Journal of Symbolic Computation 75 (2016) 244–268

Contents lists available at ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc

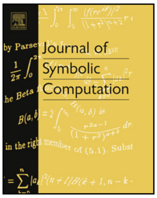

Tame decompositions and collisions

Konstantin Ziegler

B-IT, Universität Bonn, D-53113 Bonn, Germany

ARTICLE INFO

Article history:
Received 9 December 2014
Accepted 10 July 2015
Available online 4 November 2015

MSC:
11T06
12Y05
68W30

Keywords:
Computer algebra
Finite fields
Tame polynomial decomposition
Ritt's Second Theorem
Counting special polynomials
Combinatorics on polynomials

ABSTRACT

A univariate polynomial f over a field is decomposable if $f = g \circ h = g(h)$ with nonlinear polynomials g and h. It is intuitively clear that the decomposable polynomials form a small minority among all polynomials over a finite field $\mathbb{F}_q$. The tame case, where the characteristic of $\mathbb{F}_q$ does not divide $n = \deg f$, is fairly well understood, and we have reasonable bounds on the number of decomposables of degree n. However, it is not known how to determine this number exactly if n has more than two prime factors. There is an obvious inclusion–exclusion formula, but to evaluate its summands, one has to determine, under a suitable normalization, the number of collisions, where essentially different components (g, h) yield the same f. Ritt's Second Theorem classifies all tame collisions of two such pairs.

We present a normal form for tame collisions of any number of decompositions with any number of components and describe exactly the (non)uniqueness of the parameters. This yields the exact number of such collisions over a finite field. We conclude with a fast algorithm for the exact number of decomposable polynomials at degree n over a finite field $\mathbb{F}_q$ of characteristic coprime to n.

1. Introduction

The *composition* of two univariate polynomials $g, h \in F[x]$ over a field F is $f = g \circ h = g(h)$. Then (g, h) is a *decomposition* of f and f is *decomposable* if g and h have degree at least 2. In the 1920s,

E-mail address: zieglerk@uni-bonn.de.
URL: http://zieglerk.net.

http://dx.doi.org/10.1016/j.jsc.2015.11.017
0747-7171/© 2015 Elsevier Ltd. All rights reserved.

Ritt (1922, 1923), Fatou (1921), and Julia (1922) studied structural properties of these decompositions over $\mathbb{C}$, using analytic methods. Particularly important are two theorems by Ritt on the uniqueness, in a suitable sense, of decompositions, the first one for (many) indecomposable components and the second one for two components, as above. Engstrom (1941) and Levi (1942) proved Ritt's theorems over arbitrary fields of characteristic zero using algebraic methods.

The theory was extended to arbitrary characteristic by Fried and MacRae (1969), Dorey and Whaples (1974), Schinzel (1982, 2000), Zannier (1993), and others. We now have applications in cryptography, see Cade (1985, 1987) and Boucher et al. (2010), and signal processing, see Demirtas et al. (2013). In computer algebra, the decomposition method of Barton and Zippel (1985) requires exponential time, see Giesbrecht and May (2007) for a detailed analysis. A fundamental dichotomy is between the *tame case*, where the characteristic p of F does not divide $\deg g$, and the *wild case*, where p divides $\deg g$, see von zur Gathen (1990a, 1990b). (Schinzel (2000, § 1.5) uses *tame* in a different sense.) A breakthrough result of Kozen and Landau (1989) was their polynomial-time algorithm to compute tame decompositions; see also von zur Gathen et al. (1987), Kozen et al. (1996), Gutierrez and Sevilla (2006b), and the survey articles of von zur Gathen (2002) and Gutierrez and Kozen (2003) with further references.

In the wild case, considerably less is known, both mathematically and computationally. Zippel (1991) suggests that the block decompositions of Landau and Miller (1985) for determining subfields of algebraic number fields can be applied to decomposing rational functions even in the wild case. A version of Zippel's algorithm by Blankertz (2014) computes in polynomial time all decompositions of a polynomial that are minimal in a certain sense.

Zannier (2007, 2008, 2009) studies a different but related question, namely decompositions $f = g \circ h$ in $\mathbb{C}[x]$ of *sparse* (or *lacunary*) polynomials f, where the number t of terms is fixed, while the corresponding degrees and coefficients may vary. He shows that the sparsity of f implies the sparsity of h, proving a conjecture by Schinzel, and also gives a parametrization of all such f, g, h in terms of varieties (for the coefficients) and lattices (for the exponents). Fuchs and Pethő (2011) and Fuchs and Zannier (2012) follow up with complete descriptions of sparse decomposable rational functions.

It is intuitively clear that the univariate decomposable polynomials form only a small minority among all univariate polynomials over a field. There is an obvious inclusion–exclusion formula for counting them. The main issue is then to determine, under a suitable normalization, the number of *collisions*, where essentially different components (g, h) yield the same f. The number of decomposable polynomials of degree n is thus the number of all pairs (g, h) with $\deg g \cdot \deg h = n$ reduced by the ambiguities introduced by collisions. An important tool for estimating the number of collisions is Ritt's Second Theorem. The first algebraic versions of Ritt's Second Theorem in positive characteristic p required $p > \deg(g \circ h)$. Zannier (1993) reduced this to the milder and more natural requirement $g' \neq 0$ for all left components g in the collision. His proof works over algebraically closed fields, and Schinzel's (2000) monograph adapts it to finite fields.

The task of counting compositions over a finite field of characteristic p was first considered by Giesbrecht (1988). He showed that the decomposable polynomials form an exponentially small fraction of all univariate polynomials. Von zur Gathen (2014a) presents general approximations to the number of decomposable polynomials. These come with satisfactory (rapidly decreasing) relative error bounds except when p divides $n = \deg f$ exactly twice.

Ritt's First Theorem relates complete decompositions of a given polynomial, where all components are indecomposable. Zieve and Müller (2008) turn this into an applicable method and Medvedev and Scanlon (2014) combine this approach with results from model theory to describe the subvarieties of the k-dimensional affine space that are preserved by a coordinatewise polynomial map. Both works lead to slightly different canonical forms for the complete decomposition of a given polynomial. Zieve and Müller (2008) study the replacement of adjacent indecomposable g, h in a complete decomposition by g^*, h^* with the same composition, but $\deg g = \deg h^* \neq \deg h = \deg g^*$. Motivated by the analogy with the Reidemeister moves of knot theory, this is called a *Ritt move*. The ensuing collision is the theme of Ritt's Second Theorem and von zur Gathen (2014b) presents a normal form with an explicit description of the (non)uniqueness of the parameters under Zannier's assumption $g'(g^*)' \neq 0$.

This work combines the above "normalizations" of Ritt's theorems to classify collisions of two or more decompositions, not necessarily complete and of arbitrary length. We proceed as follows. In