Decomposition of Algebraic Functions*

Dexter Kozen[1] Susan Landau[2] Richard Zippel[1]

[1] Computer Science Department, Cornell University, Ithaca, NY 14853
[2] Computer Science Department, University of Massachusetts, Amherst, MA 01003

Abstract. Functional decomposition—whether a function $f(x)$ can be written as a composition of functions $g(h(x))$ in a nontrivial way—is an important primitive in symbolic computation systems. The problem of univariate polynomial decomposition was shown to have an efficient solution by Kozen and Landau [9]. Dickerson [5] and von zur Gathen [13] gave algorithms for certain multivariate cases. Zippel [15] showed how to decompose rational functions. In this paper, we address the issue of decomposition of algebraic functions. We show that the problem is related to univariate resultants in algebraic function fields, and in fact can be reformulated as a problem of *resultant decomposition*. We characterize all decompositions of a given algebraic function up to isomorphism, and give an exponential time algorithm for finding a nontrivial one if it exists. The algorithm involves genus calculations and constructing transcendental generators of fields of genus zero.

1 Introduction

Functional decomposition is the problem of representing a given function $f(x)$ as a composition of "smaller" functions $g(h(x))$. Decomposition of polynomials is useful in simplifying the representation of field extensions of high degree, and is provided as a primitive by many major symbolic algebra systems.

The first analyzed algorithms for decomposition of polynomials were provided in 1985 by Barton and Zippel [2] and Alagar and Thanh [1], who gave algorithms for the problem of decomposing univariate polynomials over fields of characteristic zero. Both solutions involved polynomial factorization and took exponential time. Kozen and Landau [9] discovered a simple and efficient polynomial time solution that does not require factorization. It works over fields of characteristic zero, and whenever the degree of h does not divide the characteristic of the underlying field, and provides NC algorithms for irreducible polynomials over finite fields and all polynomials over fields of characteristic zero. Dickerson [5] and von zur Gathen [13] gave algorithms for certain multivariate cases. In addition, von zur Gathen also found algorithms for the case in which the degree of h divides

---

* Algorithmic Number Theory Symposium, Ithaca, New York, May, 1994. Also Department of Computer Science Tech Report TR94-1410, Cornell University, February 1994.

the characteristic of the field [14]. Zippel [15] showed how to decompose rational functions.

In this paper we address the decomposition problem for algebraic functions. We show that the problem bears an interesting and useful relationship to univariate resultants over algebraic function fields, and in fact can be reformulated as a certain resultant decomposition problem: whether some power of a given irreducible bivariate polynomial $f(x,z)$ can be expressed as the resultant with respect to y of two other bivariate polynomials $g(x,y)$, $h(y,z)$. We determine necessary and sufficient conditions for an algebraic function to have a nontrivial decomposition, and classify all such decompositions up to isomorphism. We give an exponential-time algorithm for finding a nontrivial decomposition of a given algebraic function if one exists. The algorithm involves calculating the genus of certain algebraic function fields and constructing transcendental generators of fields of genus zero.

The paper is organized as follows. In §2, we review the basic properties of univariate resultants, state the decomposition problem for algebraic functions, and describe the relationship between the two. In §3 we prove a general theorem that characterizes the set of all possible decompositions of an algebraic function. In §4 we give an exponential time algorithm for the decomposition problem. We conclude in §5 with an example.

2 Resultants and Algebraic Functions

2.1 The Univariate Resultant

Here we review some basic facts about the univariate resultant; see [8, 16] for a detailed introduction.

The *resultant* of two polynomials

$$g(y) = a\prod_{i=1}^{m}(y-\alpha_i) \qquad h(y) = b\prod_{j=1}^{\ell}(y-\beta_j)$$

with respect to y is the polynomial

$$\mathbf{res}_y(g,h) = a^{\ell}b^{m}\prod_{i,j}(\beta_j-\alpha_i) = b^{m}\prod_{h(\beta)=0} g(\beta) \ . \tag{1}$$

The resultant vanishes iff g and h have a common root. It can be calculated as the determinant of the *Sylvester matrix*, a certain $(m+\ell)\times(m+\ell)$ matrix containing the coefficients of g and h.

The following are some useful elementary properties.

$$\begin{aligned}
\mathbf{res}_y(g,h) &= (-1)^{m\ell}\mathbf{res}_y(h,g) \\
\mathbf{res}_y(g_1g_2,h) &= \mathbf{res}_y(g_1,h)\cdot\mathbf{res}_y(g_2,h) \\
\mathbf{res}_y(g,h_1h_2) &= \mathbf{res}_y(g,h_1)\cdot\mathbf{res}_y(g,h_2) \\
\mathbf{res}_y(c,h) &= c^{\ell}
\end{aligned}$$

$$\mathbf{res}_y(g,1) = \mathbf{res}_y(1,h) = 1$$
$$\mathbf{res}_y(g, y-\beta) = g(\beta)$$
$$\mathbf{res}_x(f(x), \mathbf{res}_y(g(x,y), h(y))) = \mathbf{res}_y(\mathbf{res}_x(f(x), g(x,y)), h(y)) \qquad (2)$$

Property (2) is an associativity property. Because of this property, we can write

$$\mathbf{res}_{x,y}(f(x), g(x,y), h(y))$$

unambiguously for the left or right hand side of (2).

We extend the definition to pairs of rational functions as follows. If neither g_1, h_2 nor g_2, h_1 have a common root, define

$$\mathbf{res}_y(\frac{g_1}{g_2}, \frac{h_1}{h_2}) = \frac{\mathbf{res}_y(g_1,h_1) \cdot \mathbf{res}_y(g_2,h_2)}{\mathbf{res}_y(g_1,h_2) \cdot \mathbf{res}_y(g_2,h_1)} \ .$$

This definition reduces to the previous one in the case of polynomials. All the properties listed above still hold, taking $m = \deg g_1 - \deg g_2$ and $n = \deg h_1 - \deg h_2$.

2.2 Resultants and Decomposition

Let K be an algebraically closed field, and let Ω be a *universal field* over K in the sense of van der Waerden [11]; *i.e.*, an algebraically closed field of infinite transcendence degree over K.

Algebraic functions of γ are usually defined as elements of some finite extension of $K(\gamma)$, the field of rational functions of γ. We can also view algebraic functions more concretely as multivalued functions $\Omega \to 2^{\Omega}$ or as binary relations on Ω defined by their minimum polynomials. In the latter view, the decomposition problem is naturally defined in terms of ordinary composition of binary relations:

$$R \circ S = \{(u,w) \mid \exists v \ (u,v) \in R \wedge (v,w) \in S\} \ .$$

Definition 1. For $f(x,z) \in K[x,z]$, let

$$V(f) = \{(\alpha,\gamma) \mid f(\alpha,\gamma) = 0\} \subseteq \Omega^2$$

be the affine variety generated by f. A *decomposition* of f is a pair of polynomials $g(x,y) \in K[x,y]$ and $h(y,z) \in K[y,z]$ such that

$$V(f) = \overline{V(g) \circ V(h)} \ ,$$

where the overbar denotes the Zariski closure in Ω^2 (see [6]). □

The Zariski closure is taken in order to account for points at infinity in a composition. An alternative approach would be to consider f as a binary relation on the projective line.

This notion of decomposition is strongly related to the univariate resultant:

$$\begin{aligned} V(g) \circ V(h) &= \{(\alpha,\beta) \mid \exists\beta \ g(\alpha,\beta) = h(\beta,\gamma) = 0\} \\ &= \{(\alpha,\beta) \mid \mathbf{res}_y(g(\alpha,y), h(y,\gamma)) = 0\} \end{aligned}$$

by (1). The following results develop this relationship further.

Lemma 2. *Let* $g(x,y) \in K[x,y]$ *and* $h(y,z) \in K[y,z]$. *Considering* $g(x,y)$ *and* $h(y,z)$ *as polynomials in* y, *let* $g_m(x)$ *and* $h_\ell(z)$ *be their respective lead coefficients. Then*

$$V(\mathbf{res}_y(g,h)) = (V(g) \circ V(h)) \cup V(g_m, h_\ell) \ .$$

Proof. Consider the two expressions

$$\mathbf{res}_y(g(\alpha,y), h(y,\gamma)) \tag{3}$$

$$\mathbf{res}_y(g(x,y), h(y,z))[x := \alpha, z := \gamma] \ . \tag{4}$$

The difference is whether α and γ are substituted for x and z before or after the resultant is taken. We claim that for any α, γ,

(i) if $g_m(\alpha) = h_\ell(\gamma) = 0$, then (4) vanishes;
(ii) if either $g_m(\alpha) \neq 0$ or $h_\ell(\gamma) \neq 0$, then (3) and (4) vanish or do not vanish simultaneously.

In case (i), we have

$$\mathbf{res}_y(g(x,y), h(y,z)) = \det S(x,z) \ ,$$

where $S(x,z)$ is the Sylvester matrix of $g(x,y)$ and $h(y,z)$. Then

$$\mathbf{res}_y(g(x,y), h(y,z))[x := \alpha, z := \gamma] = \det S(\alpha,\gamma) = 0 \ ,$$

since the first row of $S(\alpha,\gamma)$ is the zero vector. In case (ii), say $h_\ell(\gamma) \neq 0$ (the other case is symmetric). Then

$$\begin{aligned} \mathbf{res}_y(g(x,y), h(y,z))[x := \alpha, z := \gamma] &= \mathbf{res}_y(g(x,y), h(y,\gamma))[x := \alpha] \\ &= h_\ell(\gamma)^{\deg_y g(x,y)} \prod_{h(\beta,\gamma)=0} g(\alpha,\beta) \\ \mathbf{res}_y(g(\alpha,y), h(y,\gamma)) &= h_\ell(\gamma)^{\deg_y g(\alpha,y)} \prod_{h(\beta,\gamma)=0} g(\alpha,\beta) \end{aligned}$$

thus both expressions are simultaneously zero or nonzero.

By (i) and (ii),

$$\begin{aligned} V(\mathbf{res}_y(g,h)) &= \{(\alpha,\gamma) \mid \mathbf{res}_y(g(x,y), h(y,z))[x := \alpha, z := \gamma] = 0\} \\ &= \{(\alpha,\gamma) \mid \mathbf{res}_y(g(\alpha,y), h(y,\gamma)) = 0 \vee g_m(\alpha) = h_\ell(\gamma) = 0\} \\ &= (V(g) \circ V(h)) \cup V(g_m, h_\ell) \ . \end{aligned}$$

□

Theorem 3. *Let* $g(x,y) \in K[x,y]$ *and* $h(y,z) \in K[y,z]$ *be irreducible and nondegenerate (i.e., positive degree in each variable). Then*

$$V(\mathbf{res}_y(g,h)) = \overline{V(g) \circ V(h)} \ .$$

Proof. We have $\overline{V(g)\circ V(h)} \subseteq V(\mathbf{res}_y(g,h))$ by Lemma 2 and the fact that $V(\mathbf{res}_y(g,h))$ is Zariski-closed.

Conversely, it follows from the assumption that $g(x,y)$ and $h(y,z)$ are irreducible and nondegenerate that for all α,β,γ such that $g(\alpha,\beta)=h(\beta,\gamma)=0$, either all $\alpha,\beta,\gamma\in K$ or all are transcendental over K. We use this to show that $\mathbf{res}_y(g,h)$ has no factor of the form $u(x)$. Suppose it did. Let $a\in K$ be a root of u (recall that K is algebraically closed). Then $\mathbf{res}_y(g,h)[x:=a]=0$. Let γ be transcendental over K. We have

$$\begin{aligned}
0 &= \mathbf{res}_y(g(x,y),h(y,z))[x:=a,z:=\gamma]\\
&= \mathbf{res}_y(g(x,y),h(y,\gamma))[x:=a]\\
&= h_\ell(\gamma)^m \prod_{h(\beta,\gamma)=0} g(x,\beta)[x:=a]\\
&= h_\ell(\gamma)^m \prod_{h(\beta,\gamma)=0} g(a,\beta)\ ,
\end{aligned}$$

thus $g(a,\beta)=h(\beta,\gamma)=0$ for some β. But $a\in K$ and γ is transcendental over K, which contradicts our observation above.

By symmetry, $\mathbf{res}_y(g,h)$ has no factor $v(z)$.

Thus all irreducible factors of $\mathbf{res}_y(g,h)$ are nondegenerate. Let (α,γ) be a generic point of some irreducible component C of $V(\mathbf{res}_y(g,h))$. Then α and γ are transcendental over K. By Lemma 2, $(\alpha,\gamma)\in V(g)\circ V(h)$, so $C\subseteq \overline{V(g)\circ V(h)}$. Since C was arbitrary, $V(\mathbf{res}_y(g,h))\subseteq \overline{V(g)\circ V(h)}$. □

Corollary 4. *Let $f(x,z)$, $g(x,y)$, and $h(y,z)$ be irreducible and nondegenerate. Then g,h give a decomposition of f iff $f^k=\mathbf{res}_y(g,h)$ for some $k>0$.*

Proof. If $f^k=\mathbf{res}_y(g,h)$, then by Theorem 3,

$$V(f)=V(f^k)=V(\mathbf{res}_y(g,h))=\overline{V(g)\circ V(h)}\ .$$

Conversely, if $V(f)=\overline{V(g)\circ V(h)}$, then by Theorem 3, $V(f)=V(\mathbf{res}_y(g,h))$, and $f^k=\mathbf{res}_y(g,h)$ follows immediately from the Nullstellensatz and the assumption that f is irreducible. □

In light of Corollary 4, the *decomposition problem* for algebraic functions becomes:

> Given an irreducible polynomial $f(x,z)$, find polynomials $g(x,y)$ and $h(y,z)$ and a positive integer k such that $f^k=\mathbf{res}_y(g,h)$.

This formulation directly generalizes the definition for polynomials and rational functions: for polynomials $g(y)$ and $h(z)$,

$$x-g(h(z))=\mathbf{res}_y(x-g(y),y-h(z))\ .$$

Under this definition, every bivariate polynomial f is decomposable in infinitely many ways:

$$\mathbf{res}_y(f(x,y^k), y^k - z) = \prod_{\beta^k = z} f(x,\beta^k) = \prod_{\beta^k = z} f(x,z) = f^k \ . \tag{5}$$

However, these decompositions are not optimal in a sense to be made precise. In the next section we will define a notion of *minimality* for decompositions, and show that up to isomorphism there are only finitely many nontrivial minimal decompositions.

2.3 Irreducible Decompositions

A decomposition $f = \mathbf{res}_y(g,h)$ is called *irreducible* if both g and h are irreducible as polynomials in $K[x,y]$ and $K[y,z]$, respectively. By the multiplicativity of the resultant, every decomposition factors into a product of irreducible decompositions.

2.4 Monic Decompositions

A decomposition $f = \mathbf{res}_y(g,h)$ is called *monic* if $g \in K(y)[x]$ and $h \in K(z)[y]$ are monic. The next result says that we can restrict our attention to monic decompositions without loss of generality.

Lemma 5. *Let $f \in K[x,z]$, $g \in K[x,y]$, $h \in K[y,z]$ be nondegenerate, g, h irreducible, f a power of an irreducible polynomial. Let $\widehat{f}$, $\widehat{g}$, and $\widehat{h}$ be the monic associates of f, g, h in $K(z)[x]$, $K(y)[x]$, and $K(z)[y]$ respectively. Then $f = \mathbf{res}_y(g,h)$ iff $\widehat{f} = \mathbf{res}_y(\widehat{g},\widehat{h})$.*

Proof. Let $f_n(z)$, $g_m(y)$, and $h_\ell(z)$ be the lead coefficients of f, g and h, respectively. Let

$$u(z) = \mathbf{res}_y(g_m(y), h(y,z)) \cdot h_\ell(z)^{\deg_y g - \deg_y g_m} \ .$$

Then

$$\mathbf{res}_y(g,h) = \mathbf{res}_y(g_m,h) \cdot \mathbf{res}_y(\widehat{g},h_\ell) \cdot \mathbf{res}_y(\widehat{g},\widehat{h}) = u \cdot \mathbf{res}_y(\widehat{g},\widehat{h}) \ .$$

But since $\widehat{g}$ and $\widehat{h}$ are monic, so is $\mathbf{res}_y(\widehat{g},\widehat{h})$, therefore if $f = \mathbf{res}_y(g,h) = u \cdot \mathbf{res}_y(\widehat{g},\widehat{h})$, then $u = f_n$ and $\widehat{f} = \mathbf{res}_y(\widehat{g},\widehat{h})$.

Conversely, if $\widehat{f} = \mathbf{res}_y(\widehat{g},\widehat{h})$, then $uf = f_n\mathbf{res}_y(g,h)$. Remove common factors to get $vf = w \cdot \mathbf{res}_y(g,h)$, where $v, w \in K[z]$ are relatively prime. Now f has no factor in $K[z]$, so w is a unit. Likewise, as argued in the proof of Theorem 3, $\mathbf{res}_y(g,h)$ has no factor in $K[z]$, so v is a unit. □

2.5 Inseparable Decompositions

In prime characteristic p, a decomposition $f(x,z)^k = \mathbf{res}_y(g(x,y),h(y,z))$ is *separable* if f is separable as a polynomial in $K(z)[x]$, g is separable as a polynomial in $K(y)[x]$, and h is separable as a polynomial in $K(z)[y]$. The following argument shows that we can restrict our attention to separable decompositions without loss of generality.

Any inseparable polynomial $f(x^q,z)$, $q=p^n$, has a nontrivial decomposition

$$f(x^q,z) = \mathbf{res}_y(x^q-y, f(y,z)) \ . \tag{6}$$

The polynomial x^q-y decomposes into the composition of n copies of x^p-y. Also,

$$\begin{aligned}\mathbf{res}_y(g(x,y), y^q-z) &= \mathbf{res}_y(g(x,y),(y-\sqrt[q]{z})^q)\\ &= \mathbf{res}_y(g(x,y), y-\sqrt[q]{z})^q\\ &= g(x,\sqrt[q]{z})^q\\ &= g^{[q]}(x^q,z)\end{aligned} \tag{7}$$

where $g^{[q]}(u,v)$ denotes the polynomial obtained from $g(u,v)$ by raising all the coefficients to the q^{th} power.

Once we have decomposed $f(x^q,z)$ as in (6), we can attempt to decompose $f(y,z)$ further. The following argument shows that we can take this step without loss of generality.

Lemma 6. *Let $f(x,z) = \mathbf{res}_y(g(x,y),h(y,z))$, h monic in y. Let q,r be powers of p. Then*

$$f^{[r]}(x^{qr},z) = \mathbf{res}_y(g(x^q,y),h(y^r,z)) \ .$$

Proof.

$$\begin{aligned}\mathbf{res}_y(g(x^q,y),h(y^r,z)) &= \mathbf{res}_{w,y,u}(x^q-w, g(w,y), y^r-u, h(u,z))\\ &= \mathbf{res}_{w,u}(x^q-w, g^{[r]}(w^r,u), h(u,z)) \qquad \text{by (7)}\\ &= \mathbf{res}_{w,v,u}(x^q-w, w^r-v, g^{[r]}(v,u), h(u,z))\\ &= \mathbf{res}_v(x^{qr}-v, f^{[r]}(v,z))\\ &= f^{[r]}(x^{qr},z) \ .\end{aligned}$$

□

Lemma 7. *If $f(x,z)^k = \mathbf{res}_y(g(x,y),h(y,z))$ is a nondegenerate irreducible decomposition, g is separable in x, and h is separable in y, then f is separable in x.*

Proof. Let γ be transcendental over K. Let β be a root of $h(y,\gamma)$ and let α be a root of $g(x,\beta)$. Then α is a root of $f(x,\gamma)$. Since h is separable in y, the extension $K(\beta,\gamma):K(\gamma)$ is separable. Since g is separable in x, the extension $K(\alpha,\beta,\gamma):K(\beta,\gamma)$ is separable. Combining these extensions, we have that the extension $K(\alpha,\beta,\gamma):K(\gamma)$ is separable, hence $f(x,\gamma)$ is separable. □

Theorem 8. *Let q be a power of p and let $f(x^q, z)^k = \mathbf{res}_y(g(x,y), h(y,z))$ be a monic nondegenerate irreducible decomposition, $f(x,z)$ separable. Then there exists a separable decomposition*

$$f(x,z)^k = \mathbf{res}_y(\widehat{g}^{[s]}(x,y), \widehat{h}(y,z))$$

where $g(x,y) = \widehat{g}(x^r, y)$, $h(y,z) = \widehat{h}(y^s, z)$, and $q = rs$.

Proof. Let r, s be powers of p such that g and h can be written $g(x,y) = \widehat{g}(x^r, y)$, $h(y,z) = \widehat{h}(y^s, z)$ with $\widehat{g}$, $\widehat{h}$ separable. Then $\widehat{g}$, $\widehat{h}$ are also irreducible. By Lemma 6,

$$\begin{aligned}
\mathbf{res}_y(x^q - y, f(y,z)^k) &= f(x^q, z)^k \\
&= \mathbf{res}_y(g(x,y), h(y,z)) \\
&= \mathbf{res}_y(\widehat{g}(x^r, y), \widehat{h}(y^s, z)) \\
&= \mathbf{res}_{y,w}(x^{rs} - y, \widehat{g}^{[s]}(y,w), \widehat{h}(w,z))
\end{aligned}$$

and $\mathbf{res}_w(\widehat{g}^{[s]}(y,w), \widehat{h}(w,z))$ is separable by Lemma 7. Thus $q = rs$ and

$$f(y,z)^k = \mathbf{res}_w(\widehat{g}^{[s]}(y,w), \widehat{h}(w,z)) \ .$$

□

This argument shows that in any irreducible decomposition of f, any inseparability of f must stem from the inseparability of one of the composition factors, and this inseparability ultimately emerges as a composition factor of the form $x^q - y$.

By Theorem 8, we can henceforth assume without loss of generality that all decompositions are separable.

3 A Characterization of All Decompositions

In this section we give a characterization of all possible irreducible decompositions of an algebraic function that can arise. As above, we assume that K is algebraically closed and that Ω is a universal field over K.

Let γ be transcendental over K and let α be a nonconstant algebraic function of γ with monic minimum polynomial $f(x,\gamma) \in K(\gamma)[x]$ of degree n. By results of the previous section, the functional decomposition problem reduces to the problem of finding all monic irreducible decompositions of the form

$$f(x,\gamma)^k = \mathbf{res}_y(g(x,y), h(y,\gamma)) = \prod_{h(\beta,\gamma)=0} g(x,\beta) \ .$$

Moreover, we can assume without loss of generality that $f(x,\gamma)$ is separable.

Let A be the set of conjugates of α over $K(\gamma)$, $|A| = n$. Let **Sym** A denote the field of symmetric functions of A. This is the smallest field containing all the coefficients of $f(x,\gamma)$. Note that **Sym** A properly contains K, for otherwise $f(x,\gamma)$

would factor into linear factors since K is algebraically closed, contradicting the assumption that α is nonconstant.

Now consider the following condition on algebraic functions β of γ:

Condition 9 *The monic minimum polynomial* $g(x,\beta)$ *of* α *over* $K(\beta)$ *divides* $f(x,\gamma)$.

If β is algebraic over $K(\gamma)$, then g exists, since α is algebraic over $K(\gamma)$ and γ is algebraic over $K(\beta)$. A subtle but important point to note is that Condition 9 does not imply that $f(x,\gamma)$ factors over $K(\beta)$. Indeed, $K(\beta)$ need not contain the coefficients of f or f/g. We give an example of this in Section 5. The following theorem states that any β satisfying Condition 9 uniquely determines a monic irreducible decomposition of α; moreover, all monic irreducible decompositions of α arise in this way.

Theorem 10. *Let* α *be an algebraic function of* γ *with monic minimum polynomial* $f(x,\gamma) \in K(\gamma)[x]$ *of degree* n*. Let* β *be algebraic over* $K(\gamma)$ *with monic minimum polynomial* $h(y,\gamma) \in K(\gamma)[y]$ *of degree* ℓ*. Let* $g(x,\beta) \in K(\beta)[x]$ *of degree* m *be the monic minimum polynomial of* α *over* $K(\beta)$*. If* β *satisfies Condition 9, i.e. if* $g(x,\beta)$ *divides* $f(x,\gamma)$*, then*

$$f(x,z)^{\frac{\ell m}{n}} = \mathbf{res}_y(g(x,y), h(y,z))$$

is a monic irreducible decomposition of α*. Moreover, all monic irreducible decompositions of* α *arise in this way.*

Proof. Let A be the set of roots of $f(x,\gamma)$ and let $B_\beta \subseteq A$ be the set of roots of $g(x,\beta)$. If η is a conjugate of β over $K(\gamma)$, let B_η be the set of roots of $g(x,\eta)$. The set B_η is the image of B_β under any Galois automorphism over $K(\gamma)$ mapping β to η. For any such conjugate η, $|B_\eta| = |B_\beta| = m$ and $B_\eta \subseteq A$, since the Galois group over $K(\gamma)$ preserves A setwise.

By the symmetry of the action of the Galois group on A, each $\delta \in A$ occurs in the same number of the B_η, say k. We determine k by counting in two ways the number of pairs (δ,η) such that $\delta \in B_\eta$. First, it is the number of conjugates η of β times the size of each B_η, or ℓm. Second, it is the number of $\delta \in A$ times the number of B_η containing δ, or nk. Equating these two values gives $k = \ell m/n$, the exponent in the statement of the theorem. Moreover, it follows from the same argument that

$$\begin{aligned} f(x,\gamma)^k = \prod_{\delta\in A}(x-\delta)^k &= \prod_{h(\eta,\gamma)=0}\ \prod_{\delta\in B_\eta}(x-\delta) \\ &= \prod_{h(\eta,\gamma)=0} g(x,\eta) = \mathbf{res}_y(g(x,y),h(y,\gamma)) \ . \end{aligned}$$

Since γ is transcendental over K, we might as well replace it with the indeterminate z to get

$$f(x,z)^k = \mathbf{res}_y(g(x,y),h(y,z)) \ . \tag{8}$$

The decomposition is monic and irreducible by definition.

Now we show that every monic irreducible decomposition of α arises in this way. Suppose we have such a decomposition (8). Let β be a common root of $g(\alpha, y)$ and $h(y, \gamma)$. Such a β exists, since $f(\alpha, \gamma)$ vanishes, hence so does the resultant $\mathbf{res}_y(g(\alpha, y), h(y, \gamma))$. Then β is algebraic over $K(\gamma)$ with minimum polynomial $h(y, \gamma)$, $g(x, \beta)$ is the minimum polynomial of α over $K(\beta)$, and

$$f(x, \gamma)^k = \mathbf{res}_y(g(x, y), h(y, \gamma)) = \prod_{h(\eta,\gamma)=0} g(x, \eta) \ .$$

Since $g(x, \beta)$ is one of the factors in the product, it divides $f(x, \gamma)$. $\square$

At this juncture we make a few observations about minimal decompositions and uniqueness.

3.1 Minimal decompositions

There may exist β of arbitrarily high degree over $K(\gamma)$ satisfying Condition 9. For example, for any k, $\beta = \sqrt[k]{\gamma}$ gives the decomposition

$$(x - z)^k = \mathbf{res}_y(x - y^k, y^k - z) \ .$$

This is also the situation with (5) above. However, we can bound the search for a suitable β as follows. Observe that if there exists a β satisfying Condition 9 with factor $g(x, \beta)$ of f, say with roots $B \subseteq A$, then α will have the same degree over any subfield of $K(\beta)$ containing the coefficients of g. Furthermore, any such subfield is again a purely transcendental extension of K by Lüroth's Theorem (see [12, 16]), so a transcendental generator of that subfield would give a decomposition with the same g and smaller degree h and smaller k. For a given g, the degree of h and exponent k are minimized by taking the smallest subfield containing the coefficients of g, namely $\mathbf{Sym}\ B$.

3.2 Nontrivial decompositions

If the minimum polynomial $g(x, \beta)$ of α over $K(\beta)$ is f (as would occur in the case $\beta = \gamma$), then the minimal decomposition with this g occurs when β is a transcendental generator of $\mathbf{Sym}\ A$. Since $\mathbf{Sym}\ A \subseteq K(\gamma)$, β would be a rational function of γ and h would be linear of the form $y - u(\gamma)$, $u \in K(z)$, giving the decomposition

$$f(x, z) = \mathbf{res}_y(g(x, y), y - u(z)) = g(x, u(z)) \ .$$

In this case α is the composition of an algebraic function and a rational function.

In case $g(x, \beta)$ is linear, say $g = x - v(\beta)$, the smallest field containing the coefficients of g is $K(v(\beta))$, so by using $v(\beta)$ instead of β we would obtain the trivial decomposition

$$f(x, z) = \mathbf{res}_y(x - y, h(y, z)) = h(x, z) \ .$$

To find a nontrivial decomposition, we must find a β such that $K(\beta)$ does not contain α.

3.3 Uniqueness up to linear composition factors

The decomposition determined by β essentially depends only on the field $K(\beta)$, not on the choice of transcendental generator β. Any other transcendental generator of $K(\beta)$ is related to β by a nonsingular fractional linear transformation

$$\beta \mapsto \frac{a\beta + b}{c\beta + d} \ , \quad ad - bc \neq 0 \ ,$$

which extends to an automorphism of $K(\beta)$. Any two decompositions defined with respect to two transcendental generators of the same field are equivalent up to invertible composition factors of the form $(cz + d)y - (az + b)$.

4 An Algorithm

As determined in the previous section, up to fractional linear transformations there are only finitely many minimal irreducible monic decompositions of f, at most one for each factor g of f. We have thus reduced the decomposition problem to the problem of finding a subset $B \subseteq A$ (the roots of g) such that the field **Sym** B (the field generated by the coefficients of g) is a purely transcendental extension of K, and then finding a transcendental generator β of **Sym** B. Such a β is automatically algebraic over $K(\gamma)$, since **Sym** $B \subseteq K(A)$, the splitting field of f over $K(\gamma)$.

We must first determine whether f has a factor g whose coefficients lie in a purely transcendental extension of K. Equivalently, we want to know when the field **Sym** B of symmetric functions in the roots B of g is isomorphic to a rational function field over K. This is true iff **Sym** B is of genus zero. Thus the problem reduces to the problem of determining the genus of an algebraic function field.

The following is a synopsis of our algorithm.

Algorithm 11

1. Let g be a nontrivial factor of f. The coefficients of g lie in some finite extension $K(\gamma, \eta)$ of $K(\gamma)$ over which f has a nontrivial factorization. Then g can be written

   $$g(x, \eta, \gamma) = x^m + u_{m-1}(\eta, \gamma)x^{m-1} + \cdots + u_0(\eta, \gamma) \ .$$

   For each such g, perform steps 2 and 3.
2. Construct the field $K(u_0, \ldots, u_{m-1})$. This is the field **Sym** B, where B is the set of roots of g. Pick one of the coefficients of g not in K, say u_0. We have two cases:
   (a) If $K(u_0, \ldots, u_{m-1}) = K(u_0)$, we are done: u is a transcendental generator of **Sym** B.

   (b) If $K(u_0, \ldots, u_{m-1}) \neq K(u_0)$, construct a primitive element θ of the extension such that $K(u_0, \ldots, u_{m-1}) = K(u_0, \theta)$. Compute the genus of $K(u_0, \theta)$ by the Hurwitz genus formula or in some other fashion. An efficient algorithm is given in [3]. If the genus is nonzero, then no decomposition arises from this factor of f. If the genus is zero, compute a rational generator β of $K(u_0, \theta)$. Coates [4], Trager [10], and Huang and Ierardi [7] give efficient algorithms for computing rational generators. The coefficients of g can then be written as rational functions of β. $\square$

3. Let $h(y, \gamma)$ be the minimum polynomial of β over $K(\gamma)$. Return $g(x, y)$ and $h(y, z)$ as the decomposition factors.

Under suitable assumptions about the complexity of operations in K, the complexity of the algorithm as given above is exponential in the worst case, since there are an exponentially many potential factors. For each such factor, the computation for that factor can be performed in polynomial time in the size of the representation of the algebraic numbers needed to express the result, or exponential time in the bit complexity model [7].

5 An Example

The following gives an example of a decomposition involving a β such that $g(x, \beta)$ divides $f(x, \gamma)$, but $f(x, \gamma)$ does not factor over $K(\beta)$. Consider the polynomial

$$f(x, z) = x^4 - zx^2(x+1) + z^3(x+1)^2 \ .$$

Let γ be transcendental over K, and let

$$\begin{aligned} \beta &= \frac{\gamma(1 + \sqrt{1 - 4\gamma})}{2} & \qquad \eta &= \frac{\gamma(1 - \sqrt{1 - 4\gamma})}{2} \\ g(x, y) &= x^2 - y(x+1) & \qquad h(y, z) &= y^2 - zy + z^3 \ . \end{aligned}$$

Then β and η are conjugates over $K(\gamma)$ with minimum polynomial $h(y, \gamma)$, and

$$f(x, \gamma) = g(x, \beta) \cdot g(x, \eta) \ ,$$

thus Theorem 10 says that g and h should give a decomposition of f. Indeed,

$$\mathbf{res}_y(g(x, y), h(y, z)) = \begin{vmatrix} -(x+1) & 0 & 1 \\ x^2 & -(x+1) & -z \\ 0 & x^2 & z^3 \end{vmatrix} = f(x, z) \ .$$

To show $f(x, \gamma)$ does not factor over $K(\beta)$, it suffices to show that its trace γ is not in $K(\beta)$. But γ is a root of the irreducible polynomial $h(\beta, z)$, therefore is algebraic of degree three over $K(\beta)$.

Acknowledgements

We thank John Cremona, Ming-Deh Huang, John Little, Paul Pedersen, Moss Sweedler, Barry Trager, Emil Volcheck, Gary Walsh, and an anonymous reader for valuable comments. The support of the U.S. Army Research Office through the ACSyAM branch of the Mathematical Sciences Institute of Cornell University under contract DAAL03-91-C-0027, the National Science Foundation under grants CCR-9204630 and CCR-9317320, and the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research grant N00014-92-J-1989 is gratefully acknowledged. This research was done while the second author was visiting the Cornell University Computer Science Department.

References

1. V. S. Alagar and M. Thanh, *Fast polynomial decomposition algorithms*, in Proc. EUROCAL85, Springer-Verlag Lect. Notes in Comput. Sci. 204, 1985, pp. 150–153.
2. D. R. Barton and R. E. Zippel, *Polynomial decomposition algorithms*, J. Symb. Comp., 1 (1985), pp. 159–168.
3. G. A. Bliss, *Algebraic Functions*, Amer. Math. Soc., 1933.
4. J. Coates, *Construction of rational functions on a curve*, Proc. Camb. Phil. Soc., 68 (1970), pp. 105–123.
5. M. Dickerson, *Polynomial decomposition algorithms for multivariate polynomials*, Tech. Rep. TR87-826, Comput. Sci., Cornell Univ., April 1987.
6. R. Hartshorne, *Algebraic Geometry*, vol. 52 of Graduate Texts in Mathematics, Springer, 1977.
7. M.-D. Huang and D. Ierardi, *Efficient algorithms for the effective Riemann-Roch problem and for addition in the Jacobian of a curve*, in Proc. 32nd Symp. Found. Comput. Sci., IEEE, November 1991, pp. 678–687.
8. D. Ierardi and D. Kozen, *Parallel resultant computation*, in Synthesis of Parallel Algorithms, J. Reif, ed., Morgan Kaufmann, 1993, pp. 679–720.
9. D. Kozen and S. Landau, *Polynomial decomposition algorithms*, J. Symb. Comput., 7 (1989), pp. 445–456.
10. B. M. Trager, *Integration of Algebraic Functions*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, September 1984.
11. B. L. van der Waerden, *Algebra*, vol. 2, Frederick Ungar, fifth ed., 1970.
12. ———, *Algebra*, vol. 1, Frederick Ungar, fifth ed., 1970.
13. J. von zur Gathen, *Functional decomposition of polynomials: the tame case*, J. Symb. Comput., 9 (1990), pp. 281–299.
14. ———, *Functional decomposition of polynomials: the wild case*, J. Symb. Comput., 10 (1990), pp. 437–452.
15. R. E. Zippel, *Rational function decomposition*, in International Symposium on Symbolic and Algebraic Computation, S. Watt, ed., New York, July 1991, ACM, pp. 1–6.
16. ———, *Effective Polynomial Computation*, Kluwer Academic Press, Boston, 1993.

This article was processed using the LaTeX macro package with LLNCS style