Factoring in Skew-Polynomial Rings

Mark Giesbrecht

Department of Computer Science
University of Toronto
email: mwg@theory.utoronto.ca

Abstract

Efficient algorithms are presented for factoring polynomials in the skew-polynomial ring $\mathsf{K}[x;\sigma]$, a non-commutative generalization of the usual ring of polynomials $\mathsf{K}[x]$, where K is a finite field and $\sigma\colon \mathsf{K} \to \mathsf{K}$ is an automorphism. Applications include fast functional decomposition algorithms for a class of polynomials in $\mathsf{K}[x]$ whose decompositions are "wild" and previously thought to be difficult to compute. Also presented is a fast probabilistic algorithm for finding zero divisors in any finite associative algebra over K.

1 Introduction

A central problem in computer algebra is factoring polynomials in $\mathsf{K}[x]$, where K is a finite field and x is an indeterminate. In this paper we present efficient factorization algorithms in a natural non-commutative generalization of the ring $\mathsf{K}[x]$, the skew-polynomial ring $\mathsf{K}[x;\sigma]$, where $\sigma\colon \mathsf{K} \to \mathsf{K}$ is a field automorphism. $\mathsf{K}[x;\sigma]$ is the ring of all polynomials in $\mathsf{K}[x]$ under the usual component-wise addition, and multiplication defined by $xa = \sigma(a)x$ for any $a \in \mathsf{K}$. For example, if

$$\begin{aligned} f &= x^2 + a_1 x + a_0 \in \mathsf{K}[x;\sigma], \\ g &= x + b_0 \in \mathsf{K}[x;\sigma], \end{aligned}$$

then

$$\begin{aligned} f + g &= x^2 + (a_1 + 1)x + (a_0 + b_0), \\ fg &= x^3 + (a_1 + \sigma^2(b_0))x^2 + (a_1\sigma(b_0) + a_0)x + a_0 b_0, \\ gf &= x^3 + (\sigma(a_1) + b_0)x^2 + (a_1 b_0 + \sigma(a_0))x + a_0 b_0, \end{aligned}$$

where $\sigma^2(a) = \sigma(\sigma(a))$ for any $a \in \mathsf{K}$. When $\sigma = id$, the identity automorphism on K, the ring $\mathsf{K}[x;\sigma]$ is the usual ring of polynomials $\mathsf{K}[x]$ with $xa = ax$ for all $a \in \mathsf{K}$. Skew-polynomial rings have been studied since Ore (1933) and complete treatments are found in Jacobson (1943), McDonald (1974), and Cohn (1985).

Assume throughout that K has size p^ξ, where p is a prime and $\xi \geq 1$. For any $f, g \in \mathsf{K}[x;\sigma]$, $\deg(fg) = \deg f + \deg g$, where $\deg\colon \mathsf{K}[x;\sigma] \setminus \{0\} \to \mathsf{N}$ is the usual polynomial degree function. This implies $\mathsf{K}[x;\sigma]$ is integral (zero is the only zero divisor), and while

not in general a unique factorization domain, it is a principal left ideal ring endowed with a right Euclidean algorithm. As in the commutative case, a non-zero $f \in \mathsf{K}[x;\sigma]$ is *irreducible* if whenever $f = gh$ for some non-zero $g, h \in \mathsf{K}[x;\sigma]$, then either $\deg g = 0$ or $\deg h = 0$. It follows that any $f \in \mathsf{K}[x;\sigma]$ can be written as $f = f_1 \cdots f_k$, where $f_1, \ldots, f_k \in \mathsf{K}[x;\sigma]$ are irreducible. This factorization may not be unique, and adjacent factors may not be interchangeable. Consider two factoring problems:

(i) The complete factorization problem: given any $f \in \mathsf{K}[x;\sigma]$, find irreducible $f_1, \ldots, f_k \in \mathsf{K}[x;\sigma]$ such that $f = f_1 \cdots f_k$.
(ii) The bi-factorization problem: given $f \in \mathsf{K}[x;\sigma]$ and $s < \deg f$, determine if there exist $g, h \in \mathsf{K}[x;\sigma]$ with $f = gh$ and $\deg h = s$, and if so, find such g and h.

This separation of the factoring problem into two cases more completely captures the full complexity of factoring in a non-commutative ring without unique factorization. We reduce the bi-factorization problem to the complete factorization problem. The complete factorization problem is in turn reduced to the problem of determining whether a finite dimensional associative algebra A, over a finite extension field F of F_p, possesses a non-trivial zero divisor, and if so, finding one. Both these reductions are deterministic and polynomial-time. The problem of determining whether A has any non-trivial zero divisors, and producing one if it does, is shown by Rónyai (1987) to be reducible (in deterministic polynomial-time) to factoring polynomials in $\mathsf{F}_p[x]$. Berlekamp's (1970) factoring algorithms for $\mathsf{F}_p[x]$ yield deterministic algorithms for complete and bi-factorization in $\mathsf{K}[x;\sigma]$ requiring time $(n\xi p)^{O(1)}$, and probabilistic algorithms requiring expected time $(n\xi \log p)^{O(1)}$, on input $f \in \mathsf{K}[x;\sigma]$ of degree n.

Our faster algorithm for finding zero divisors in any associative algebra A yields a faster probabilistic algorithm for factoring in skew-polynomial rings. Rónyai's method for finding zero divisors in A is an application of his more general algorithm for computing an explicit decomposition of A, a considerably more complicated problem. His algorithm is quite involved, and Rónyai only shows it to be polynomial-time and does not calculate the running time explicitly. In Section 5 we present a simple, fast, and practical probabilistic solution to the problem of determining whether or not A has any non-trivial zero divisors, and producing a pair multiplying to zero if it does. The algorithm relies on an upper bound on the density of elements in A whose minimal polynomials are irreducible. This algorithm for finding zero divisors yields faster probabilistic algorithms for complete and bi-factorizations of $f \in \mathsf{K}[x;\sigma]$ of degree n, which require expected time $n^4 \cdot (\xi \log p \log n)^{O(1)}$. Rónyai (1990) also presents a number of applications of finding zero divisors in finite associative algebras to problems in computational linear algebra and group theory.

Applications of Skew-Polynomial Rings

Linearized polynomials represent a difficult or "wild" case for algorithms which functionally decompose polynomials, for which no general algorithms are known (Zippel (1991) presents recent progress on this problem, which we discuss below). We present very fast algorithms for the functional decomposition of linearized polynomials.

The linearized polynomials over K, in an indeterminate λ, are those of the form $\sum_{0 \le i \le n} a_i \lambda^{p^i}$ (where $a_0, \ldots, a_n \in \mathsf{K}$). The set A_K of all linearized polynomials in $\mathsf{K}[\lambda]$ forms a ring under the usual polynomial addition ($+$), and functional composition ($\circ$) —

if $f, g \in A_K$ with

$$f = \sum_{0 \le i \le n} a_i \lambda^{p^i}, \quad \text{and } g = \sum_{0 \le j \le r} b_j \lambda^{p^j}, \quad \text{then } f \circ g = f(g(\lambda)) = \sum_{0 \le i \le n} \sum_{0 \le j \le r} a_i b_j^{p^i} \lambda^{p^{i+j}}.$$

The ring $K[x; \psi]$, where $\psi(a) = a^p$ for any $a \in K$, is a skew-polynomial ring with $xa = a^p x$ for all $a \in K$. It is isomorphic to the ring A_K under the map $\Phi: A_K \to K[x; \psi]$, which acts as the identity on K and sends λ^{p^i} to x^i for $i \ge 0$ (see McDonald (1974), Theorem 2.13). Note that if $f \in A_K$, then $\deg \Phi(f) = \log_p(\deg f)$. Computationally, Φ just maps between two interpretations of the input, and is free of charge.

The functional decomposition problem for general polynomials in $K[\lambda]$ comes in two flavours analogous to our complete factorization and bi-factorization problems for $K[x; \sigma]$. Given a polynomial $f \in K[\lambda]$ of degree N, the (functional) complete decomposition problem asks for functionally indecomposable $f_1, \ldots, f_k \in K[\lambda]$ such that $f = f_1 \circ \cdots \circ f_k$ (any $h \in K[\lambda] \setminus K$ is functionally indecomposable if all its bi-decompositions contain a linear composition factor). When $p \nmid N$, the so-called "tame" case for complete decomposition, fast deterministic algorithms for complete decomposition are presented in von zur Gathen *et al.* (1987). When $p \mid N$, the "wild" case, a recent algorithm of Zippel (1991) apparently solves the complete decomposition problem in time $(\deg f)^{O(1)}$, though the exact running time is not calculated. All polynomials $f \in A_K$ have degree p^n for some $n \in \mathbb{N}$, so the complete decomposition problem for linearized polynomials is certainly in the wild case. Given $f \in K[\lambda]$ and $s \in \mathbb{N}$, the (functional) bi-decomposition problem asks if there exist $g, h \in K[\lambda]$ such that $f = g \circ h$ and $\deg h = s$, and if so, find such g, h. The tame case, when $p \nmid (N/s)$, is solved efficiently in von zur Gathen *et al.* (1987). When $p \mid (N/s)$, the wild case, no algorithm is known to solve this problem in time $(\deg f)^{O(1)}$, though a partial solution is provided in von zur Gathen (1990). All non-trivial bi-decompositions of linearized polynomials fall into the wild case, since, if $f \in A_K$ and $f = g \circ h$ for $g, h \in K[\lambda]$, then Dorey & Whaples (1974) show that $\deg g = p^r$ for some $r \in \mathbb{N}$.

When $f \in A_K$, we can solve both the bi-decomposition and complete decomposition problems using our algorithms for complete factorization and bi-factorization in $K[x; \psi]$. The key observation is that we need only consider decompositions of $f \in A_K$ into linearized polynomials: Dorey & Whaples (1974) show that if $f = \bar{f}_1 \circ \cdots \circ \bar{f}_k$ for any $\bar{f}_1, \ldots, \bar{f}_k \in K[\lambda]$, then there exist $f_1, \ldots, f_k \in A_K$ such that $f = f_1 \circ \cdots \circ f_k$ and $\deg f_i = \deg \bar{f}_i$ for $1 \le i \le k$. A complete decomposition of any $f \in A_K$ of degree p^n can be found by finding a complete factorization of $\Phi(f)$ in $K[x; \psi]$. Similarly, the bi-decomposition problem on input $f \in A_K$ of degree p^n and $s \in \mathbb{N}$, is equivalent to the bi-factorization problem in $K[x; \psi]$ on inputs $\Phi(f) \in K[x; \psi]$ and $\log_p s$. Probabilistic versions of these algorithms require expected time $n^4 \cdot (\xi \log p \log n)^{O(1)}$, i.e., they run in time polynomial in $\log(\deg f)$. Deterministic versions require time $(n \xi p)^{O(1)}$.

The Computational Model and Input Specification

We now characterize explicitly any skew-polynomial ring $K[x; \sigma]$ over a finite field K. The automorphism $\sigma: K \to K$ fixes some maximum subfield F of K, and if $[F : F_p] = \eta$ then $F \cong F_q$ where $q = p^\eta$. The only automorphisms of K fixing F are iterates of the Frobenius map $\tau: K \to K$ of K/F, defined by $\tau(a) = a^q$ for all $a \in K$. Thus σ must have the form $\sigma(a) = \tau^\kappa(a) = a^{q^\kappa}$ for all $a \in K$, where $\kappa < \mu = [K : F]$. Furthermore, since F is the largest subfield of K fixed by σ, $\gcd(\mu, \kappa) = 1$.

Part of the input to our algorithms is some auxiliary information to describe $\mathsf{K}[x;\sigma]$: a prime p, the integers η and μ such that $[\mathsf{K}:\mathsf{F}]=\mu$ and $[\mathsf{F}:\mathsf{F}_p]=\eta$, and a *description* of the fields F and K. The description of F consists of a polynomial $\Gamma_\mathsf{F}\in\mathsf{F}_p[x]$ of degree η which is irreducible over F_p. We identify $\mathsf{F}=\mathsf{F}_p[x]/(\Gamma_\mathsf{F})\cong\mathsf{F}_q$, so that F has basis $\mathcal{B}_\mathsf{F}=\{1,\Theta_\mathsf{F},\Theta_\mathsf{F}^2,\ldots,\Theta_\mathsf{F}^{\eta-1}\}$ as an F_p-vector space, where $\Theta_\mathsf{F}=x \bmod \Gamma_\mathsf{F}$ and $\mathsf{F}=\mathsf{F}_p[\Theta_\mathsf{F}]$. The field K is described as an extension of F by a polynomial $\Gamma_\mathsf{K}\in\mathsf{F}[x]$ of degree μ, which is irreducible over F. Identify $\mathsf{K}=\mathsf{F}[x]/(\Gamma_\mathsf{K})$, so K has basis $\mathcal{B}_\mathsf{K}=\{1,\Theta_\mathsf{K},\Theta_\mathsf{K}^2,\ldots,\Theta_\mathsf{K}^{\mu-1}\}$ as an F-vector space, where $\Theta_\mathsf{K}=x \bmod \Gamma_\mathsf{K}$ and $\mathsf{K}=\mathsf{F}[\Theta_\mathsf{K}]$. For reasons described later in this section, we also require the element $\Theta_\mathsf{K}^q=\tau(\Theta_\mathsf{K})$, represented with respect to this basis. Such an element can be computed with $\log q$ operations in K by repeated squaring, though for convenience we consider it pre-computation and do not count it in our complexity analyses. The cost of computing $\tau(\Theta_\mathsf{K})$ is dominated by other costs in our algorithms for both complete and bi-factorization. Note that $\mathsf{K}[x;\sigma]$ is an associative F-algebra with basis $\{\Theta_\mathsf{K}^i x^j \mid 0\le i<\mu,\ j\ge 0\}$. It is not in general a K-algebra, since K is not, in general, in the centre of $\mathsf{K}[x;\sigma]$.

Input size is counted in terms of bits and cost in terms of operations in F. Multiplication in K can be done with $O(\mathrm{M}(\mu))$ operations in F, where $\mathrm{M}(\mu)=\mu^2$ using the usual "school" method, or $\mathrm{M}(\mu)=\mu\log\mu\log\log\mu$ with the algorithm of Cantor & Kaltofen (1987). We can also compute a^{-1} for any $a\in\mathsf{K}$ with $O(\mathrm{M}(\mu)\log\mu)$ operations in F. Using a new algorithm of von zur Gathen & Shoup (1991), for any $a\in\mathsf{K}$ we can compute all conjugates $a,\tau(a),\tau^2(a),\ldots,\tau^{\mu-1}(a)$, of a with $O(\mu\mathrm{M}(\mu)\log\mu+(\mu\log\mu)^2)$ operations in F, assuming that we have computed $\tau(\Theta_\mathsf{K})$ as described above. For convenience we assume throughout the paper that $\mathrm{M}(\mu)=\Omega(\mu\log\mu)$, whence von zur Gathen & Shoup's algorithm requires $O(\mu\mathrm{M}(\mu)\log\mu)$ operations in F. Two $n\times n$ matrices over an arbitrary field L can be multiplied with $O(\mathrm{MM}(n))$ operations in L, where $\mathrm{MM}(n)=n^3$ using the standard algorithm, or $\mathrm{MM}(n)=n^{2.376}$ with the algorithm of Coppersmith & Winograd (1990). With $O(\mathrm{MM}(n))$ operations in L we can also solve a system of n linear equations in n unknowns over L.

2 Basic Operations in $\mathsf{K}[x;\sigma]$

A brief development of the theory of skew-polynomial rings follows, along with algorithms implementing aspects of this theory when appropriate. We begin with an easy observation on the complexity of addition and multiplication in $\mathsf{K}[x;\sigma]$. Let

$$f=\sum_{0\le i\le n} a_i x^i,\qquad g=\sum_{0\le j\le r} b_j x^{p^j}, \tag{1}$$

with $a_0,\ldots,a_n,b_0,\ldots,b_r\in\mathsf{K}$ and $a_n,b_r\neq 0$. To compute fg we expand

$$fg=\sum_{0\le i\le n}\sum_{0\le j\le r} a_i x^i b_j x^j=\sum_{0\le i\le n}\sum_{0\le j\le r} a_i\sigma^i(b_j)x^{i+j},$$

and using the fast method of von zur Gathen & Shoup (1991) to compute all conjugates of b_j for $0\le j\le r$, we obtain the next lemma.

Lemma 1. *Given $f,g\in\mathsf{K}[x;\sigma]$, each of degree at most n, we can compute $f+g$ with $O(n\mu)$ operations in* F, *and fg with $O(n^2\mathrm{M}(\mu)+n\mu\mathrm{M}(\mu)\log\mu)$ operations in* F.

The skew-polynomial ring $\mathsf{K}[x;\sigma]$ has a right division algorithm, and in fact a (right) Euclidean algorithm. The right division algorithm is analogous to the usual one in $\mathsf{K}[x]$. Let $f, g \in \mathsf{K}[x;\sigma]$ be as in (1) with $g \neq 0$: we want to find $Q, R \in \mathsf{K}[x;\sigma]$ such that $f = Qg + R$ and $\deg R < \deg g$ or $R = 0$. The algorithm is trivial if $n < r$ — we know $Q = 0$ and $R = f$ — so assume $n \geq r$. Let $f^{(n)} = f$, and for $n \geq i \geq r$ define $h^{(i)} = (\bar{a}_i/\sigma^{i-r}(b_r)) \cdot x^{i-r}$, where $\bar{a}_i$ is the coefficient of x^i in $f^{(i)}$. Next define $f^{(i-1)} = f^{(i)} - h^{(i)}g \in \mathsf{K}[x;\sigma]$, whence $f^{(i)} = h^{(i)}g + f^{(i-1)}$ and $\deg f^{(i-1)} < \deg f^{(i)}$. Computing $h^{(n)}$, $f^{(n-1)}$, $h^{(n-1)}$, $f^{(n-1)}, \ldots, h^{(r)}$, $f^{(r-1)}$ in sequence, we get $f = Qg + R$ where $Q = h^{(n)} + h^{(n-1)} + \cdots + h^{(r)}$ and $R = f^{(r-1)}$, with $\deg R < \deg g$ or $R = 0$. The Q and R obtained in the division algorithm are unique, as they are in $\mathsf{K}[x]$.

Lemma 2. *If $f, g \in \mathsf{K}[x;\sigma]$ have degree at most n and $g \neq 0$, then computing $Q, R \in \mathsf{K}[x;\sigma]$ such that $f = Qg + R$ and $\deg R < \deg g$ or $R = 0$ requires $O(n^2\mathrm{M}(\mu) + n\mu\mathrm{M}(\mu)\log\mu)$ operations in F.*

Using the above division algorithm, modular equivalence can be meaningfully defined: given $f_1, f_2, g \in \mathsf{K}[x;\sigma]$, we write $f_1 \equiv f_2 \bmod g$ if there exists a $Q \in \mathsf{K}[x;\sigma]$ such that $f_1 - f_2 = Qg$.

Ore (1933) proves the main structure theorem on complete factorizations in $\mathsf{K}[x;\sigma]$, a somewhat simplified version of which is stated below.

Theorem 3 (Ore). *If $f \in \mathsf{K}[x;\sigma]$ factors completely in two ways as $f = f_1 f_2 \cdots f_k = g_1 g_2 \cdots g_t$, where $f_1, \ldots, f_k, g_1, \ldots, g_t \in \mathsf{K}[x;\sigma]$ are irreducible, then $k = t$ and there exists a permutation φ of $\{1, \ldots, k\}$ such that for $1 \leq i \leq k$, $\deg f_i = \deg g_{\varphi(i)}$.*

Common Multiples and Divisors in $\mathsf{K}[x;\sigma]$

From the existence of a right division algorithm in $\mathsf{K}[x;\sigma]$ follows the existence of a right Euclidean scheme in the usual way (see van der Waerden (1970), pp. 55). Assume $f_1, f_2 \in \mathsf{K}[x;\sigma]$ with $\deg f_1 \geq \deg f_2$. At each stage $i > 2$, let f_i be the remainder of f_{i-2} divided on the right by f_{i-1}. We obtain a Euclidean scheme defined by $f_{i-2} = Q_{i-2}f_{i-1} + f_i$ for $1 \leq i \leq t$, where $Q_i, f_i \in \mathsf{K}[x;\sigma]$, $\deg f_i < \deg f_{i-1}$, and $f_{t-1} = Q_{t-1}f_t$. The polynomial $w = af_t \in \mathsf{K}[x;\sigma]$, where $a \in \mathsf{K}$ is chosen such that af_t is monic, is the Greatest Common Right Divisor (GCRD) of f_1 and f_2, denoted $\mathrm{gcrd}(f_1, f_2)$. The GCRD w is the unique monic polynomial of highest degree such that there exist $u_1, u_2 \in \mathsf{K}[x;\sigma]$ with $f_1 = u_1 w$ and $f_2 = u_2 w$. In the usual polynomial ring $\mathsf{K}[x] = \mathsf{K}[x;id]$ we have $\mathrm{gcrd}(f_1, f_2) = \gcd(f_1, f_2)$, the usual greatest common divisor of $f_1, f_2 \in \mathsf{K}[x]$.

Lemma 4. *If $f_1, f_2 \in \mathsf{K}[x;\sigma]$ with $n = \deg f_1 \geq \deg f_2$, then we can compute $\mathrm{gcrd}(f_1, f_2)$ with $O(n^2\mu^2)$ operations in F.*

The existence of a right Euclidean algorithm implies $\mathsf{K}[x;\sigma]$ is a principal left ideal ring, that is, each left ideal is generated by a single polynomial in $\mathsf{K}[x;\sigma]$. If $\mathsf{K}[x;\sigma]f$ and $\mathsf{K}[x;\sigma]g$ are the two left ideals generated by $f, g \in \mathsf{K}[x;\sigma]$ respectively, then the ideal $\mathsf{K}[x;\sigma]\,\mathrm{gcrd}(f, g) = \mathsf{K}[x;\sigma]f + \mathsf{K}[x;\sigma]g$ (see Jacobson (1943), chapter 3).

The set $\mathsf{K}[x;\sigma]f \cap \mathsf{K}[x;\sigma]g$ is also a left ideal, consisting of all polynomials in $\mathsf{K}[x;\sigma]$ which are left multiples of both f and g. Since this left ideal is principal, it is generated by a unique monic $h = \mathrm{lclm}(f, g) \in \mathsf{K}[x;\sigma]$, the Least Common Left Multiple (LCLM)

of f and g. The LCLM h is the unique monic polynomial in $\mathsf{K}[x;\sigma]$ of lowest degree such that there exist $u_1, u_2 \in \mathsf{K}[x;\sigma]$ with $h = u_1 f$ and $h = u_2 g$. Ore (1933) shows that $\deg \operatorname{lclm}(f,g) = \deg f + \deg g - \deg \operatorname{gcrd}(f,g)$. In $\mathsf{K}[x] = \mathsf{K}[x;id]$ the LCLM is simply the usual least common multiple in $\mathsf{K}[x]$, and $\deg \operatorname{lcm}(f,g) = \deg f + \deg g - \deg \gcd(f,g)$.

The LCLM $h \in \mathsf{K}[x;\sigma]$ of $f, g \in \mathsf{K}[x;\sigma]$ is computed by first computing the degree s of $\operatorname{lclm}(f,g)$ using Ore's formula and the algorithm for the GCRD. We use the fact that h is the (non-zero) left multiple of f (in $\mathsf{K}[x;\sigma]$) of lowest degree which is equivalent to zero modulo g. Assuming f and g have degrees n and r respectively, compute the sequence $x^i f = Q_i g + R_i$ for $0 \le i \le s-n$, where $Q_i, R_i \in \mathsf{K}[x;\sigma]$ and $\deg R_i < \deg g = r$. If $u = \sum_{0 \le i \le s-n} c_i x^i \neq 0$, for some $c_0, \ldots, c_{s-n} \in \mathsf{K}$ such that $\sum_{0 \le i \le s-n} c_i R_i = 0$, then $uf \equiv 0 \bmod g$, and uf is a scalar multiple of the LCLM h. Solve for u with linear algebra to obtain the LCLM.

Lemma 5. *Given $f, g \in \mathsf{K}[x;\sigma]$ such that $n = \deg f \ge \deg g$, we can compute $h = \operatorname{lclm}(f,g)$ with $O(\mathrm{MM}(n) \cdot \mathrm{M}(\mu) + n^2\mu^2)$ operations in F.*

A polynomial can also be "decomposed" with respect to LCLM's as follows. Two polynomials $f_1, f_2 \in \mathsf{K}[x;\sigma]$ are *co-prime* if $\operatorname{gcrd}(f_1, f_2) = 1$. Extending this to more polynomials, say $f_1, \ldots, f_\ell \in \mathsf{K}[x;\sigma]$ are *mutually co-prime* if

$$\operatorname{gcrd}(f_i, \operatorname{lclm}(f_1, \ldots, f_{i-1}, f_{i+1}, \ldots, f_\ell)) = 1$$

for $1 \le i \le \ell$, i.e., each f_i is co-prime to the LCLM of the remaining components. This is stronger than the usual pairwise co-primality often seen for $\mathsf{K}[x]$, though the two notions are equivalent in a commutative domain. An *(LCLM) decomposition* of $f \in \mathsf{K}[x;\sigma]$ is a list $(f_1, \cdots, f_\ell) \in \mathsf{K}[x;\sigma]^\ell$ of mutually co-prime polynomials such that $f = \operatorname{lclm}(f_1, \cdots, f_\ell)$; f is *(LCLM) indecomposable* if it admits no non-trivial (LCLM) decompositions.

The following result of Ore (1933) captures the uniqueness of polynomial decompositions in any skew-polynomial ring.

Theorem 6 (Ore). *Let $f \in \mathsf{K}[x;\sigma]$ be monic such that $f = \operatorname{lclm}(f_1, f_2, \ldots, f_\ell)$, where $f_1, \ldots, f_\ell \in \mathsf{K}[x;\sigma]$ are (LCLM) indecomposable and mutually co-prime.*

(i) If $f = \operatorname{lclm}(g_1, g_2, \ldots, g_m)$, where $g_1, \ldots, g_m \in \mathsf{K}[x;\sigma]$ are (LCLM) indecomposable and mutually co-prime, then $\ell = m$ and there exists a permutation φ of $\{1, \ldots, \ell\}$ such that $\deg f_i = \deg g_{\varphi(i)}$ for $1 \le i \le \ell$.

(ii) If, for $1 \le i \le \ell$, we completely factor $f_i = f_{i,1} f_{i,2} \cdots f_{i,s_i}$, where each $f_{i,j} \in \mathsf{K}[x;\sigma]$ is irreducible for $1 \le j \le s_i$, and completely factor $f = h_1 h_2 \cdots h_k$, where $h_1, \ldots, h_k \in \mathsf{K}[x;\sigma]$ are irreducible, there exists a bijection φ from $\{1, \ldots k\}$ to $\{(i,j) \mid 1 \le i \le \ell,\ 1 \le j \le s_i\}$ such that $\deg h_e = \deg f_{\varphi(e)}$ for $1 \le e \le k$.

3 Finding Complete Factorizations

To completely factor a polynomial $f \in \mathsf{K}[x;\sigma]$, we construct a small finite associative algebra D with the property that each non-trivial zero divisor in D yields a non-trivial factorization of f. A candidate for D is the quotient $\mathsf{K}[x;\sigma]/\mathsf{K}[x;\sigma]f$, but it is in general only a $\mathsf{K}[x;\sigma]$-module, and not an algebra. It is only an algebra when $\mathsf{K}[x;\sigma]f$ is a two-sided ideal in $\mathsf{K}[x;\sigma]$. To regain some of the desirable structure of finite algebras, we follow Cohn (1985), Section 0.7, and introduce the concept of an eigenring. Define

$I(\mathsf{K}[x;\sigma]f) = \{u \in \mathsf{K}[x;\sigma] \mid fu \equiv 0 \bmod f\}$, the *idealizer* of $\mathsf{K}[x;\sigma]f$. The set $I(\mathsf{K}[x;\sigma]f)$ is the largest subalgebra of $\mathsf{K}[x;\sigma]$ in which $\mathsf{K}[x;\sigma]f$ is a two-sided ideal. The *eigenring* $E(\mathsf{K}[x;\sigma]f)$ of $\mathsf{K}[x;\sigma]f$ is defined as the quotient $E(\mathsf{K}[x;\sigma]f) = I(\mathsf{K}[x;\sigma]f)/\mathsf{K}[x;\sigma]f$, a finite F-algebra since $\mathsf{K}[x;\sigma]$ is an F-algebra and $\mathsf{K}[x;\sigma]f$ a two-sided ideal in $I(\mathsf{K}[x;\sigma]f)$. If $\deg f = n$, the eigenring $E(\mathsf{K}[x;\sigma]f)$ is isomorphic to the F-algebra

$$D = \{u \in I(\mathsf{K}[x;\sigma]f) \mid \deg u < n\} = \{u \in \mathsf{K}[x;\sigma] \mid fu \equiv 0 \bmod f \text{ and } \deg u < n\},$$

under addition in A and multiplication in A reduced modulo f (i.e., each element in $E(\mathsf{K}[x;\sigma]f)$ is represented by its unique residue modulo f). A basis for D is easily found using linear algebra. The next two theorems demonstrate the usefulness of the eigenring.

Theorem 7. *For $f \in \mathsf{K}[x;\sigma]$, the eigenring $E(\mathsf{K}[x;\sigma]f)$ is a (finite) field if and only if f is irreducible in $\mathsf{K}[x;\sigma]$.*

Theorem 8. *For $f \in \mathsf{K}[x;\sigma]$, if $u, v \in D \setminus \{0\}$ with $uv \equiv 0 \bmod f$, then $\gcrd(f,u) \neq 1$.*

The problem of finding complete factorizations in $\mathsf{K}[x;\sigma]$ is reduced to the problem of finding zero divisors in finite algebras by the following algorithm.

Algorithm: `Complete-Factorization`

Input: $f \in \mathsf{K}[x;\sigma]$ of degree n;

Output: $f_1, \ldots, f_k \in \mathsf{K}[x;\sigma]$ irreducible, with $f = f_1 \cdots f_k$.

(1) Compute a basis for D (above) as an F-algebra;

(2) If D is a field Then Return f;

Else

(3) Find a non-trivial left zero divisor $u \in D$;

(4) Compute $h = \gcrd(f,u)$ and $g \in \mathsf{K}[x;\sigma]$ with $f = gh$.

(5) Recursively factor $g = g_1 \cdots g_r$ and $h = h_1 \cdots h_s$
with $g_1, \ldots, g_r, h_1, \ldots, h_s \in \mathsf{K}[x;\sigma]$ irreducible;

(6) Return $g_1, \ldots, g_r, h_1, \ldots h_s$;

End.

The polynomial $f \in \mathsf{K}[x;\sigma]$ is irreducible if and only if D is a field, and the algorithm halts correctly in this case. If $f \in \mathsf{K}[x;\sigma]$ is reducible then Theorem 7 implies D is not a field, and therefore possesses non-trivial zero divisors. By Theorem 8 each of these zero divisors has a non-trivial GCRD with f, yielding a proper factorization in step 4.

Using Rónyai's reduction from finding zero divisors in finite associative algebras to factoring polynomials in $\mathsf{F}_p[x]$ (see Theorem 15), and our fast probabilistic algorithm for this same problem (see Corollary 20 below), we obtain the following theorem.

Theorem 9. *Let $f \in \mathsf{K}[x;\sigma]$ have degree n. The algorithm* `Complete-Factorization` *correctly finds a complete factorization of f in $\mathsf{K}[x;\sigma]$, and proves:*

(i) the complete factorization problem is deterministically reducible, with $(n\mu \log q)^{O(1)}$ operations in F, to the problem of factoring polynomials in $\mathsf{F}_p[x]$ of degree $(n\xi)^{O(1)}$, and is solvable by a deterministic algorithm requiring $(n\xi p)^{O(1)}$ operations in F.

(ii) the complete factorization problem is solvable by a probabilistic algorithm with an expected $O(n^4\mu \mathrm{M}(\mu) + n^3\mu^2 \mathrm{M}(\mu)\log\mu + n\mathrm{MM}(n\mu) + n^2\mu\log q)$ operations in F.

4 Bi-Factorization With Central Multiples

Finding a polynomial $\hat{f} \equiv 0 \bmod f$, where $\hat{f}$ is in the centre C of $\mathsf{K}[x;\sigma]$, provides the key to bi-factorization. McDonald (1974) shows $C = \mathsf{F}[x^\mu;\sigma] \subseteq \mathsf{K}[x;\sigma]$, the polynomials in x^μ with coefficients in F (recall $\mu = [\mathsf{K} : \mathsf{F}]$). Letting $y = x^\mu$, we identify $C = \mathsf{F}[y]$, the usual ring of polynomials over F in the indeterminate y. In particular, $\mathsf{F}[y]$ is a commutative unique factorization domain. Clearly, the degree (in x) of any element in $\mathsf{F}[y] = \mathsf{F}[x^\mu]$ will always be a multiple of μ. The following theorem demonstrates how a left multiple of f in $\mathsf{F}[y]$ can be used to factor f in $\mathsf{K}[x;\sigma]$.

Theorem 10. *Let $f \in \mathsf{K}[x;\sigma]$ and $\hat{f} \in \mathsf{F}[y]\backslash\{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. If $\hat{f} = \hat{f}_1 \cdots \hat{f}_\ell$ for pairwise co-prime $\hat{f}_1, \ldots, \hat{f}_\ell \in \mathsf{F}[y]$, then $f = \operatorname{lclm}(h_1, \ldots, h_\ell)$, where $h_i = \operatorname{gcrd}(\hat{f}_i, f)$ for $1 \leq i \leq \ell$. Furthermore $h_1, \ldots, h_\ell$ are mutually co-prime, and $\deg f = \sum_{0 \leq i \leq \ell} \deg h_i$.*

The above theorem is used to get a partial decomposition of f by factoring $\hat{f} \in \mathsf{F}[y]$, as a polynomial in y, into pairwise co-prime polynomials in $\mathsf{F}[y]$, and then taking GCRD's between f and each of these factors. However, it is not clear per se that, for every $f \in \mathsf{K}[x;\sigma]$, there exists a non-zero $\hat{f} \in \mathsf{F}[y]$ with $\hat{f} \equiv 0 \bmod f$. In fact, such an $\hat{f}$ does always exist, and can be found efficiently. Call the non-zero $\hat{f} \in \mathsf{F}[y]$ of minimal degree such that $\hat{f} \equiv 0 \bmod f$ the *minimal central left multiple* of f. To find $\hat{f}$, compute the sequence $x^{i\mu} = Q_i f + R_i$ for $0 \leq i \leq n\mu$, where $Q_i, R_i \in \mathsf{K}[x;\sigma]$ and $\deg R_i < \deg f = n$. The set of all polynomials in $\mathsf{K}[x;\sigma]$ of degree less than n forms an F-vector space of dimension $n\mu$. Since there are $n\mu + 1$ polynomials $R_0, \ldots R_{n\mu}$, there exists a minimal $t \leq n\mu$ and $\alpha_0, \ldots, \alpha_t \in \mathsf{F}$, not all zero, such that $\sum_{0 \leq i \leq t} \alpha_i R_i = 0$. The minimal central left multiple $\hat{f}$ of f is then $\hat{f} = \alpha_t^{-1} \sum_{0 \leq i \leq t} \alpha_i x^{\mu i}$, and is found using linear algebra.

Lemma 11. *Given $f \in \mathsf{K}[x;\sigma]$ of degree n, we can find the minimal central left multiple of f with $O(n^3\mu\mathrm{M}(\mu) + n^2\mu^2\mathrm{M}(\mu)\log\mu + \mathrm{MM}(n\mu))$ operations in F.*

A distinct degree factorization (in $\mathsf{F}[y]$), of the minimal central left multiple $\hat{f}$ of $f \in \mathsf{K}[x;\sigma]$, yields the degrees of all factors in any complete factorization of f as shown in the next theorem.

Theorem 12. *Let $f \in \mathsf{K}[x;\sigma]$ and $\hat{f} \in \mathsf{F}[y]\backslash\{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. Furthermore, suppose $\hat{f} = \hat{u}^e$ for some $\hat{u} \in \mathsf{F}[y]\backslash\{0\}$ and $e \geq 1$, where $\hat{u}$ is irreducible as a polynomial in $\mathsf{F}[y]$, and $\deg_x \hat{u} = d\mu$. Then for all complete factorizations $f = f_1 \cdots f_k$, with $f_1, \ldots, f_k \in \mathsf{K}[x;\sigma]$ irreducible in $\mathsf{K}[x;\sigma]$, we have $\deg f_i = d$.*

Corollary 13. *Let $f \in \mathsf{K}[x;\sigma]$ and $\hat{f} \in \mathsf{F}[y]\backslash\{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. Furthermore, suppose $\hat{f} = \hat{u}_1^{e_1}\hat{u}_2^{e_2}\cdots\hat{u}_t^{e_t}$ where $e_1, \ldots, e_t \geq 1$ and $\hat{u}_1, \ldots, \hat{u}_t \in \mathsf{F}[y]$ are distinct and irreducible as polynomials in $\mathsf{F}[y]$, all with the same degree $d\mu$ in x. Then for any complete factorization $f = f_1 \cdots f_k$, with $f_1, \ldots, f_k \in \mathsf{K}[x;\sigma]$ irreducible, we have $\deg f_i = d$.*

Corollary 13 yields an efficient reduction from the bi-factorization problem to the complete factorization problem.

Algorithm: `Bi-Factorization`

Input: $f \in \mathsf{K}[x;\sigma]$ and $s \leq \deg f = n$;

Output: $g, h \in \mathsf{K}[x;\sigma]$ with $\deg h = s$, and $f = gh$, or a message that no such h exists;

(1) Compute the minimal central left multiple $\hat{f} \in \mathsf{F}[y]$ of f;

(2) Find a distinct degree factorization of $\hat{f}$ as $\hat{f} = \hat{f}_1 \hat{f}_2 \cdots \hat{f}_n$, where $\hat{f}_i \in \mathsf{F}[y]$ is such that if $\hat{u} \in \mathsf{F}[y]$ divides $\hat{f}_i$, and $\hat{u}$ is irreducible as a polynomial in $\mathsf{F}[y]$, then $\deg_y \hat{u} = i$, for $1 \leq i \leq n$ (some f_i's may have degree zero).

(3) Find $h_i = \gcrd(\hat{f}_i, f) \in \mathsf{K}[x;\sigma]$ for $1 \leq i \leq n$. Assume $\deg h_i = ie_i$ for some $e_i \in \mathbf{N}$;

(4) Factor each h_i completely in $\mathsf{K}[x;\sigma]$ as $h_i = h_{i,1} h_{i,2} \cdots h_{i,e_i}$, where $\deg h_{i,j} = i$ for $1 \leq j \leq e_i$, and $1 \leq i \leq n$;

(5) Determine if there exists a set $d_1, \ldots, d_n \in \mathbf{N}$ with $d_i \leq e_i$ for $1 \leq i \leq n$, such that $\sum_{0 \leq i \leq n} i d_i = s$;

If such $d_1, \ldots, d_n$ exist then return $g, h \in \mathsf{K}[x;\sigma]$, where $h = \operatorname{lclm}(\bar{h}_1, \bar{h}_2, \ldots, \bar{h}_n)$, and $\bar{h}_i = h_{i,e_i-d_i+1} h_{i,e_i-d_i+2} \cdots h_{i,e_i}$ for $1 \leq i \leq n$, and $g \in \mathsf{K}[x;\sigma]$ is such that $f = gh$;

Otherwise, return "f has no right factor of degree s in $\mathsf{K}[x;\sigma]$";

End.

Any pair $g, h \in \mathsf{K}[x;\sigma]$ produced by the algorithm has $\deg h = s$ and $f = gh$, and if such a bi-factorization exists, this algorithm produces one. To see the former, we note that by Theorem 10, $f = \operatorname{lclm}(h_1, \ldots, h_n)$, where $h_i = \gcrd(\hat{f}_i, f)$ as computed in step 3. By Corollary 13, all complete factorizations of $h_i = h_{i,1} h_{i,2} \cdots h_{i,e_i}$ into irreducible $h_{i,j} \in \mathsf{K}[x;\sigma]$, are such that $\deg h_{i,j} = i$ for $1 \leq j \leq e_i$ and $1 \leq i \leq n$. If $h = \operatorname{lclm}(\bar{h}_1, \ldots, \bar{h}_n)$, then Theorem 6 implies $\deg h = s$. The computed h is a right factor of f since each $\bar{h}_i$ is a right factor of h_i and each h_i is a right factor of f for $1 \leq i \leq n$.

If $f = uv$ for some $u, v \in \mathsf{K}[x;\sigma]$ and $\deg v = s$, this algorithm finds some right factor h of f of degree s. Suppose $v = v_1 v_2 \cdots v_t$, with $v_1, \ldots, v_t \in \mathsf{K}[x;\sigma]$ irreducible. If exactly d_i of the factors $v_1, \ldots, v_t$ have degree i for $1 \leq i \leq n$, then $d_i \leq e_i$ by Theorem 6. Hence $h = \operatorname{lclm}(\bar{h}_1, \ldots, \bar{h}_n)$, computed in step 5, has degree s.

Theorem 14. *Let $f \in \mathsf{K}[x;\sigma]$ have degree n and $s < n$. The algorithm* `bi-factorization` *correctly solves the problem of determining if there exist $g, h \in \mathsf{K}[x;\sigma]$ with $f = gh$, and $\deg h = s$, and if so, find such g, h, and proves:*

(i) the bi-factorization problem is deterministically reducible, with $(n\mu \log q)^{O(1)}$ operations in F, to the problem of factoring polynomials in $\mathsf{F}_p[x]$ of degree $(n\xi)^{O(1)}$, and is solvable with a deterministic algorithm requiring $(n\xi p)^{O(1)}$ operations in F;

(ii) the bi-factorization problem is solvable by a probabilistic algorithm with an expected $O(n^4 \mu \mathrm{M}(\mu) + n^3 \mu^2 \mathrm{M}(\mu) \log \mu + n \mathrm{MM}(n\mu) + n^2 \mu \log q)$ operations in F.

5 Zero Divisors in Finite Associative Algebras

Let A be a finite dimensional associative algebra over a finite field F. That is, A is a finite dimensional vector space over F, with a product $\times : A \to A$, such that A is a ring under $+$ and $\times$ (we write ab for $a \times b$ for any $a, b \in A$). Our goal in this section is to determine (efficiently) if A has any non-trivial zero divisors, and if it does, to produce a pair $b_1, b_2 \in A \setminus \{0\}$ such that $b_1 b_2 = 0$. Familiarity with the basic theorems

and terminology of finite dimensional associative algebras is assumed. A good reference for this information is Pierce (1982), Chapters 1-4. Throughout this section, the word "algebra" (over F) will mean a finite dimensional associative algebra (over F) with a multiplicative identity $1 \in A$. Algebras may or may not be commutative.

The problem of finding zero divisors is known to be polynomial-time reducible to factoring univariate polynomials over a finite field. Let $\mathsf{F} \cong \mathsf{F}_q$ and $q = p^\eta$ for some prime p and $\eta \geq 1$. An algebra A over F is described computationally as an F-vector space of dimension $\nu > 0$ with a supplied basis; elements of A are represented as vectors in F^ν. A representation of the identity element in this basis is assumed to be supplied (alternatively, it may be computed within the allowed time). Addition in A is componentwise and a "black box" algorithm for multiplication in A, requiring a polynomial number of operations in F, is assumed to be provided.

Theorem 15 (Rónyai, 1987). *Let A be an algebra over F with dimension ν, presented as above. The problem of determining if there exist zero divisors in A, and if so finding a pair $b_1, b_2 \in A \setminus \{0\}$ with $b_1 b_2 = 0$, is reducible with $(\nu\eta \log p)^{O(1)}$ operations in F to the problem of factoring polynomials in $\mathsf{F}_p[x]$ of degree $(\nu\eta)^{O(1)}$*

Making use of the factoring algorithms of Berlekamp (1970), this theorem implies the problem of determining whether non-trivial zero divisors exist in A, and finding a pair multiplying to zero if they do, is solvable with a deterministic algorithm requiring $(\nu\eta p)^{O(1)}$ operations in F_p, or a probabilistic algorithm requiring an expected $(\nu\eta \log p)^{O(1)}$ operations in F_p.

Rónyai's algorithm computes an explicit decomposition of A, finding the radical of A, computing a Wedderburn decomposition if A is semi-simple, and computing an isomorphism with a full matrix ring over an extension field of F if A is simple. Needless to say, this algorithm is quite complicated, and while not necessarily inefficient, Rónyai only shows it to be polynomial time and does not calculate the running time explicitly.

We propose a simple algorithm which provides a fast probabilistic algorithm to solve this same problem. Assume A is represented as above. Addition in A requires ν operations in F, while the supplied multiplication algorithm for A requires χ operations in F.

Algorithm: `Zero-Divisor`

Input: an algebra A of dimension ν over F (see above), and a description of F/F_p;
Output: $b_1, b_2 \in A \setminus \{0\}$ with $b_1 b_2 = 0$, or a report that A is a field, or failure;
(1) Choose random $a_1, a_2 \in A$;
  For $b \in \{a_1,\ a_2,\ a_1 a_2 - a_2 a_1\} \setminus \{0\}$ Do
(2) Compute $f = \min_A(b) \in \mathsf{F}[x]$;
(3) Factor f over $\mathsf{F}[x]$;
  If f is reducible (say $f = gh$ for $g, h \in \mathsf{F}[x] \setminus \{0\}$)
(4) Return $g(b)$, $h(b)$;
  Else if $\deg f = \nu$ (and f is irreducible)
(5) Return "A is a field (and has no zero divisors)";
  End For;
(6) Return "Failure";
End.

The minimal polynomial $f = \min_A(b)$ of $b \in A$ is the monic polynomial in $\mathsf{F}[x] \setminus \{0\}$ of minimal degree such that $f(b) = 0$. Clearly f has degree at most ν since $\dim A = \nu$.

To see that the algorithm is correct, examine two cases: when A has (non-zero) zero divisors, and when A is a field. These cases are sufficient by Wedderburn's Theorem (see Lidl & Niederreiter (1983), Section 2.6) which shows any finite algebra whose only zero divisor is zero, is a field. If A is not a field, let $b \in A$ have a reducible minimal polynomial $f \in \mathsf{F}[x]$ (we shall show that there are many such elements). Factoring $f = gh$, for some $g, h \in \mathsf{F}[x] \setminus \mathsf{F}$, yields $f(b) = 0 = g(b)h(b)$, and $g(b)$, $h(b)$ are non-zero since f is the minimal polynomial of b. If some $b \in A$ has a minimal polynomial $f \in \mathsf{F}[x]$ which is irreducible of degree ν, then $A = \mathsf{F}[b]$ and $\mathsf{F}[b]$ is isomorphic to the finite field $\mathsf{F}[x]/(f) \cong \mathsf{F}_{q^\nu}$ under the isomorphism mapping b to $x \bmod f$.

While determining the complexity of this algorithm, assume failure probability $\varrho < 1$. In the sequel we show that $\varrho \leq 1/2$. Computing f in step 2 can be accomplished by first computing the sequence $1, b, b^2, \ldots, b^\nu \in A$, requiring $O(\nu\chi)$ operations in F. Using linear algebra f can then be found with $O(\mathrm{MM}(\nu))$ additional operations in F. Factoring f can be done using the probabilistic algorithm of Berlekamp (1970), with an expected $O(\mathrm{MM}(\nu) + \nu \log q)$ operations in F. Evaluating $g(b)$ and $h(b)$ in step 4 can be done with $O(\nu^2)$ operations in F, using the powers of b computed in step 2.

Theorem 16. *Let A be an algebra with dimension ν over $\mathsf{F} = \mathsf{F}_q$.* `Zero-Divisor` *requires an expected $O((1-\varrho)^{-1} \cdot (\nu\chi + \mathrm{MM}(\nu) + \nu \log q))$ operations in F to determine whether A is a field extension of F, or to produce $b_1, b_2 \in A \setminus \{0\}$ with $b_1 b_2 = 0$.*

The proof that $\varrho \leq 1/2$ for any algebra A is quite involved, the hardest case being when A has a non-trivial zero divisor. When A is a field extension of F, the number of elements in A which generate A is well known to be very high.

Theorem 17. *Let A be field of dimension ν over F. The algorithm* `Zero-Divisor` *with input A reports failure with probability $\varrho \leq 1/2$.*

Now let A be an algebra with at least one non-trivial zero divisor. Steps 2-4 are executed at most 3 times, with $b = a_1$, $b = a_2$, and $b = a_1 a_2 - a_2 a_1$. Except when A is a local algebra (i.e., $A/\mathrm{Rad}(A)$ is a finite field — see Theorem 19), we ignore the possibility of success with $b = a_1 a_2 - a_2 a_1$. For a random selection of b in a non-local A, we bound above by $1/\sqrt{2}$ the probability that steps 2-4 fail to find a zero divisor. When two random choices of b are made, as in the main algorithm, the probability ϱ of failure is at most $1/2$. Let $P(A)$ be the number of elements in A whose minimal polynomial in $\mathsf{F}[x]$ is irreducible. The failure probability ϱ of the algorithm is at most $(P(A)/q^\nu)^2$, so it is sufficient to show $P(A) \leq q^\nu/\sqrt{2}$.

Theorem 18. *Let A be an algebra of dimension ν over F which is not local and has a multiplicative identity. The algorithm* `Zero-Divisor` *fails to find a non-trivial zero divisor in A with probability $\varrho \leq 1/2$.*

Proof (outline). Let $b \in A$ have minimal polynomial $f \in \mathsf{F}[x]$, as computed in step 2. The proof depends upon the type of algebra A:

- A is simple (all two-sided ideals are trivial): It is well known that A is isomorphic to a full matrix algebra over an extension field K of F. Making use of an explicit

formula of Hodges (1958) for the number of matrices over K of given size which are annihilated by a given polynomial in $\mathsf{K}[x]$, we achieve the desired bound.

- A is semi-simple (a direct sum of simple algebras): For an element $b \in A$ to have an irreducible minimal polynomial, its image in each of the simple components of A must have the same irreducible minimal polynomial. This is shown to be unlikely.
- A is not semi-simple (and not local): use the Wedderburn-Malcev Principal Theorem (see McDonald (1974), Theorem 8.28) to write A as the sum of its radical and a semi-simple subalgebra. The larger the radical, the more likely it is the minimal polynomial of a randomly selected element is reducible. □

The case when A is local must be considered separately because there exist non-commutative local algebras (with non-trivial zero divisors) such that the minimal polynomial of every unit is irreducible. Only non-zero zero divisors in A have reducible minimal polynomials, and the number of these is small. When A is not commutative, we know $a_1a_2 - a_2a_1 \in \mathrm{Rad}(A)$. Thus $a_1a_2 - a_2a_1$ is nilpotent, and hence only has an irreducible minimal polynomial when it is zero. In a non-commutative local algebra, the probability that $a_1a_2 - a_2a_1 = 0$ is shown to be small. When A is a commutative local algebra, it is straightforward to show that $P(A) \leq q^{\nu}/\sqrt{2}$, as in Theorem 18.

Theorem 19. *Let A be a local algebra of dimension ν over F, with at least one non-trivial zero divisor. The algorithm* `Zero-Divisor` *fails to find a non-trivial zero divisor with probability $\varrho \leq 1/2$.*

Applying Theorem 17 when A is a field, and Theorems 18 and 19 when it is not, completes our analysis:

Corollary 20 to Theorem 16. *Let A be an algebra of dimension ν over $\mathsf{F} = \mathsf{F}_q$. The algorithm* `Zero-Divisor` *requires an expected $O(\nu\chi + \mathrm{MM}(\nu) + \nu \log q)$ operations in F to determine whether A is a field extension of F, or to produce $b_1, b_2 \in A \setminus \{0\}$ with $b_1b_2 = 0$ (where χ is the number of operations in F required for a single multiplication in A)*

References

E. R. Berlekamp. Factoring polynomials over large finite fields. *Math. Comp.* **24**, pp. 713–735, 1970.

D. G. Cantor and E. Kaltofen. Fast multiplication of polynomials over arbitrary rings. Technical Report 87-35, Dept. of Computer Science, Rensselaer Polytechnic Institute, 1987. *Acta Inform.*, to appear.

P. Cohn. *Free Rings and their Relations.* Academic Press (London), 1985.

D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comp.* **9**, pp. 251–280, 1990.

F. Dorey and G. Whaples. Prime and composite polynomials. *J. Algebra* **28**, pp. 88–101, 1974.

J. von zur Gathen. Functional decomposition of polynomials: the wild case. *J. Symb. Comp.* **10**, pp. 437–452, 1990.

J. von zur Gathen and V. Shoup. Computing frobenius maps and factoring polynomials. Manuscript, 31 pp., 1991.

J. von zur Gathen, D. Kozen, and S. Landau. Functional decomposition of polynomials. In *Proc. 28th Ann. IEEE Symp. Foundations of Computer Science*, pp. 127–131, Los Angeles CA, 1987.

J. H. Hodges. Scalar polynomial equations for matrices over a finite field. *Duke Math. J.* **25**, pp. 291–296, 1958.

N. Jacobson. *The Theory of Rings.* American Math. Soc. (New York), 1943.

R. Lidl and H. Niederreiter. *Finite Fields*, vol. 20 of *Encyclopedia of Mathematics and its Applications.* Addison-Wesley (Reading MA), 1983.

B. McDonald. *Finite Rings with Identity.* Marcel Dekker, Inc. (New York), 1974.

O. Ore. Theory of non-commutative polynomials. *Annals of Mathematics* **34**(22), pp. 480–508, 1933.

R. Pierce. *Associative Algberas.* Springer-Verlag (Heidelberg), 1982.

L. Rónyai. Simple algebras are difficult. In *Proc. 19th ACM Symp. on Theory of Comp.*, pp. 398–408, New York, 1987.

L. Rónyai. Computing the structure of finite algebras. *J. Symb. Comp.* **9**, pp. 355–373, 1990.

B. L. van der Waerden. *Algebra*, vol. 1. Frederick Ungar Publishing Co. (New York), 7th edition, 1970.

R. Zippel. Decomposition of rational functions, 1991. Preprint: Extended abstract in Proc. ISSAC'91.