

Polynomial Decomposition Algorithms

Dexter Kozen*

Department of Computer Science
Cornell University
Ithaca, New York 14853

Susan Landau†

Department of Mathematics
Wesleyan University
Middletown, Connecticut 06457

Abstract

We examine the question of when a polynomial f over a commutative ring has a nontrivial functional decomposition $f = g \circ h$. Previous algorithms [2, 3, 1] are exponential-time in the worst case, require polynomial factorization, and only work over fields of characteristic 0. We present an $O(n^2)$ -time algorithm. We also show that the problem is in NC . The algorithm does not use polynomial factorization, and works over any commutative ring containing a multiplicative inverse of r . Finally, we give a new structure theorem that leads to necessary and sufficient algebraic conditions for decomposibility over any field. We apply this theorem to obtain an NC algorithm for decomposing irreducible polynomials over finite fields, and a subexponential algorithm for decomposing irreducible polynomials over any field admitting efficient polynomial factorization.

1 Introduction

In this paper, we address the following question: given a polynomial f with coefficients in a commutative ring K , does it have a nontrivial functional decomposition $f = g \circ h$? If so, can the coefficients of g and h be obtained efficiently?

This problem arises in certain computations in symbolic algebra. The following example, due to Barton and Zippel [2, 3], shows that decomposition can simplify symbolic root-finding. Since the polynomial

$$f = x^6 + 6x^4 + x^3 + 9x^2 + 3x - 5$$

decomposes into $f = g \circ h$, where

$$g = x^2 + x - 5 \text{ and } h = x^3 + 3x ,$$

*Supported by NSF grant DCR-8602663.

†Supported by NSF grants DCR-8402175 and DCR-8301766.

one can solve for the roots α of g , and then solve for the roots of $h - \alpha$ to obtain the roots of f . Many symbolic algebra systems (e.g., MACSYMA, MAPLE, and SCRATCHPAD II) support polynomial decomposition for purposes such as this.

Barton and Zippel [2, 3] present two decomposition algorithms. Some simplifications are suggested by Alagar and Thanh [1]. All these algorithms require K to be a field of characteristic 0, they all use polynomial factorization, and they all take exponential time in the degree of f in the worst case. No decomposition algorithms over fields of finite characteristic, or over more general rings, were known.

Our first result is a polynomial-time algorithm for computing a decomposition $f = g \circ h$ over any commutative ring K containing a multiplicative inverse of the degree of g . The algorithm has several advantages over the algorithms of [2, 3, 1]:

1. It is polynomial-time. A straightforward implementation uses $O(n^2r)$ algebraic operations, where r is the degree of g . All previously known algorithms [2, 3, 1] were exponential. With a bit more care, the complexity can be further reduced to $O(n^2)$ algebraic operations. It also follows from our algorithm and from a general result of Valiant *et al.* [14] that the problem is in NC .
2. It is more general. The algorithms of [2, 3, 1] require K to be a field of characteristic 0. Our algorithm works in any commutative ring K containing a multiplicative inverse of r , the degree of g .
3. It is simpler. The algorithm itself is only a few lines, and requires little more than high school algebra for its analysis, in contrast to [2, 3, 1]. Most significantly, it does not use polynomial factorization.
4. It is faster in practice. A rough implementation has been coded in about 100 lines of MAPLE by Bruce Char at the University of Waterloo. Although no performance figures are yet available, Char reports that it is faster than Barton and Zippel on all examples he has tried. The algorithm has also been implemented by Arash Baratloo, an undergraduate at Cornell. His implementation decomposes a polynomial of degree 10 in less than 1 second of real time.

This algorithm is described in §3.

In the remainder of the paper, we restrict K to be a field, but place no restriction on its characteristic. In §4, we generalize the notion of *block decomposition* of the Galois group of an irreducible polynomial f over a field to the case where f need not be irreducible or separable. We then establish necessary and sufficient conditions in terms of this construct for the existence of a nontrivial decomposition of any polynomial f over any field K , and show how the decomposition problem effectively reduces to the problem of polynomial factorization over K . The complexity of decomposition will depend on the complexity of factoring over K , and in general will be at least exponential.

This result gives more efficient algorithms in the following special cases. In §5, we give an NC algorithm for decomposing irreducible polynomials over any finite field. In §6, we

consider arbitrary fields, and give an $O(n^{\log n})$ sequential-time algorithm for decomposing irreducible polynomials over any field admitting a polynomial-time factorization algorithm.

2 Algebraic Preliminaries

Let

$$f = x^n + a_{n-1}x^{n-1} + \cdots + a_1x + a_0$$

be a monic polynomial of degree greater than one with coefficients in a commutative ring K . A (*functional*) *decomposition* of f is a sequence g_1, \dots, g_k such that $f = g_1 \circ g_2 \circ \cdots \circ g_k$, i.e., $f(x) = g_1(g_2(\cdots g_k(x) \cdots))$. The g_i are called *components* of the decomposition. For any unit a of the ring K , the linear polynomials $ax + b$ and $\frac{1}{a}(x - b)$ are inverses under composition, thus f always has trivial decompositions in which all but one of the components are linear. If all decompositions of f are trivial, then f is said to be *indecomposable*. A *complete decomposition* is one in which each of the components is of degree greater than one and is indecomposable.

Complete decompositions are not unique. Consider the examples

$$\begin{aligned} f \circ g &= f(x + b) \circ (g - b) \\ x^n \circ x^m &= x^m \circ x^n \\ T_n \circ T_m &= T_m \circ T_n \end{aligned}$$

where T_n and T_m are Chebyshev polynomials. In the first example, one decomposition is obtained from the other by inserting linear components. In the second and third, one decomposition is obtained from the other by interchanging commuting components. *Ritt's first theorem* states that a complete decomposition of f is unique up to these ambiguities, provided the characteristic of K is either 0 or finite but greater than the degree of f [13, 6, 8].

If f is monic and has a decomposition $f = g \circ h$, where

$$\begin{aligned} g &= b_r x^r + b_{r-1} x^{r-1} + \cdots + b_0 \\ h &= c_s x^s + c_{s-1} x^{s-1} + \cdots + c_0, \end{aligned}$$

then $g(c_s x)$ and $h(x)/c_s$ are monic and also give a decomposition of f . The leading coefficients of g and h are invertible, since $b_r c_s^r = 1$. It follows that any complete decomposition is equivalent to a decomposition in which all the components are monic.

It will follow from Algorithm 3.4 below that if $f \in K[x]$ is decomposable over any extension of K , then it is decomposable over K , provided K contains a multiplicative inverse of r . This generalizes a result of Fried and MacRae [8], who show that if K is a field and $f \in K[x]$ is decomposable over an algebraic extension of K , then it is decomposable over K .

3 Decomposition over Commutative Rings

The equation $f = g \circ h$, where g , h , and f are monic of degree r , s , and $n = rs$ respectively, results in rs equations in $r + s$ unknowns. When $r \geq s$, the longest of these equations can be

shown to have at least as many terms as the number of partitions of s , which is $2^{\Omega(\sqrt{s})}$ [10]; thus a naive implementation results in an exponential algorithm. This led Barton and Zippel to develop two other algorithms which are more efficient in practice, but still exponential in the degree of f in the worst case. These algorithms involve polynomial factorization and assume that K is a field of characteristic 0.

In this section we give an elementary algorithm requiring at most $O(n^2r)$ algebraic operations in the worst case. The algorithm does not involve polynomial factorization, and works in any commutative ring containing a multiplicative inverse of r . We then show how an alternative implementation using interpolation further reduces the complexity to $O(n^2)$ algebraic operations.

We begin with an elementary lemma.

Lemma 3.1 *Let K be a commutative ring, and let $f_1, f_2, g \in K[x]$ be monic. If f_1 and f_2 agree on their first k coefficients, then so do f_1g and f_2g .*

Let

$$\begin{aligned} f &= x^{rs} + a_{rs-1}x^{rs-1} + \cdots + a_0 \\ g &= x^r + b_{r-1}x^{r-1} + \cdots + b_0 \\ h &= x^s + c_{s-1}x^{s-1} + \cdots + c_0 \end{aligned}$$

such that $r, s < rs = n$ and $f = g \circ h$. We can assume without loss of generality that $c_0 = 0$, since if $f = g \circ h$, then $f = g(x + c_0) \circ (h - c_0)$.

Let g factor in an extension of K as

$$g = \prod_{i=1}^r (x - \beta_i) .$$

Then

$$f = g(h) = \prod_{i=1}^r (h - \beta_i) .$$

By r applications of Lemma 3.1, we have

Lemma 3.2 *The products h^r and $f = \prod_{i=1}^r (h - \beta_i)$ agree on their first s coefficients.*

Let

$$q_k = x^s + c_{s-1}x^{s-1} + \cdots + c_{s-k}x^{s-k} , \quad 0 \leq k \leq s .$$

Then $q_0 = x^s$, $q_s = q_{s-1} = h$, and

$$q_k = q_{k-1} + c_{s-k}x^{s-k} , \quad 1 \leq k \leq s .$$

Again by Lemma 3.1, we have

Lemma 3.3 *The powers h^r and q_k^r agree on their first $k + 1$ coefficients, $0 \leq k \leq s$.*

The $k + 1^{\text{st}}$ coefficient of q_k^r is the coefficient of x^{rs-k} ; but since

$$\begin{aligned} q_k^r &= (q_{k-1} + c_{s-k}x^{s-k})^r \\ &= q_{k-1}^r + rc_{s-k}x^{s-k}q_{k-1}^{r-1} + p \end{aligned}$$

where $p \in K[x]$ is of degree at most $rs - 2k$, this coefficient is just $d_k + rc_{s-k}$, where d_k is the coefficient of x^{rs-k} in q_{k-1}^r and rc_{s-k} is the coefficient of x^{rs-k} in $rc_{s-k}x^{s-k}q_{k-1}^{r-1}$. By Lemmas 3.2 and 3.3, this agrees with a_{rs-k} , the $k + 1^{\text{st}}$ coefficient of f , $1 \leq k \leq s - 1$. Thus if the earlier coefficients $c_{s-1}, \dots, c_{s-k+1}$ of h are known, then c_{s-k} can be determined by computing

$$c_{s-k} = \frac{a_{rs-k} - d_k}{r}, \quad 1 \leq k \leq s - 1,$$

provided K contains a multiplicative inverse of r . This gives rise to the following decomposition algorithm:

Algorithm 3.4

/ coefficients of h */*

set $q_0 := x^s$ and compute $q_0^i := x^{is}$, $0 \leq i \leq r$;

for $k := 1$ to $s - 1$ do

/ assume $q_{k-1}^0, \dots, q_{k-1}^r$ are known from the previous step */*

set $d_k :=$ the coefficient of x^{rs-k} in q_{k-1}^r ;

set $c_{s-k} := \frac{1}{r}(a_{rs-k} - d_k)$;

/ then c_{s-k} = the coefficient of x^{s-k} in h */*

calculate c_{s-k}^j , $0 \leq j \leq r$;

set $q_k^j := \sum_{i=0}^j \binom{j}{i} c_{s-k}^i x^{i(s-k)} q_{k-1}^{j-i}$, $0 \leq j \leq r$;

/ then $q_k^j = (q_{k-1} + c_{s-k}x^{s-k})^j$ */*

/ coefficients of g */*

Solve the triangular system

$$Ab = a, \tag{1}$$

where A_{ij} is the coefficient of x^{is} in h^j , $0 \leq i, j \leq r$, $b = (b_0, \dots, b_r)$, and $a = (a_0, a_s, \dots, a_{rs})$.

/ consistency check */*

Compute $g \circ h$ and compare with f .

The calculation of $h = q_{s-1}$ takes $O(n^2r)$ algebraic operations. This also produces the entries of the matrix A which is used in the calculation of g . The matrix A is triangular with all diagonal elements 1, thus is easily invertible over K in time $O(r^2)$.

The consistency check is necessary, since the system is overconstrained (there are rs equations in $r + s$ unknowns), and the g and h produced by the above algorithm may not compose to give f . This is easily done in $O(nr)$ time, using the powers of h already computed.

It is easily shown that the above procedure produces an expression for each coefficient of h and g that is of degree at most r in each coefficient of f . Since these coefficients are calculated by a straight-line program involving only algebraic operations, it follows from a general simulation result of Valiant *et al.* [14] that the algorithm can be parallelized so as to run in $\log^{O(1)} n$ parallel time, using polynomially many processors, therefore the problem is in the complexity class NC .

We now show that the $O(n^2r)$ sequential bound can be improved to $O(n^2)$, provided the ring contains at least $n + 1$ elements, using polynomial interpolation. Let ξ_0, \dots, ξ_n be $n + 1$ distinct elements of K , and let V be the Vandermonde matrix $V_{ij} = \xi_i^j$. Computing the transform $\widehat{f} = Vf$ amounts to evaluating f at the points ξ_i (here f is represented by its vector of coefficients $(a_0, \dots, a_{n-1}, 1)$, and Vf is just the matrix-vector product). We carry out the computations of Algorithm 3.4 in the transformed space to obtain a vector of values \widehat{h} , and then interpolate to obtain h by solving the system $Vh = \widehat{h}$.

In the following, let e_i denote the vector whose i^{th} entry is 1 and whose other entries are 0. We will need to perform the analog of the operation, “let d be the coefficient of x^i in the polynomial p ” in the transformed space. This is equivalent to computing the inner product $e_i \cdot p$, so if $\widehat{p} = Vp$, then

$$\begin{aligned} d &= e_i \cdot (V^{-1}\widehat{p}) \\ &= (e_i^T V^{-1}) \cdot \widehat{p}, \end{aligned}$$

where $e_i^T V^{-1}$ is the i^{th} row of V^{-1} . This can be done in linear time, provided we have the matrix V^{-1} .

Algorithm 3.5

compute the transform and inverse transform matrices V and V^{-1} ;

set $f := Vf$;

/ coefficients of h */*

set $\widehat{q}_0 := (\xi_0^s, \xi_1^s, \dots, \xi_n^s)$;

/ then $\widehat{q}_0 = Vq_0$ */*

for $k := 1$ to $s - 1$ do

/ assume $\widehat{q}_{k-1} = Vq_{k-1}$ is known from the previous step */*

compute $\widehat{q}_{k-1}^r =$ the pointwise r^{th} power of \widehat{q}_{k-1} ;

set $d_k := (e_{rs-k}^T V^{-1}) \cdot \widehat{q}_{k-1}^r$;

/ then $d_k =$ the coefficient of x^{rs-k} in q_{k-1}^r */*

set $c_{s-k} := \frac{1}{r}(a_{rs-k} - d_k)$;

/ then $c_{s-k} =$ the coefficient of x^{s-k} in h */*

set $\widehat{q}_k := \widehat{q}_{k-1} + c_{s-k} V e_{s-k}$;

/* note that Ve_{s-k} is the $s - k^{\text{th}}$ column of V , and $c_{s-k}Ve_{s-k} =$ the transform of $c_{s-k}x^{s-k}$ */
 set $h = V^{-1}\widehat{q}_{s-1}$;
 /* the calculation of the coefficients of g and the consistency check are the same as in Algorithm 3.4 */

The matrices V and V^{-1} can be computed in time $O(n^2)$ [12]. It is also straightforward to show that these computations can be performed in NC . The only nonlinear step in the loop is the calculation of the pointwise r^{th} power of \widehat{q}_{k-1} ; this takes $O(n \log r)$ steps. Thus the total time in the loop is $O(ns \log r) \leq O(n^2)$. We have shown

Theorem 3.6 *Let K be a commutative ring containing a multiplicative inverse of r . There is an $O(n^2)$ algorithm to determine whether a given monic polynomial f of degree n over K has a nontrivial decomposition $f = g \circ h$ over K , $\deg g = r$, and to produce the coefficients of g and h if such a decomposition exists.*

4 A Structure Theorem

The algorithms of §3 require that the ring K contain certain multiplicative inverses. Barton and Zippel [2, 3] and Alagar and Thanh [1] consider only fields of characteristic 0. In this section we give a structure theorem that provides a necessary and sufficient condition for the decomposability of any polynomial over any field. These conditions involve a generalization of the notion of block decomposition of the Galois group of an irreducible polynomial over a field.

Let K be a field of arbitrary characteristic. Let $f \in K[x]$ be monic of degree $n = rs$, not necessarily irreducible or separable. Let \widehat{K} denote the splitting field of f . Let \mathcal{G} denote the Galois group of \widehat{K} over K . The following definition reduces to the usual notion of block decomposition for f irreducible and separable.

Definition 4.1 A *block decomposition* for f is a multiset Δ of multisets of elements of \widehat{K} such that

1. $f = \prod_{A \in \Delta} \prod_{\alpha \in A} (x - \alpha)$
2. if $\alpha \in A \in \Delta$, $\beta \in B \in \Delta$, and $\sigma \in \mathcal{G}$ such that $\sigma(\alpha) = \beta$, then

$$B = \{\sigma(\gamma) \mid \gamma \in A\} .$$

A block decomposition Δ is an $r \times s$ *block decomposition* if $|\Delta| = r$ and $|A| = s$ for all $A \in \Delta$. □

Note that non-irreducible polynomials may have block decompositions that are not $r \times s$ for any r, s .

Let c_k^m denote the k^{th} elementary symmetric function on m -element multisets:

$$c_k^m(A) = \sum_{B \subseteq A, |B|=k} \prod B .$$

By convention, $c_0^m = 1$.

Theorem 4.2 *Let $f \in K[x]$ be monic of degree $n = rs$. The following two statements are equivalent:*

- (i) $f = g \circ h$ for some $g, h \in K[x]$ of degree r and s , respectively;
- (ii) there exists an $r \times s$ block decomposition Δ for f such that

$$c_k^s(A) = c_k^s(B) \in K, \text{ for all } A, B \in \Delta, 0 \leq k \leq s - 1 .$$

Proof. (i) \rightarrow (ii). Let Γ be the multiset of roots of g . Then

$$g = \prod_{\gamma \in \Gamma} (x - \gamma) ,$$

and

$$f = g(h) = \prod_{\gamma \in \Gamma} (h - \gamma) .$$

For each $\gamma \in \Gamma$, let A_γ be the multiset of roots of $h - \gamma$. Let $\Delta = \{A_\gamma \mid \gamma \in \Gamma\}$. Then

$$f = \prod_{\gamma \in \Gamma} (h - \gamma) = \prod_{\gamma \in \Gamma} \prod_{\alpha \in A_\gamma} (x - \alpha) .$$

For each $\gamma \in \Gamma$, $(-1)^k c_k^s(A_\gamma)$ is the coefficient of x^{s-k} in $h - \gamma$, $0 \leq k \leq s - 1$, therefore

$$c_k^s(A_\gamma) \in K, \quad 0 \leq k \leq s - 1 .$$

It remains to show that Δ is a block decomposition. Suppose that $\gamma, \delta \in \Gamma$, $\alpha \in A_\gamma$, $\beta \in A_\delta$, $\sigma \in \mathcal{G}$, and $\sigma(\alpha) = \beta$. Since $h(\alpha) = \gamma$ and $h(\beta) = \delta$,

$$\sigma(\gamma) = \sigma(h(\alpha)) = h(\sigma(\alpha)) = h(\beta) = \delta .$$

Extending $\sigma : \widehat{K} \rightarrow \widehat{K}$ uniquely to $\sigma : \widehat{K}[x] \rightarrow \widehat{K}[x]$ by taking $\sigma(x) = x$, we have

$$\begin{aligned} \prod_{\alpha \in A_\gamma} (x - \sigma(\alpha)) &= \prod_{\alpha \in A_\gamma} (\sigma(x) - \sigma(\alpha)) \\ &= \sigma\left(\prod_{\alpha \in A_\gamma} (x - \alpha)\right) \\ &= \sigma(h - \gamma) \\ &= h - \sigma(\gamma) \\ &= h - \delta \\ &= \prod_{\beta \in A_\delta} (x - \beta) . \end{aligned}$$

Thus $A_\delta = \{\sigma(\alpha) \mid \alpha \in A_\gamma\}$.

(ii) \rightarrow (i). Given Δ , let $A \in \Delta$ and let

$$h = \sum_{k=0}^{s-1} (-1)^k c_k^s(A) x^{s-k} \in K[x].$$

The choice of A does not matter, by the assumption of (ii). For each $A \in \Delta$, let

$$\gamma_A = (-1)^{s+1} \prod_{\alpha \in A} \alpha = (-1)^{s+1} c_s^s(A) \in \widehat{K},$$

and let

$$g = \prod_{A \in \Delta} (x - \gamma_A).$$

Then

$$h - \gamma_A = \sum_{k=0}^s (-1)^k c_k^s(A) x^{s-k} = \prod_{\alpha \in A} (x - \alpha),$$

and

$$\begin{aligned} f &= \prod_{A \in \Delta} \prod_{\alpha \in A} (x - \alpha) \\ &= \prod_{A \in \Delta} (h - \gamma_A) \\ &= g(h). \end{aligned}$$

The coefficients of h are in K by the assumption of (ii), and the coefficients of g are in K since they are solutions of the nonsingular system (1). \square

If f is irreducible, then we need only check the condition of Theorem 4.2 for one $A \in \Delta$; if it holds for one, then it holds for all, since \mathcal{G} is transitive on Δ . The coefficients of h will be the $c_k^s(A)$, $1 \leq k \leq s-1$. The constant coefficient of h is 0, without loss of generality. The roots of g are $c_s^s(A)$, $A \in \Delta$. The coefficients of g are obtained by solving the triangular linear system (1).

Theorem 4.2 gives an algebraic condition that can be used to test decomposability of any f over any field K , provided one can factor over K and thereby construct the splitting field of f . The complexity of the algorithm depends on the complexity of factoring over K .

5 Decomposition over Finite Fields

The conditions of Theorem 4.2 give an algorithm for decomposing any polynomial over any field, provided one can factor polynomials over that field. In this section, we show that for irreducible polynomials over finite fields, we obtain an *NC* algorithm (parallel $\log^{O(1)} mnp$ time on $(mnp)^{O(1)}$ processors), where $K = GF(p^m)$ and n is the degree of f .

Block decompositions of irreducible polynomials over finite fields are particularly easy to compute. If $K = GF(p^m)$ and f is irreducible over K , then $\widehat{K} = K(\alpha)$ where α is any root

of f , and the only automorphisms of $K(\alpha)$ over K are of the form $\beta \mapsto \beta^{p^{km}}$, $0 \leq k \leq n-1$. It follows that there is a *unique* $r \times s$ block decomposition

$$\begin{aligned} \Delta &= \{A_i \mid 0 \leq i \leq r-1\}, \\ A_i &= \{\alpha^{p^{(i+jr)m}} \mid 0 \leq j \leq s-1\}, \quad 0 \leq i \leq r-1 \end{aligned}$$

for f . It thus suffices to construct the conjugates $\alpha^{p^{km}}$, $0 \leq k \leq n-1$, so that one of the polynomials

$$h - \gamma_i = \prod_{\beta \in A_i} (x - \beta), \quad 0 \leq i \leq r-1$$

can be computed and the condition of Theorem 4.2 checked. These conjugates are represented by the polynomials

$$x^{p^{km}} \bmod f, \quad 0 \leq k \leq n-1,$$

which can be computed in NC (parallel time $\log^{O(1)} mnp$ on $(mnp)^{O(1)}$ processors) by the method of Fich and Tompa [7].

Theorem 5.1 *An irreducible polynomial f over a finite field K can be tested for the existence of a nontrivial decomposition $g \circ h$ in NC . If such a decomposition exists, the coefficients of g and h can be computed in NC .*

Proof. First we compute representations of the roots of f in NC , as described above. For any r, s such that $n = rs$, these roots can be arranged in the unique $r \times s$ block decomposition Δ , and for some $A \in \Delta$, the product

$$\prod_{\alpha \in A} (x - \alpha)$$

computed, and tested for the condition of Theorem 4.2. This computation also gives the coefficients of h , if it exists. Finally, the triangular linear system (1) can be solved in NC [4] to obtain the coefficients of g . \square

Example 5.2 There are exactly two primitive irreducible polynomials of degree 4 over $GF(2)$, namely

$$\begin{aligned} p_1 &= x^4 + x + 1 \\ p_2 &= x^4 + x^3 + 1. \end{aligned}$$

We will show that p_1 has a nontrivial decomposition and p_2 does not. First consider p_1 . The unique 2×2 block decomposition is

$$\Delta = \{\{\alpha, \alpha^4\}, \{\alpha^2, \alpha^8\}\},$$

where α is a root of p_1 . This reduces modulo p_1 to

$$\{\{\alpha, \alpha + 1\}, \{\alpha^2, \alpha^2 + 1\}\}.$$

Computing the two polynomials whose roots are the two multisets of Δ , we obtain

$$\begin{aligned}(x - \alpha)(x - (\alpha + 1)) &= x^2 + x + (\alpha^2 + \alpha) \\ (x - \alpha^2)(x - (\alpha^2 + 1)) &= x^2 + x + (\alpha^2 + \alpha + 1),\end{aligned}$$

thus the requirements of Theorem 4.2 are met, and we may take

$$h = x^2 + x.$$

The coefficients of g are obtained by solving the nonsingular triangular system (1). In this case, we have

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$$

with solution $(1, 1, 1)$, giving

$$g = x^2 + x + 1.$$

Computing $g \circ h$, we find

$$g \circ h = (x^2 + x)^2 + (x^2 + x) + 1 = p_1.$$

For p_2 , the unique 2×2 block decomposition is

$$\Delta = \{\{\beta, \beta^4\}, \{\beta^2, \beta^8\}\} = \{\{\beta, \beta^3 + 1\}, \{\beta^2, \beta^3 + \beta^2 + \beta\}\}$$

where β is a root of p_2 . Computing the two polynomials whose roots are the two multisets of Δ , we obtain

$$\begin{aligned}(x - \beta)(x - (\beta^3 + 1)) &= x^2 + (\beta^3 + \beta + 1)x + (\beta^3 + \beta + 1) \\ (x - \beta^2)(x - (\beta^3 + \beta^2 + \beta)) &= x^2 + (\beta^3 + \beta)x + (\beta^3 + \beta),\end{aligned}$$

neither of whose second coefficients are in $GF(2)$, therefore h does not exist, and p_2 has no nontrivial decomposition. \square

6 Decomposition over Arbitrary Fields

Theorem 4.2 also gives the following decomposition result for arbitrary fields:

Theorem 6.1 *A complete decomposition of an irreducible polynomial f over any field K admitting a polynomial-time polynomial factorization algorithm can be constructed in time $O(n^{\log n})$.*

Proof. First we dispose of the problem of inseparability. If K is a field of characteristic p , and if $f \in K[x]$ is inseparable and irreducible, then $f = g(x^{p^k})$ for some irreducible and separable g , thus we have an immediate nontrivial decomposition. Hence we will ignore this case and assume that f is separable.

In Landau and Miller [11], it was shown how to determine minimal blocks of the roots of f in polynomial time. A similar procedure can be used to determine *all* blocks.

By Theorem 4.2, any nontrivial decomposition $f = g \circ h$ corresponds to an $r \times s$ block decomposition Δ , where $s = \deg h$, such that $c_i^s(A) \in K$, $0 \leq i \leq s - 1$, for one (and hence all) $A \in \Delta$. We first use the algorithm of [11] to compute all minimal blocks and test this condition. For each block A which satisfies this condition, the coefficients of h are obtained immediately from the $c_i^s(A)$, and the coefficients of g are obtained by solving the triangular linear system (1). The decomposition algorithm is then applied to g . Note that g is irreducible and h is indecomposable; if g were reducible, then f would have been, and if h were decomposable, then the block A would not have been minimal.

If no minimal block decomposition satisfying Theorem 4.2 is found, then the algorithm is repeated using blocks whose only proper subblock is minimal. If any of these satisfy Theorem 4.2, then the appropriate g and h are constructed, and the decomposition algorithm is then performed on g . If not, the next level of minimal blocks is examined. Each time a decomposition is found, the polynomial g is irreducible, and the polynomial h is indecomposable. The algorithm halts when the only block left to examine is Ω or when g is indecomposable.

If there is a polynomial-time algorithm for factoring over the base field K , then the algorithm of [11] for finding a block is polynomial in the size of f . The difficulty is that there are potentially $n^{\log n}$ blocks, and we must search through all of them to find a decomposition. For each block, the condition of Theorem 4.2 can be tested and the coefficients of g and h obtained in polynomial time. \square

Finally we observe that over \mathcal{Q} , most irreducible polynomials are indecomposable. This is because Theorem 4.2 implies that an irreducible polynomial f over \mathcal{Q} can be decomposed *only if* the Galois group of the splitting field of f over \mathcal{Q} acts imprimitively on the roots of f . But Gallagher [9] has shown that almost all irreducible polynomials over \mathcal{Q} have Galois group S_n , which acts primitively.

7 Recent work

Some improvements and extensions have appeared since the results of this paper were first reported. M. Atkinson [correspondence] and independently J. von zur Gathen (see [15]) have observed that coefficients of h in Algorithm 3.4 can be obtained in time $O(n \log n)$, using a technique related to Hensel lifting. M. Dickerson [5] has obtained partial results for multivariate polynomials.

Acknowledgments

We are grateful to Bruce Char and Arash Baratloo for their implementation of Algorithms 3.4 and 3.5, and to Mike Atkinson for insightful comments.

References

- [1] V. S. Alagar and M. Thanh. Fast polynomial decomposition algorithms. In *Proc. EURO-CAL85*, pages 150–153. Springer-Verlag Lect. Notes in Comput. Sci. 204, 1985.
- [2] D. R. Barton and R. E. Zippel. Polynomial decomposition. In *Proceedings SYMSAC '76*, pages 356–358, 1976.
- [3] D. R. Barton and R. E. Zippel. Polynomial decomposition algorithms. *J. Symb. Comp.*, 1:159–168, 1985.
- [4] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5:618–623, 1976.
- [5] Matthew Dickerson. Polynomial decomposition algorithms for multivariate polynomials. Technical Report TR87-826, Comput. Sci., Cornell Univ., April 1987.
- [6] H. T. Engstrom. Polynomial substitutions. *Amer. J. Math.*, 63:249–255, 1941.
- [7] Faith E. Fich and Martin Tompa. The parallel complexity of exponentiating polynomials over finite fields. In *Proc. 17th Symp. Theory of Comput.*, pages 38–47. ACM, May 1985.
- [8] M. D. Fried and R. E. MacRae. On the invariance of chains of fields. *Ill. J. Math.*, 13:165–171, 1969.
- [9] P. X. Gallagher. Analytic number theory. In *Proc. Symp. in Pure Math.*, pages 91–102. Amer. Math. Soc., 1972.
- [10] G. H. Hardy and S. Ramanujan. *Proc. London Math. Soc.*, 2(17):75–115, 1918.
- [11] Susan Landau and Gary Miller. Solvability by radicals is in polynomial time. *J. Comput. Syst. Sci.*, 30:179–208, 1985.
- [12] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University, 1986.
- [13] J. F. Ritt. Prime and composite polynomials. *Trans. Amer. Math. Soc.*, 23:51–66, 1922.
- [14] L. G. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Comput.*, 12(4):641–644, 1983.
- [15] Joachim von zur Gathen, Dexter Kozen, and Susan Landau. Functional decomposition of polynomials. In *Proc. 28th Symp. Found. Comput. Sci.*, pages 127–131. IEEE, November 1987.