac{f}{y} \left(\sum (1 - e_i/n) \frac{P_i'}{P_i} - j/nv' - \frac{d'}{nd} \right) = \frac{g}{y}$$

After performing a partial-fraction decomposition of the integrand, we can deal with all denominators other than θ just as in the previous cases. Thus we will now assume an integrand of the form $C/(\theta^k y)$, and we again write $C = \sum c_i \theta^i$. We are again trying to decrease k and letting B be an arbitrary polynomial we compute:

$$\left(\frac{Bf}{\theta^k y}\right)' = \frac{B'f + Bg - kv'Bf}{\theta^k y}$$

Requiring the numerator to be congruent to C modulo θ is the same as equating constant terms.

$$c_0 = b_0'f_0 + b_0g_0 - kv'b_0f_0$$

f not divisible by θ implies f_0 is nonzero, thus we can divide through by f_0 . Again θ a monomial will force this equation to have at most one solution in F for $k \geq 0$. As long as the equation continues to have solutions, we can reduce k to 0.

Finally we must deal with integrands of the form C/y in the exponential case. Here $\deg(g) = \deg(f)$ always, and we again assume a solution of the form $\frac{Bf}{y}$ and equate leading terms.

$$c_{k+n} = b_k' + kv'b_k + b_k \text{lcf}(g) \quad (6)$$

$$\text{lcf}(g) = \frac{-j}{n}v' - \frac{d'}{nd} + \sum (1 - e_i/n)\deg(P_i)v'$$

Equation (6) will have at most one solution as long as the coefficient of v' is nonzero. This coefficient is $\frac{k-j}{n} + \sum (1 - e_i/n)\deg(P_i)$. Since the third term is always positive, $k > 0$ or if $j = 0$ then $k \geq 0$ is sufficient to guarantee that v' is present in equation (6).

SECTION 3. SUMMARY AND CONCLUSIONS

The reduction formulae in section 2 have enabled us to find the rational part of our integral if it exists. If the original problem was integrable, all the remaining integrands must generate the transcendental portion of the integral. Note that cases 1 and 2 will always reduce any integrand whether it be integrable or not. In particular, for the case $\theta' = 1$ we see that any integral can be reduced to $\int A/(By)$ where B is square-free and

$$\deg(A) - \deg(B) < \deg(f) - 1.$$

We have shown that the question of computing integrals in $F(\theta, y)$ can be reduced to the problems of computing integrals in F and that of solving first order linear differential equations over F . If F was constructed as a tower of monomial extensions, then [38] shows that these problems are solvable.

We claim that the algorithms presented here form a natural extension to those presented in [38]. By restricting ourselves to this special but very important case, we are able to generate the integral using nothing more than simple polynomial arithmetic. Although our formulas are somewhat more complicated, we require no additional machinery than those necessary in Risch's original approach. [38] We have not worked out the details of the logarithmic part of the integral, but the techniques introduced in chapters 5 and 6 of this thesis should be generalizable to deal with this situation. The major problems involve dealing with poles at infinity which were finessed by a changed of variables in the purely algebraic case. The restriction to unnested radicals guarantees no simple poles at branch places as in the purely algebraic case. Formulae for the residue at infinity similar to equation (7) of Chapter 5 are needed. The fact that the principal parts of a function now only determines it up to a function in one fewer variables instead of up to a constant could cause some additional difficulty. We have presented what we hope are very usable practical algorithms, and we intend to implement them in the near future.

REFERENCES

1. Abhyankar, Sheeram, "Historical Ramblings in Algebraic Geometry and Related Algebra," *Amer. Math. Monthly*, Vol. 83, 6 (1976) pp. 409-448.
2. Atiyah, M.F., and MacDonald, I.G., *Introduction to Commutative Algebra*, Addison Wesley Pub. Co., Reading, Massachusetts, (1969)
3. Artin, Emil, *Algebraic Numbers and Algebraic Functions*, Gordon and Breach, N.Y., (1967)
4. Baldassarri, F. and Dwork, B., "On second order linear differential equations with algebraic solutions," *American Journal of Mathematics*, Vol. 101, 1 (1979) pp. 42-76.
5. Berry, T.G., "On Coates Algorithm," *Sigsam Bulletin*, Vol. 17, 2 (1983) pp. 12-17.
6. Bliss, G.A., *Algebraic Functions*, Dover, (1966), first published as AMS Colloquium vol. XVI (1933)
7. Bois, G. Petit, *Tables of Indefinite Integrals*, Dover, N.Y., (1961)
8. Chevalley, Claude, *Algebraic Functions of One Variable*, Math. Surveys Number VI, American Math. Society, N.Y., (1951)
9. Chou, Tsu-Wu Joseph, *Algorithms for the Solution of Systems of Linear Diophantine Equations*, Ph.D. thesis, University of Wisconsin, (1979)
10. Chow, Wei-Liang, "On the Principle of Degeneration in Algebraic Geometry," *Annals of Mathematics*, Vol. 66, No. 1, pp. 70-79, (1957)
11. Chow, W.-L., Lang, S., "On the Birational Equivalence of Curves under Specialization," *American Journal of Math.*, Vol. 79, pp. 649-652, (1957)
12. Coates, J., "Construction of rational functions on a curve," *Proc. Camb. Phil. Soc.*, Vol 68, pp. 105-123, (1970)
13. Cohn, Harvey, *A Classical Invitation to Algebraic Numbers and Class Fields*, Springer-Verlag, N.Y., (1979)
14. Davenport, James, *On the Integration of Algebraic Functions*, Lecture Notes in Computer Science No. 102, Springer-Verlag, N.Y., (1981)
15. Davenport, James, *Integration Formelle*, (1983).
16. Davenport, J. and Singer, M., Private communication, (1984)
17. Dieudonné, J., "The Historical Development of Algebraic Geometry," *American Mathematical Monthly*, Vol. 79, 8 pp. 827-866, (1972)
18. Duval, Dominique, *Une méthode géométrique de factorisation des polynomes en deux indéterminés*, Institute Fourier, Grenoble, (1983).

19. Eichler, Martin, *Introduction to the Theory of Algebraic Numbers and Functions*, Academic Press, N.Y., (1966)
20. Ford, David James, *On the Computation of the Maximal Order in a Dedekind Domain*, Ph.D. thesis, Ohio State University, Dept. of Mathematics, (1978)
21. Hardy, G.H., *The Integration of Functions of a single variable (2nd ed.)*, Cambridge Tract 2, Cambridge U. Press, (1916)
22. Hartshorne, Robin, *Algebraic Geometry*, Springer-Verlag, (1977)
23. Hasse, Helmut, *Number Theory*, Springer-Verlag, (1980)
24. Hermite, E., "Sur l'intégration des fractions rationnelles," *Nouvelles Annales de Mathématiques*, 2 Sér., 11(1872) pp. 145-148.
25. Herstein, I. N., *Topics in Algebra*, Blaisdell, Waltham, Mass., (1964)
26. Horowitz, E., *Algorithms for Symbolic Integration of Rational Functions*, Ph.D. Thesis, U. of Wisconsin, (1970)
27. Jacobson, Nathan, *Basic Algebra II*, W.H. Freeman and Co., San Francisco, (1980)
28. Lang, Serge, *Algebra*, Addison-Wesley, Reading, Mass., (1971)
29. Lang, Serge, *Algebraic Number Theory*, Addison-Wesley, Reading, Mass., (1970)
30. Lang, Serge, *Introduction to Algebraic Geometry*, Addison-Wesley, Reading, Mass., (1958)
31. Liouville, J., "Premier Memoire sur la Determination des Integrales dont la Valeur est Algebrique," *Journal de l'Ecole Polytechnique*, Vol. 22, (1833) pp. 124-148.
32. Lichtenbaum, Stephen, "Curves over Discrete Valuation Rings," *American Journal of Mathematics*, Vol. 90, pp. 380-405, (1968)
33. Mack, D., *On Rational Integration*, Computer Science Dept., Utah Univ., UCP-38, 1975.
34. Matsuda, Michihiko, *First Order Algebraic Differential Equations*, Lecture Notes in Mathematics, No. 804, Springer-Verlag, (1980)
35. Moses, Joel, "Symbolic Integration: The Stormy Decade," *Communications of the ACM*, Vol. 14, no. 8, pp. 548-560, (1971)
36. Nering, Evar, "Reduction of an Algebraic Function Field Modulo a Prime in the Constant Field," *Annals of Mathematics*, Vol. 67, No. 3, (1958) pp. 590-606.
37. Newman, Morris, *Integral Matrices*, Academic Press, New York, (1972)
38. Risch, R.H., "The Problem of Integration in Finite Terms," *Trans. AMS*, Vol. 139, (1969) pp. 167-189.
39. Risch, R.H., *On the Integration of Elementary Functions which are built up using Algebraic Operations*, Rep. SP-2801/002/00, System Development Corp., Santa Monica, Calif., (1968)

40. Risch, R.H., "The Solution of the Problem of Integration in Finite Terms," *Bulletin A.M.S.*, Vol. 76, (1970) pp. 605-608.
41. Ritt, J.R., *Integration in Finite Terms*, Columbia U. Press, N.Y., (1948)
42. Rosenlicht, Maxwell, "Differential Extension Fields of Exponential Type," *Pacific Journal of Mathematics*, Vol. 57, 1 (1975)
43. Rosenlicht, Maxwell, "Integration in Finite Terms," *American Mathematical Monthly*, Vol. 79, 9 pp. 963-972. (1972)
44. Serre, J.-P., Tate, J., "Good reduction of abelian varieties," *Annals of Math.*, pp. 492-517. (1968)
45. Schmidt, Wolfgang M., *Equations over Finite Fields an Elementary Approach*, Lecture Notes in Mathematics, No. 536, Springer-Verlag, (1976)
46. Shafarevich, I. R., *Basic Algebraic Geometry*, Die Grundlehren der mathematischen Wissenschaften, Band 213, Springer-Verlag, (1974)
47. Shimura, G., Taniyama, Y., *Complex Multiplication of Abelian Varieties and its Applications to Number Theory*, Mathematical Society of Japan, (1961)
48. Trager, B.M., *Algorithms for Manipulating Algebraic Functions*, S.M. thesis, M.I.T., (1976)
49. Trager, B.M., "Algebraic Factoring and Rational Function Integration," *Proc. 1976 ACM Symposium on Symbolic and Algebraic Manipulation*, pp. 219-226, (1976)
50. van der Waerden, B.L., *Modern Algebra*, vol 1, tr. Fred Blum, Frederick Ungar Publishing Co., New York, (1953)
51. Waiker, Robert J., *Algebraic Curves*, Dover, New York, (1962)
52. Weil, André, *Courbes algébriques et variétés abéliennes*, Hermann, Paris, (1971), first published (1948) in *Actualités Scientifiques et Industrielles* nos. 1041 and 1064.
53. Weil, André, *Foundations of Algebraic Geometry*, Amer. Math. Society, N.Y., (1946)
54. Zariski, Oscar, and Pierre Samuel, *Commutative Algebra*, Vol. 1&2, D. Van Nostrand Co., Princeton, N.J., (1958)
55. Zassenhaus, Hans, *Symposia Mathematica*, Vol. XV, "On Hensel Factorization II," Academic Press, London and N.Y., (1975) pp. 499-513.
56. Zassenhaus, Hans, "On the Second Round of the Maximal Order Program," *Applications of Number Theory to Numerical Analysis*, Academic Press, New York, (1972) pp. 389-431.