

A polynomial time algorithm for computing all minimal decompositions of a polynomial

Raoul Blankertz
rblankertz@uni-bonn.de

Abstract

The composition of two polynomials $g(h) = g \circ h$ is a polynomial. For a given polynomial f we are interested in finding a functional decomposition $f = g \circ h$. In this paper an algorithm is described, which computes all minimal decompositions in polynomial time. In contrast to many previous decomposition algorithms this algorithm works without restrictions on the degree of the polynomial and the characteristic of the ground field. The algorithm can iteratively applied compute all decompositions. It is based on ideas of Landau & Miller (1985) and Zippel (1991). Additionally, an upper bound on the number of minimal decompositions is given.

1. Introduction

A decomposition of an univariate polynomial f over a field is a pair (g, h) of polynomials such that $f = g \circ h$ and $\deg g, \deg h \geq 2$. The computational problem to compute a decomposition of a given polynomial is much investigated. A major challenge in this task is the case where the characteristic of the field divides the degree of the input polynomial. Barton & Zippel (1985) give an exponential time algorithm, which works independently from the characteristic, while the algorithm of Kozen & Landau (1989) runs in polynomial time, but does not work unconditionally in positive characteristic.

We discuss an algorithm that computes all minimal decompositions of a polynomial in polynomial time, where a decomposition (g, h) is called minimal if h has no decomposition. In contrast to many previous decomposition algorithms—see also von zur Gathen (1990a, Theorem 2.4)—the algorithm described here works over fields of positive characteristic as well. We estimate its runtime for finite fields. This gives an explicit polynomial bound on the complexity of the problem of computing a decomposition of an univariate polynomial over a finite field. By applying the algorithm iteratively, we can compute all decompositions in quasi-polynomial time. This is the best we can expect, since there are polynomials with quasi-polynomially many decompositions; see Giesbrecht (1988, Theorem 3.9).

The main idea for the algorithm is to relate decompositions of $f \in F[x]$ to certain partitions of the set of roots of $f - t$, where t is transcendental over $F[x]$, and to find a way to efficiently compute these partitions. To specify this idea, let (g, h) be a decomposition of f . For each root λ of $g - t$, the roots of $h - \lambda$ form a subset of the roots of $f - t$. Furthermore, two different roots of $g - t$ yield two disjoint subsets. In this way one can partition the set of roots of $f - t$ with respect to a decomposition of f . These partitions are related to blocks of imprimitivity of a certain permutation group. This relation will be made precise in Section 2.

In Section 3 two sharp upper bounds on the number of minimal decompositions of a polynomial are given. After introducing briefly the necessary results from Landau & Miller (1985) in Section 4, we describe an algorithm for polynomial decomposition in Section 5 which is based on ideas of Zippel (1991).

2. Minimal decompositions and blocks of imprimitivity

In this section we first give some general definitions and properties of decompositions. For a polynomial $f \in F[x]$, we define a related polynomial $\varphi = f - t$, where t is transcendental over $F[x]$. Then we establish the connection between the decompositions of f and certain subsets of the set of roots of φ . We do this in three steps. First, we consider the field extension by a root of φ and relate the decompositions to the intermediate fields of this field extension. Second, we establish the notion of blocks of imprimitivity of a permutation group and state their relation to certain subgroups of the permutation group. Finally, we consider the Galois group G of φ as permutation group on the set of roots of φ and connect the decompositions of f to the blocks of imprimitivity of G via Galois theory.

Let F be an arbitrary field. In the runtime considerations of the algorithm in Section 5 we restrict F to a field in which one can compute efficiently. One can think of F being a finite field, which is the most interesting case.

DEFINITION 2.1. *A polynomial $f \in F[x]$ is decomposable if there are g and $h \in F[x]$, both of degrees at least two, such that $f = g \circ h$. The pair (g, h) is called a decomposition of f . In a decomposition (g, h) , we call g the left component and h the right component. A polynomial is indecomposable if it is not decomposable. We call a polynomial original if its graph passes through the origin—or, equivalently, its constant term is zero. A polynomial is monic original if it is monic and original. A decomposition (g, h) is called minimal if h is monic original and indecomposable.*

In a decomposition (g, h) of f , g is uniquely determined by f and h , since the ring homomorphism $F[x] \rightarrow F[x]$ with $x \mapsto h$ is injective. Furthermore, g is easy to compute by the generalized Taylor expansion; see von zur Gathen (1990a, Section 2). Let a be the leading coefficient of h and c be its constant term. For $\ell = ax - c$ and $\ell^* = a^{-1}(x + c)$, we have $f = g \circ h = g \circ \ell^* \circ \ell \circ h$ and hence $(g \circ \ell^*, \ell \circ h)$ is a decomposition of f where $\ell \circ h$ is monic original. Thus we may restrict ourselves without loss of generality to decompositions (g, h) where h is monic original.

Functional decomposition is related to intermediate fields of certain field extensions in the following way. Let t be transcendental over $F[x]$ and $F(t)$ be the rational function field in t over F . Then for a given polynomial $f \in F[x]$ let φ be the irreducible polynomial $f - t \in F(t)[x]$. If we assume that the derivative f' of f is not zero, then the derivative of φ with respect to x is not zero and thus φ is separable. In this case, for a root α of φ in an algebraic closure of $F(t)$, the field $F(t)[\alpha] = F(\alpha)$ is a separable field extension of $F(t)$.

If the characteristic of F is $p > 0$ and $f' = 0$, then there exists a polynomial \tilde{f} and a natural number r such that $f = \tilde{f}(x^{p^r})$ and $\tilde{f}' \neq 0$. If F is perfect—for instance if F is finite—then the Frobenius endomorphism $x \mapsto x^p$ is an automorphism of F . In this case, by knowing all decompositions of \tilde{f} one knows all decompositions of f ; see Giesbrecht (1988, Section 4.6). In general the Frobenius endomorphism is not an automorphism, for example on function fields. From now on we assume that $f' \neq 0$. This assumption excludes some cases in general, but we lose no generality if F is perfect.

Now the following fact states a correspondence between decompositions of f and intermediate fields of $F(\alpha) \mid F(t)$. Let $\mathcal{R} = \{h \in F[x] : h \text{ is monic original and } f = g \circ h \text{ for some } g \in F[x]\}$ be the set of right components of decompositions of f and let \mathcal{M} be the set of intermediate fields between $F(\alpha)$ and $F(t)$.

FACT 2.2 (Fried & MacRae 1969, Proposition 3.4). *Let $f \in F[x]$ with $f' \neq 0$. Then the map $\mathcal{R} \rightarrow \mathcal{M}$ with $h \mapsto F(h(\alpha))$ is bijective.*

The minimal polynomial of α over $F(h(\alpha))$ is $h(x) - h(\alpha)$. Thus $[F(\alpha) : F(h(\alpha))] = \deg(h)$. Furthermore, if $h = u \circ h^*$ for some $u \in F[x]$, then $F(h(\alpha)) \subseteq F(h^*(\alpha))$. Thus, if we take $h^* \leq h$ to mean that $h = u \circ h^*$ for some $u \in F[x]$, then the bijection in Fact 2.2 is an order-reversing bijection of partially ordered sets. Hence, \mathcal{R} equipped with \leq is a lattice.

DEFINITION 2.3. *We call \mathcal{R} the lattice of decompositions of f .*

We introduce the notion of blocks of imprimitivity and its relation to decompositions. For this purpose consider a finite permutation group G on a finite set Z , that is, G is a subgroup of the symmetric group on Z .

DEFINITION 2.4. *A block of G is a subset $B \subseteq Z$ such that for all $\sigma \in G$ the set $\sigma(B) \cap B$ is empty or equals B .*

Equivalently, B is a block of G if for all $\sigma \in G$ the sets B and $\sigma(B)$ are disjoint or equal. If B is a block, then any $\sigma(B)$ is a block. If G is transitive and $B \neq \emptyset$, then $\{\sigma(B)\}_{\sigma \in G}$ is a partition of Z .

LEMMA 2.5. *If B and C are blocks of G , then $B \cap C$ is a block of G .*

PROOF. Let $\sigma \in G$. Then $\sigma(B \cap C) \cap (B \cap C) = (\sigma B \cap B) \cap (\sigma C \cap C)$ is empty if and only if $\sigma B \cap B$ or $\sigma C \cap C$ is empty. If both are nonempty, then $\sigma(B \cap C) \cap (B \cap C) = B \cap C$, since B and C are blocks. \square

DEFINITION 2.6. *The blocks \emptyset , Z , and $\{\gamma\}$, for $\gamma \in Z$, are called trivial blocks. A nontrivial block is called block of imprimitivity. A permutation group G on Z is called primitive if there are only trivial blocks. It is called imprimitive, otherwise.*

For a subgroup $U \subseteq G$ and $\gamma \in Z$, the orbit of γ under U is the subset $U(\gamma) = \{\sigma(\gamma) : \sigma \in U\} \subseteq Z$. For a subset $S \subseteq Z$, the (setwise) stabilizer of S is the subgroup $G_S = \{\sigma \in G : \sigma(S) = S\}$. We write G_γ for $G_{\{\gamma\}}$.

The following theorem is essential for the link between the decomposition of polynomials and blocks of imprimitivity.

FACT 2.7 (Wielandt 1964, Theorem 7.5). *Let G be a finite transitive permutation group on a finite set Z and let $\gamma \in Z$. Then the mapping $U \mapsto U(\gamma)$ is an isomorphism from the lattice of subgroups between G_γ and G to the lattice of blocks of G containing γ . The inverse mapping is $B \mapsto G_B$.*

We fix the following notation. Let $f \in F[x]$ be of degree $n > 1$ with $f' \neq 0$. As before, we define $\varphi = f - t \in F(t)[x]$. Let L be a splitting field of φ over $F(t)$, G be its Galois group, and Z the set of roots of φ in L . Then G acts transitively on Z and we consider G as a permutation group on Z . Furthermore, we fix $\alpha \in Z$ that on the one hand plays the role of a root of φ as in the beginning of the section and on the other hand defines the lattice of blocks of G containing α as γ does in Fact 2.7.

COROLLARY 2.8. *Let $f \in F[x]$ be of degree $n > 1$ with $f' \neq 0$ and \mathcal{R} be the lattice of decompositions of f .*

- (i) *Then \mathcal{R} and the lattice of blocks of G containing α are isomorphic.*
- (ii) *Let $h \in \mathcal{R}$ and B be the block corresponding to h . Then $\deg(h) = |B|$.*

PROOF. The lattice of decompositions of f is isomorphic to the lattice of intermediate fields of $F(\alpha) \mid F(t)$; see Fact 2.2. This in turn is by Galois theory isomorphic to the lattice of subgroups between G_α and G . Thus, by Fact 2.7, there is an isomorphism between the lattice of decompositions of f and the lattice of blocks containing α .

Let U be the subgroup corresponding to h and B be the corresponding block, that is, $F(h(\alpha)) = L^U$ and $U(\alpha) = B$, where L^U is the subfield of L that is fixed by U . Then $\deg(h) = [F(\alpha) : F(h(\alpha))] = [L^{G_\alpha} : L^U] = (U : G_\alpha) = |U(\alpha)| = |B|$. \square

DEFINITION 2.9. *We call a block B of G minimal if it contains α and all blocks properly contained in B are trivial.*

The minimal blocks of G correspond to the minimal decompositions of f , by Corollary 2.8(i).

3. An upper bound

We deduce two sharp upper bounds on the number of minimal decompositions of a polynomial. These bounds coincide partly with results in von zur Gathen, Giesbrecht & Ziegler (2010).

The intersection of two distinct minimal blocks is a block, by Lemma 2.5, and therefore trivial. Hence the minimal blocks minus $\{\alpha\}$ are distinct sets in $Z \setminus \{\alpha\}$. Therefore, the sum of the cardinality of all minimal blocks minus $\{\alpha\}$ is less than or equal to $n - 1$. Since the cardinality of a block equals the degree of the right component of the corresponding decomposition, we get the following result.

COROLLARY 3.1. *Let f be a decomposable polynomial of degree n with $f' \neq 0$.*

- (i) *Let d divide n . Then there are at most $(n - 1)/(d - 1)$ minimal decompositions (g, h) of f with $\deg(h) = d$.*
- (ii) *Let q be the smallest prime divisor of n . Then there are at most $(n - 1)/(q - 1)$ minimal decompositions of f .*

EXAMPLE 3.2. Let p be the characteristic of F and let f be a separable additive polynomial of degree p^r with $r \geq 2$, that is, f is of the form $\sum_{i=0}^r a_i x^{p^i}$ with $a_0 \neq 0$, see Lidl & Niederreiter (1997, Chapter 4.3). Furthermore, assume that f splits completely over F . Then the roots of f form a group $G \subseteq F$ which is isomorphic to $(\mathbb{Z}/p\mathbb{Z})^r$. If α is a root of $f - t$, so is $\alpha + a$ for all roots a of f . Thus $F(\alpha) | F(t)$ is Galois and its Galois group is isomorphic to G . But G has exactly $(p^r - 1)/(p - 1)$ subgroups of order p . Thus f has exactly $(p^r - 1)/(p - 1)$ minimal decompositions. This shows that both bounds are sharp. \diamond

The numbers $(n - 1)/(d - 1)$ and $(n - 1)/(q - 1)$ in Corollary 3.1 are not necessarily integers. Even though the bounds are sharp, in many cases there are much fewer minimal decomposition than stated in Corollary 3.1. For more details on the number of decomposition of polynomials see von zur Gathen (2009), von zur Gathen, Giesbrecht & Ziegler (2010), and Blankertz, von zur Gathen & Ziegler (2012).

4. Finding minimal blocks

In this section we will discuss a method to compute minimal blocks of the Galois group G . The astonishing result here is that one can compute these blocks without computing the Galois group itself. This result and all intermediate results were introduced in Landau & Miller (1985) for the ground field \mathbb{Q} . In our case we have the ground field $F(t)$, but the proofs are essentially the same. Therefore we will discuss this method here only briefly. A detailed discussion can be found in Landau & Miller (1985), Landau (1993), Zippel (1996), and Blankertz (2011).

We factor φ over $F(\alpha)$ into monic irreducible factors ψ_i such that

$$(4.1) \quad \varphi = \prod_{i=1}^s (x - \alpha_i) \cdot \psi_{s+1} \cdot \dots \cdot \psi_r$$

with $\alpha = \alpha_1$, $\alpha_i \in F(\alpha)$, and $\psi_i = x - \alpha_i$ for $1 \leq i \leq s$, and $\deg(\psi_i) \geq 2$ for $s < i \leq r$. Since $\alpha_i \in F(\alpha)$ for $1 \leq i \leq s$, there are rational functions ℓ_i such that $\alpha_i = \ell_i(\alpha)$. Since α is transcendental over F , from the equation $f(\alpha) = t = f(\ell_i(\alpha))$ follows that ℓ_i is a linear polynomial.

Then $H = (\{\ell_i: 1 \leq i \leq s\}, \circ)$ is a group and $H(\alpha) = \{\alpha_i: 1 \leq i \leq s\}$ is a block of G . Let U_α be the subgroup corresponding to $H(\alpha)$. Then $H \cong U_\alpha/G_\alpha$. We first show that we can compute all blocks that are contained in $H(\alpha)$.

Suppose $s > 1$. Then the induced action of U_α on $H(\alpha)$ is determined by the action of H on $H(\alpha)$, since G_α acts trivial on $H(\alpha)$. If there are minimal blocks of U_α containing α , then one can find all of them in polynomial time in the size of H by an algorithm of Atkinson (1975). These blocks are also minimal blocks of G . If there are no nontrivial blocks of U_α , then $H(\alpha)$ is a minimal block of G . Thus we have the following lemma.

LEMMA 4.2. *All minimal blocks that are contained in $H(\alpha)$ can be computed in polynomial time in n by the algorithm of Atkinson (1975).*

If $s = 1$, then $B_\alpha = \{\alpha\}$ is trivial. Thus the method above cannot be applied to find minimal blocks of G . Even more generally, if $s < n$, there may exist minimal blocks containing α but being not contained in $H(\alpha)$. For such a block Λ , we have $\Lambda \cap H(\alpha) = \{\alpha\}$, by Lemma 2.5 and the minimality of Λ . In the following we show how to compute such blocks.

THEOREM 4.3. *Let Λ be a minimal block of G with $\alpha \in \Lambda$ and $\Lambda \cap H(\alpha) = \{\alpha\}$. Then for all $\beta \in \Lambda$ distinct from α the orbit $\langle G_\alpha, G_\beta \rangle(\alpha)$ equals Λ .*

Fix $k > s$ and $\beta \in Z$ such that β is a root of ψ_k . Then $\beta \notin B_\alpha$ and thus $\langle G_\alpha, G_\beta \rangle(\alpha)$ is a block, which is minimal if there is a minimal block containing α and β . Let $\sigma \in G$ such that $\sigma(\alpha) = \beta$ and set $\psi_i^* = \sigma(\psi_i)$ for all $1 \leq i \leq r$. Then the polynomials $\psi_i^* \in F(\beta)[x]$ are the polynomials ψ_i with β substituted for α and the irreducible factors of φ over $F(\beta)$ are precisely the polynomials ψ_i^* .

THEOREM 4.4. *Consider the bipartite graph Γ_β with the set of vertices consisting of ψ_i and ψ_i^* for $1 \leq i \leq r$ and with an undirected edge between ψ_i and ψ_j^* if $\gcd(\psi_i, \psi_j^*) \neq 1$. Let C_β be the set of roots of those ψ_i that are connected to ψ_1 . Then $\langle G_\alpha, G_\beta \rangle(\alpha) = C_\beta$.*

We can compute C_β by performing at most r^2 gcd computations in $F(\alpha, \beta)[x]$. In particular, we do not need to compute G . Thus by Lemma 4.2 and the previous theorems, we can compute all minimal blocks of G in polynomial time.

5. The algorithm

In this section we first state a lemma that tells us how to compute the corresponding decomposition from a given block. Then we see two examples to illustrate the involved computations and the connection between the Galois group, its blocks, and the decompositions. Finally, we describe the algorithm and prove its properties.

LEMMA 5.1. *Let B be a block and h be the right component of a decomposition of f corresponding to B . Then $h(x) - h(\alpha) = \prod_{\gamma \in B} (x - \gamma)$.*

PROOF. The block B corresponds to the intermediate field L^{G_B} and by Fact 2.2 there is a decomposition of f with right component h such that $L^{G_B} = F(h(\alpha))$. For $\lambda = h(\alpha)$, the minimal polynomial of α over $F(\lambda)$ is $h - \lambda$. Let $u = \prod_{\gamma \in B} (x - \gamma)$. For all $\sigma \in G_B$, we have $\sigma(B) = B$ and therefore $\sigma(u) = \prod_{\gamma \in B} (x - \sigma\gamma) = \prod_{\gamma \in \sigma(B)} (x - \gamma) = u$. Since $F(\lambda) = L^{G_B}$, this proves that u is in $F(\lambda)[x]$. Since $\alpha \in B$, we find $u(\alpha) = 0$ and the gcd of u and $h - \lambda$ in $F(\lambda)[x]$ is not constant. Since $h - \lambda$ is irreducible over $F(\lambda)[x]$ and since both polynomials are monic and have the same degree, we find $u = h - \lambda$. \square

The constant term of $h(x) - h(\alpha)$ is $\prod_{\gamma \in B} (-\gamma)$. Since h is monic original, we get $h = \prod_{\gamma \in B} (x - \gamma) - \prod_{\gamma \in B} (-\gamma)$, as explicit formula.

EXAMPLE 5.2. Let p be an odd prime and F be a finite field of characteristic p . Let $f = x^2 \circ (x^p - x)$, a be an element of the prime field \mathbb{F}_p of F and ζ be either 1 or -1 . Then $f(\zeta x + a) = (\zeta^p x^p + a^p - \zeta x - a)^2 = f(x)$. Thus, for a root α of $f - t$ also $\zeta\alpha + a$ is a root of f . Hence we have $2p$ roots of $f - t$ in $F(\alpha)$ and therefore $F(\alpha) \mid F(t)$ is Galois. Its Galois group is isomorphic to $\{(\zeta x + a) \mid \zeta \in \{-1, 1\}, a \in \mathbb{F}_p\} \cong \mathbb{F}_p \rtimes \mathbb{Z}/2\mathbb{Z} \cong D_{2p}$. The dihedral group D_{2p} has one subgroup of order p and p subgroups of order two. Hence f has $p+1$ decompositions. A block with two elements is of the form $\{\alpha, -\alpha + a\}$. Then $h(x) - h(\alpha) = (x - \alpha)(x - (-\alpha + a)) = x^2 - ax - (\alpha^2 - a\alpha)$ and we have found the right component of a decomposition of f , namely $h = x^2 - ax$. \diamond

EXAMPLE 5.3. Let $p = 3$ and $f = x^9 - x$ over \mathbb{F}_3 . Let α be a root of $f - t$. One checks that

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

is the factorization of $f - t$ into irreducible polynomials in $\mathbb{F}_3(\alpha)[x]$. As stated above, $\{\alpha, \alpha - 1, \alpha + 1\}$ forms a block. We compute $h(x) - h(\alpha) = (x - \alpha)(x - \alpha + 1)(x - \alpha - 1) = x^3 - x - (\alpha^3 - \alpha)$ and we find that $h = x^3 - x$ is the right component of a decomposition of f . The corresponding left component is $g = x^3 + x$. We have found a decomposition according to Lemma 4.2. Next we compute a decomposition according to Theorem 4.3. Let

$$\begin{aligned} \psi_1 &= x^2 + \alpha x + \alpha^2 + 1, \\ \psi_2 &= x^2 + \alpha x + x + \alpha^2 - \alpha - 1, \\ \psi_3 &= x^2 + \alpha x - x + \alpha^2 + \alpha - 1. \end{aligned}$$

With this notation we have $f - t = (x - \alpha)(x - \alpha + 1)(x - \alpha - 1)\psi_1\psi_2\psi_3$. Now let ζ be a root of $x^2 + x - 1$ in \mathbb{F}_9 . Then we find

$$\begin{aligned} \psi_1 &= (x - (\alpha + \zeta + 1))(x - (\alpha - \zeta - 1)), \\ \psi_2 &= (x - (\alpha + \zeta - 1))(x - (\alpha - \zeta)), \\ \psi_3 &= (x - (\alpha + \zeta))(x - (\alpha - \zeta + 1)). \end{aligned}$$

Let $\beta_1 = \alpha + \zeta + 1$. Then we have that ψ_1 with β_1 substituted for α is $\psi_1^* = (x - \alpha)(x - (\alpha - \zeta - 1))$ and thus $C_{\beta_1} = \{\alpha, \alpha + \zeta + 1, \alpha - \zeta - 1\}$ is a minimal block. We find $h^*(x) - h^*(\alpha) = (x - \alpha)\psi_1 = x^3 + x - (\alpha^3 + \alpha)$ and hence another right component $h^* = x^3 + x$. Then the corresponding left component is $x^3 - x$.

Next, for $\beta_2 = \alpha + \zeta - 1$ we get $\psi_3^* = (x - \alpha)(x - (\alpha - \zeta - 1))$ is ψ_3 with β_2 substituted for α . Thus $C_{\beta_2} = C_{\beta_1}$. In the same way $\alpha - \zeta + 1$ does not yield any further block. Therefore, f has all in all exactly two decompositions over \mathbb{F}_3 .

Note that going to the extension \mathbb{F}_9 of \mathbb{F}_3 unveils more structure. Actually, f has four decompositions over \mathbb{F}_9 as we have seen in Example 3.2. \diamond

In the following we give an algorithm that computes all minimal decompositions of a polynomial, whose derivative does not vanish. It is based on Landau & Miller (1985) and on Zippel (1991).

Algorithm 5.4 calls a subroutine $\text{Atkinson}(G, Z, \alpha)$ which returns a list of all minimal blocks of G acting on Z which contain α . If G is primitive, this list consists of Z only.

THEOREM 5.5. *Algorithm 5.4 correctly computes all minimal decompositions of f .*

PROOF. Let (g, h) be a minimal decomposition and Λ be the corresponding block. Then either $\Lambda \subseteq B_\alpha$ or $\Lambda \cap B_\alpha = \{\alpha\}$, by Lemma 2.5 and the minimality of Λ . In the first case, Λ is computed in step 5; see Lemma 4.2. Then h is recovered from Λ in step 7, by Lemma 5.1. In the second case, let $\beta \in \Lambda \setminus \{\alpha\}$ and k such that $\psi_k(\beta) = 0$. By Theorem 4.3 and Theorem 4.4, we have $\Lambda = C_\beta = \{\gamma : \exists i \in I_k : \psi_i(\gamma) = 0\}$, where C_β is as in Theorem 4.4 and I_k is computed in step 13. Then in step 15, we have $\prod_{i \in I_k} \psi_i = \prod_{\gamma \in \Lambda} (x - \gamma) = h(x) - h(\alpha)$, from which we can recover h . \square

ALGORITHM 5.4. Computing minimal decompositions.

Input: A monic polynomial $f \in F[x]$ of degree n with $f' \neq 0$.

Output: A list of decompositions (g, h) of f . This list is empty if f is indecomposable.

1. Set $List = \{\}$ and let $F(\alpha)$ be the rational function field in α .
2. Factor $f(x) - f(\alpha) \in F(\alpha)[x]$ into $\prod_{i=1}^s (x - \alpha_i) \cdot \psi_{s+1} \cdot \dots \cdot \psi_r$ as in (4.1).
3. If $s > 1$ then
 4. Set $H(\alpha) = \{\alpha_i : 1 \leq i \leq s\}$ and $H = (\{\ell_i : 1 \leq i \leq s\}, \circ)$ where $\alpha_i = \ell_i(\alpha)$.
 5. Set $AtkinsonBlocks = Atkinson(H, H(\alpha), \alpha)$.
 6. For $\Lambda \in AtkinsonBlocks$ with $|\Lambda| < n$ do 7–9
 7. Compute $h(x) = \prod_{\gamma \in \Lambda} (x - \gamma) - \prod_{\gamma \in \Lambda} (-\gamma)$.
 8. Compute g such that $f = g \circ h$.
 9. Attach (g, h) to $List$.
10. For $k \in \{s + 1, \dots, r\}$ do 11–17
11. Let β be a root of ψ_k and let ψ_i^* be ψ_i with β substituted for α for all $1 \leq i \leq r$.
12. Compute the graph Γ_β as in Theorem 4.4.
13. Compute $I_k = \{i : \psi_i \text{ is connected to } \psi_1 \text{ in } \Gamma_\beta\}$.
14. If $I_k \neq \{1, \dots, r\}$ then
 15. Compute $h(x) - h(\alpha) = \prod_{i \in I_k} \psi_i$, where $h(x) \in F[x]$ is monic original.
 16. Compute g such that $f = g \circ h$.
 17. Attach (g, h) to $List$.
18. Return $List$.

In Algorithm 5.4 we iterate though all ψ_i with $s < i \leq r$ and—as seen in Example 5.3—we might hit a block more than once. Furthermore, for a root β of ψ_i , C_β is a block even if there is no minimal block containing α and β . Then either C_β is minimal and $\beta \notin C_\beta$ or C_β is not minimal and contains minimal blocks. Hence the output list of Algorithm 5.4 does not necessarily consist of distinct decompositions and may contain decompositions that are not minimal. An algorithm that computes *only* minimal decompositions and outputs each decomposition only once should keep track of this. This is easily done by checking in each iteration if the new computed I_k is contained in or contains a previously computed I_j for all $s < j < k$.

For the runtime consideration, let F be a field over which one can factor bivariate polynomials in polynomial time. Then F is in particular computable. For instance, consider F being a finite field. The gcd computation in step 12 can be done by computing the resultant which can be done in polynomial time. Thus the algorithm runs in polynomial time. Zippel (1991) proposes an algorithm to compute decompositions of rational functions. This algorithm runs in polynomial time and can be used to compute decompositions of polynomials, as well. Hence we know that the problem of computing a decomposition of a polynomial over F is in polynomial time. Results on explicit complexity bounds were not given. The following runtime estimation for finite fields gives such an explicit complexity bound.

THEOREM 5.6. *Let F be a finite field and n be the degree of the input polynomial f . Denote by $CP(n)$ the complexity of testing if two polynomials over a function field of degree at most n are coprime. Then Algorithm 5.4 runs in $\mathcal{O}(n^3 CP(n))$.*

PROOF. The factorization in step 2 can be done in $\mathcal{O}^\sim(n^{\omega+1})$ field operations, where $2 \leq \omega \leq 3$ is the matrix multiplication exponent; see Bostan, Lecerf, Salvy, Schost & Wiebelt (2004) and Lecerf (2007). Atkinson's algorithm takes $\mathcal{O}(n^3)$ bit operations; see Atkinson (1975) and Butler (1992, Section 2). Butler (1992, Section 3) improved the runtime of Atkinson's algorithm to $\mathcal{O}(n^2 \log n)$. Let $M(n)$ denote the complexity of multiplying two polynomials over F of degree at most n . Then in step 8 and 16, for each right component h , the appropriate left component g can be computed in $\mathcal{O}(M(n) \log n)$ field operations by the generalized Taylor expansion; see von zur Gathen (1990a, Section 2). Since there are at most s minimal blocks computed by the algorithm of Atkinson, step 8 is called at most s times. Step 16 is called at most $r - s$ times. Thus we get $\mathcal{O}(rM(n) \log n)$ field operations for step 8 and 16.

To compute the graph in step 12 we need at most r^2 gcd computations. We have to compute $r - s$ such graphs. Thus in total we have at most $(r - s)r^2 \leq n^3$ such gcd computations. Actually, we do not need to compute the gcds, but just test if two polynomials are coprime. Since one coprimality test takes at least n field operations, step 12 dominates the runtime in the worst case. \square

REMARK 5.7. (i) *The field arithmetic of $F(\alpha, \beta)$ is quite costly, thus one should use a modular algorithm that tests if two polynomials in $F(\alpha, \beta)[x]$ are coprime. Blankertz (2011) proposes such an algorithm. With fast multiplication it has expected runtime $\mathcal{O}^\sim(n^3 \log(q))$, where q is the size of F , and an error probability of at most $(4n)^{-1}$. If we repeat this coprimality test c times, then we get for all n^3 computations an error probability of at most $n^3(4n)^{-c}$. Then Algorithm 5.4 takes an expected number of $\mathcal{O}^\sim(cn^6 \log(q))$ operations in F .*

(ii) *The bottleneck of Algorithm 5.4 clearly is step 12. But even if one could improve this step, one can never get faster than the factorization in step 2.*

(iii) *If we want to compute the lattice of decompositions of a polynomial, we can apply Algorithm 5.4 iteratively. In each iteration there are at most n minimal decompositions and the iteration depth is in $\mathcal{O}(\log n)$. Thus we get an algorithm with quasi-polynomial runtime. Actually, we can have quasi-polynomially many decompositions, see Giesbrecht (1988, Theorem 3.9), which shows that we cannot get better than quasi-polynomial without “a totally new approach” (von zur Gathen 1990b).*

6. Conclusion

Computing a decomposition of a polynomial over a finite field can be done in polynomial time. Compared with other decomposition algorithms, see for instance von zur Gathen (1990a, Theorem 2.4), Algorithm 5.4 is quite slow. The advantage of this algorithm is that it has not such a strong assumption on the input polynomial. As pointed out in Section 2, the assumption $f' \neq 0$ is easily handled if F is finite.

Iteratively applied, Algorithm 5.4 computes the lattice of decompositions of a polynomial in quasi-polynomial time. Another approach to compute the lattice of decompositions is to use the subfield finding algorithm from van Hoeij, Klüners & Novocin (2011). From a given finite field extension this algorithm computes a set of subfields—so called generating subfields—such that any other subfield is an intersection of some of these generating subfields. One can apply Fact 2.2 to

compute the corresponding decompositions. A detailed runtime analysis for this approach is still outstanding.

7. Acknowledgments

This work was founded in part by the B-IT Foundation and the Land Nordrhein-Westfalen. Parts of this work have also appeared in Blankertz (2011).

References

- M. D. ATKINSON (1975). An Algorithm for Finding the Blocks of a Permutation Group. *Mathematics of Computation* **29**(131), 911–913. ISSN 0378-4754. URL <http://www.jstor.org/stable/2005304>.
- DAVID R. BARTON & RICHARD ZIPPEL (1985). Polynomial Decomposition Algorithms. *Journal of Symbolic Computation* **1**, 159–168.
- RAOUL BLANKERTZ (2011). *Decomposition of Polynomials*. Diplomarbeit, Universität Bonn, Bonn. Modified version available at <http://arxiv.org/abs/1107.0687>.
- RAOUL BLANKERTZ, JOACHIM VON ZUR GATHEN & KONSTANTIN ZIEGLER (2012). Compositions and collisions at degree p^2 . In *Proceedings of the 2012 International Symposium on Symbolic and Algebraic Computation ISSAC2012*, Grenoble, France, 91–98. ACM Press, New York, USA. Full version available at <http://arxiv.org/abs/1202.5810>.
- A. BOSTAN, G. LECERF, B. SALVY, É. SCHOST & B. WIEBELT (2004). Complexity issues in bivariate polynomial factorization. In *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation ISSAC2004*, Santander, Spain, 42–49. ACM Press. ISBN 1-58113-827-X. URL <http://dx.doi.org/10.1145/1005285.1005294>.
- GREG BUTLER (1992). An analysis of Atkinson’s algorithm. *ACM SIGSAM Bulletin* **26**(2), 1–9. ISSN 0163-5824. URL <http://dx.doi.org/10.1145/130933.130935>.
- MICHAEL D. FRIED & R. E. MACRAE (1969). On the invariance of chains of Fields. *Illinois Journal of Mathematics* **13**, 165–171.
- JOACHIM VON ZUR GATHEN (1990a). Functional Decomposition of Polynomials: the Tame Case. *Journal of Symbolic Computation* **9**, 281–299. URL [http://dx.doi.org/10.1016/S0747-7171\(08\)80014-4](http://dx.doi.org/10.1016/S0747-7171(08)80014-4).
- JOACHIM VON ZUR GATHEN (1990b). Functional Decomposition of Polynomials: the Wild Case. *Journal of Symbolic Computation* **10**, 437–452. URL [http://dx.doi.org/10.1016/S0747-7171\(08\)80054-5](http://dx.doi.org/10.1016/S0747-7171(08)80054-5).
- JOACHIM VON ZUR GATHEN (2009). The number of decomposable multivariate polynomials. In *Abstracts of the Ninth International Conference on Finite Fields and their Applications*, 21–22. Claude Shannon Institute, Dublin. URL <http://www.shannoninstitute.ie/fq9/AllFq9Abstracts.pdf>.
- JOACHIM VON ZUR GATHEN, MARK GIESBRECHT & KONSTANTIN ZIEGLER (2010). Composition collisions and projective polynomials. Statement of results. In *Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation ISSAC2010*, Munich, Germany, STEPHEN WATT, editor, 123–130. ACM Press. URL <http://dx.doi.org/10.1145/1837934.1837962>. Preprint available at <http://arxiv.org/abs/1005.1087>.

- MARK WILLIAM GIESBRECHT (1988). Complexity Results on the Functional Decomposition of Polynomials. Technical Report 209/88, University of Toronto, Department of Computer Science, Toronto, Ontario, Canada. Available as <http://arxiv.org/abs/1004.5433>.
- DEXTER KOZEN & SUSAN LANDAU (1989). Polynomial Decomposition Algorithms. *Journal of Symbolic Computation* **7**, 445–456. An earlier version was published as Technical Report 209/88, University of Toronto, Department of Computer Science, Toronto, Ontario, Canada, 1988.
- S. LANDAU & G. L. MILLER (1985). Solvability by Radicals is in Polynomial Time. *Journal of Computer and System Sciences* **30**, 179–208.
- SUSAN LANDAU (1993). Finding maximal subfields. *ACM SIGSAM Bulletin* **27**(3), 4–8. ISSN 0163-5824. URL <http://dx.doi.org/10.1145/170906.170907>.
- GRÉGOIRE LECERF (2007). Improved dense multivariate polynomial factorization algorithms. *Journal of Symbolic Computation* **42**(4), 477–494. ISSN 0747-7171.
- RUDOLF LIDL & HARALD NIEDERREITER (1997). *Finite Fields*. Number 20 in Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, UK, 2nd edition. First published by Addison-Wesley, Reading MA, 1983.
- MARK VAN HOEIJ, JÜRGEN KLÜNERS & ANDREW NOVOCIN (2011). Generating subfields. In *Proceedings of the 2011 International Symposium on Symbolic and Algebraic Computation ISSAC2011*, San Jose CA, 345–352. ACM Press, New York, USA. ISBN 978-1-4503-0675-1. URL <http://dx.doi.org/10.1145/1993886.1993937>.
- HELMUT WIELANDT (1964). *Finite permutation groups*. Academic Press, New York. ISBN 0-127-49656-4. Translated from the German by R. Bercov.
- RICHARD ZIPPEL (1991). Rational Function Decomposition. In *Proceedings of the 1991 International Symposium on Symbolic and Algebraic Computation ISSAC '91*, Bonn, Germany, STEPHEN M. WATT, editor, 1–6. ACM Press, Bonn, Germany. ISBN 0-89791-437-6.
- RICHARD ZIPPEL (1996). Functional Decomposition. online. URL <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.51.3154>. Last visited 31 May 2013.