J. Symbolic Computation (1997) **0**, 1–000. To appear.

Factoring in Skew-Polynomial Rings over Finite Fields[†]

MARK GIESBRECHT

Department of Computer Science, University of Manitoba

Winnipeg, Manitoba, Canada, R3T 2N2

`mwg@cs.umanitoba.ca`

(Received April 10, 1994)

Efficient algorithms are presented for factoring polynomials in the skew-polynomial ring $\mathsf{F}[x;\sigma]$, a non-commutative generalization of the usual ring of polynomials $\mathsf{F}[x]$, where F is a finite field and $\sigma\colon \mathsf{F} \to \mathsf{F}$ is an automorphism (iterated Frobenius map). Applications include fast functional decomposition algorithms for a class of polynomials in $\mathsf{F}[x]$ whose decompositions are "wild" and previously thought to be difficult to compute.

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

$$\begin{aligned} f =& x^2 + a_1x + a_0 \in \mathsf{F}[x;\sigma],\\ g =& x + b_0 \in \mathsf{F}[x;\sigma], \end{aligned}$$

then

$$\begin{aligned} f + g =& x^2 + (a_1 + 1)x + (a_0 + b_0),\\ fg =& x^3 + (a_1 + \sigma^2(b_0))x^2 + (a_1\sigma(b_0) + a_0)x + a_0b_0,\\ gf =& x^3 + (\sigma(a_1) + b_0)x^2 + (a_1b_0 + \sigma(a_0))x + a_0b_0, \end{aligned}$$

where $\sigma^2(a) = \sigma(\sigma(a))$ for any $a \in \mathsf{F}$. When $\sigma = id$, the identity automorphism on F, the ring $\mathsf{F}[x;\sigma]$ is the usual ring of polynomials $\mathsf{F}[x]$ with $xa = ax$ for all $a \in \mathsf{F}$.

Skew-polynomial rings (over more general fields) have been studied since Ore (1933) and complete treatments are found in Jacobson (1943), McDonald (1974), and Cohn (1985). Computationally such polynomials have appeared in the context of uncoupling and solving systems of linear differential and difference equations in closed form (see Grigoriev (1990), Bronstein & Petkovšek (1994, 1996), Singer (1996)). Skew-polynomial

[†] Research was supported in part by Natural Sciences and Engineering Research Council of Canada research grant OGP0155376.

0747–7171/90/000000 + 00 \$03.00/0

rings most generally allow both an automorphism σ of F and a *derivation* $\delta : \mathsf{F} \to \mathsf{F}$, a linear function such that $\delta(ab) = \sigma(a)\delta(b) + \delta(a)b$ for any $a, b \in \mathsf{F}$. The skew-polynomial ring $\mathsf{F}[x; \sigma, \delta]$ is then defined such that $xa = \sigma(a)x + \delta(a)$ for any $a \in \mathsf{F}$. In this paper we only consider the case when $\delta = 0$ and F is finite.

Assume throughout this paper that F has size p^ω, where p is a prime and $\omega \geq 1$. For any $f, g \in \mathsf{F}[x; \sigma]$ we find that $\deg(fg) = \deg f + \deg g$, where $\deg\colon \mathsf{F}[x; \sigma] \setminus \{0\} \to \mathbb{N}$ is the usual polynomial degree function. This implies $\mathsf{F}[x; \sigma]$ is integral (i.e., 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 (see Section 1). As in the commutative case, a non-zero $f \in \mathsf{F}[x; \sigma]$ is *irreducible* if whenever $f = gh$ for some non-zero $g, h \in \mathsf{F}[x; \sigma]$, then either $\deg g = 0$ or $\deg h = 0$. It follows that any $f \in \mathsf{F}[x; \sigma]$ can be written as $f = f_1 \cdots f_k$, where $f_1, \ldots, f_k \in \mathsf{F}[x; \sigma]$ 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 non-constant $f \in \mathsf{F}[x; \sigma]$, find irreducible $f_1, \ldots, f_k \in \mathsf{F}[x; \sigma]$ such that $f = f_1 \cdots f_k$.
(ii) The bi-factorization problem: given any non-constant $f \in \mathsf{F}[x; \sigma] \setminus \{0\}$ and a positive integer $s < \deg f$, determine if there exist $g, h \in \mathsf{F}[x; \sigma]$ with $f = gh$ and $\deg h = s$, and if so, find such g and h.

In a commutative unique factorization domain these two notions of factorizations are computationally equivalent by polynomial-time reductions. However, when we have neither commutativity nor unique factorization (as is the case with skew-polynomial rings), this separation of the factoring problem into two cases more completely captures the full complexity of factoring.

In Sections 2 and 3 we give a reduction from the complete factorization problem for $f \in \mathsf{F}[x; \sigma]$ to the problem of determining whether a finite dimensional associative algebra $\mathfrak{A}$ over a finite field possesses any non-zero zero divisors, and if so, finding a pair multiplying to zero. This reduction is deterministic and requires a number of operations in F which is polynomial in $\deg f$ and $\omega \log p$.

The bi-factorization problem in $\mathsf{F}[x; \sigma]$ is reduced in Section 4 to the complete factorization problem: given $f \in \mathsf{F}[x; \sigma]$ and $s < n = \deg f$, we can determine if there exist $g, h \in \mathsf{F}[x; \sigma]$ such that $f = gh$ and $\deg h = s$ with $(n\omega \log p)^{O(1)}$ operations in F plus the cost of completely factoring polynomials in $\mathsf{F}[x; \sigma]$ of total degree $O(n)$. This yields algorithms for bi-factorization which require $(n\omega p)^{O(1)}$ operations in F, and Las Vegas type probabilistic algorithms which require $(n\omega \log p)^{O(1)}$ operations in F.

In Section 5 we present a fast new algorithm for finding zero divisors in any finite associative algebra. This algorithm is probabilistic of the Las Vegas type and, for an algebra $\mathfrak{A}$ of dimension ν over $\mathbb{F}_q$, requires $O(\nu)$ multiplications in $\mathfrak{A}$ plus about $O(\nu^3 + \nu^2 \log q)$ operations in $\mathbb{F}_q$ to determine whether $\mathfrak{A}$ is a field or to produce a zero divisor in $\mathfrak{A}$. This yields algorithms for complete and bi-factorization in skew-polynomial rings which require $n^4 \cdot (\omega \log p \log n)^{O(1)}$ operations in F.

A paper containing some of this work (with many of the proofs omitted), first appeared in the LATIN'92 conference (Giesbrecht, 1992).

Applications of Skew-Polynomial Rings

An application of skew-polynomials is to the problem of functionally decomposing a class of polynomials which had previously defied polynomial-time decomposition algorithms. Algorithms which functionally decompose polynomials have received considerable attention lately. Given $f \in \mathsf{F}[\lambda]$ in an indeterminate λ, the problem is to determine polynomials $g, h \in \mathsf{F}[\lambda]$ of given degree such that $f = g \circ h = g(h(\lambda))$. Kozen & Landau (1989) and von zur Gathen *et al.* (1987) present polynomial-time (in $\deg f$) solutions to this problem in the "tame" case, when the characteristic p of F does not divide $\deg g$ (see also von zur Gathen (1990a)). In the "wild" case, when $p \mid \deg g$, no general algorithm is known, though partial solutions are given in von zur Gathen (1990b) and Zippel (1991). A very wild type of polynomial is the set of *linearized polynomials* over F, those of the form $\sum_{0 \le i \le n} a_i \lambda^{p^i}$ (where $a_0, \ldots, a_n \in \mathsf{F}$). It turns out that whenever $g, h \in \mathsf{F}[\lambda]$ are such that $f = g \circ h$ then $\deg g = p^r$ for some $r \in \mathbb{N}$, i.e., all functional decompositions of linearized polynomials are wild. In Section 6 we present very fast algorithms for the functional decomposition of linearized polynomials, which run in time polynomial in $\log \deg f$.

Representing Skew-Polynomial Rings

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

Part of the input to our algorithms is some auxiliary information to describe $\mathsf{F}[x; \sigma]$: a prime p, the integers η and μ such that $[\mathsf{F} : \mathsf{K}] = \mu$ and $[\mathsf{K} : \mathbb{F}_p] = \eta$, and a *description* of the fields K and F. The description of K consists of a polynomial $\Gamma_\mathsf{K} \in \mathbb{F}_p[x]$ of degree η which is irreducible over $\mathbb{F}_p$. We identify $\mathsf{K} = \mathbb{F}_p[x]/(\Gamma_\mathsf{K}) \cong \mathbb{F}_q$, so that K has basis $\mathcal{B}_\mathsf{K} = \{1, \Theta_\mathsf{K}, \Theta_\mathsf{K}^2, \ldots, \Theta_\mathsf{K}^{\eta-1}\}$ as an $\mathbb{F}_p$-vector space, where $\Theta_\mathsf{K} = x \bmod \Gamma_\mathsf{K}$ and $\mathsf{K} = \mathbb{F}_p[\Theta_\mathsf{K}]$. The field F is described as an extension of K by a polynomial $\Gamma_\mathsf{F} \in \mathsf{K}[x]$ of degree μ, which is irreducible over K. Identify $\mathsf{F} = \mathsf{K}[x]/(\Gamma_\mathsf{F})$, so F has basis $\mathcal{B}_\mathsf{F} = \{1, \Theta_\mathsf{F}, \Theta_\mathsf{F}^2, \ldots, \Theta_\mathsf{F}^{\mu-1}\}$ as a K-vector space, where $\Theta_\mathsf{F} = x \bmod \Gamma_\mathsf{F}$ and $\mathsf{F} = \mathsf{K}[\Theta_\mathsf{F}]$. We also require the element $\Theta_\mathsf{F}^q = \tau(\Theta_\mathsf{F})$, represented with respect to this basis. This will allow us to make use of von zur Gathen & Shoup's (1992) algorithm to quickly compute all conjugates of an element in F over K (see below). 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 this cost in algorithms using this technique. The cost of computing $\tau(\Theta_\mathsf{F})$ is dominated by other costs in our algorithms for both complete and bi-factorization. Note that $\mathsf{F}[x; \sigma]$ is an associative K-algebra with basis $\{\Theta_\mathsf{F}^i x^j \mid 0 \le i < \mu,\ j \ge 0\}$. It is not in general an F-algebra, since F is not, in general, in the centre of $\mathsf{F}[x; \sigma]$.

Input size is counted in terms of elements in K, and cost in terms of operations in K. For convenience we sometimes use the "soft O" notation in summarizing results: for any $g, h\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$, $g = O\tilde{~}(h)$ if and only if there exists a constant $k > 0$ such that $g = O(h(\log h)^k)$. Multiplication in F can be done with $O(\mathrm{M}(\mu))$ operations in K, where $\mathrm{M}(\mu) = \mu^2$ using the usual "school" method, or $\mathrm{M}(\mu) = \mu \log \mu \log\log \mu$ with the algorithms of Schönhage & Strassen (1971) and Schönhage (1977), or Cantor & Kaltofen (1991). For convenience we assume throughout the paper that $\mathrm{M}(\mu) = \Omega(\mu \log \mu)$. We can also compute a^{-1} for any $a \in \mathsf{F}$ with $O(\mathrm{M}(\mu) \log \mu)$ operations in K. Using an algorithm of von zur Gathen & Shoup (1992), for any $a \in \mathsf{F}$ we can compute all conjugates $a, \tau(a), \tau^2(a), \ldots, \tau^{\mu-1}(a)$ of a with $O(\mu \mathrm{M}(\mu) \log \mu)$ operations in K, assuming that we have computed $\tau(\Theta_{\mathsf{F}})$ as described above. Two $n \times n$ matrices over any 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 asymptotically best known 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.

1. Basic Operations in $\mathsf{F}[x;\sigma]$

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 costs of addition and multiplication in $\mathsf{F}[x;\sigma]$. Let

$$f = \sum_{0 \le i \le n} a_i x^i, \qquad g = \sum_{0 \le j \le r} b_j x^j, \tag{1.1}$$

with $a_0, \ldots, a_n, b_0, \ldots, b_r \in \mathsf{F}$ and $a_n, b_r \ne 0$. Without loss of generality we can assume that $r \le n$. Obviously $f + g$ can be computed with $O(n)$ operations in F or $O(n\mu)$ operations in K. To compute fg we expand

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

Compute $\sigma^i(b_j)$ for $0 \le i < \mu$ and $0 \le j \le r$ with $O(r\mu \mathrm{M}(\mu)\log\mu)$ operations in K, as described in the introduction. Next compute the rn products in F to obtain fg.

Lemma 1.1. *Given $f, g \in \mathsf{F}[x;\sigma]$, each of degree n and r respectively, we can compute $f+g$ with $O(n\mu)$ operations in K, and fg with $O(rn\mathrm{M}(\mu) + r\mu\mathrm{M}(\mu)\log\mu)$ or $O\tilde{~}(rn\mu + r\mu^2)$ operations in K.*

The skew-polynomial ring $\mathsf{F}[x;\sigma]$ has a right division algorithm and a (right) Euclidean algorithm. The right division algorithm is analogous to the usual one in $\mathsf{F}[x]$. Let $f, g \in \mathsf{F}[x;\sigma]$ be as in (1.1) with $g \ne 0$: we want to find $Q, R \in \mathsf{F}[x;\sigma]$ 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 $n \ge r$. Let $f^{(n)} = f$, and for $n \ge i \ge r$ 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 $f^{(i-1)} = f^{(i)} - h^{(i)} g \in \mathsf{F}[x;\sigma]$, whence $f^{(i)} = h^{(i)} g + f^{(i-1)}$ and $\deg f^{(i-1)} < \deg f^{(i)}$. Computing $h^{(n)}$, $f^{(n-1)}$, $h^{(n-1)}$, $f^{(n-2)}, \ldots, h^{(r)}$, $f^{(r-1)}$ 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 $\mathsf{F}[x]$.

Lemma 1.2. *If $f, g \in \mathsf{F}[x;\sigma]$ with $n = \deg f$, $r = \deg g$, and $g \ne 0$, then computing $Q, R \in \mathsf{F}[x;\sigma]$ such that $f = Qg + R$ and $\deg R < \deg g$ or $R = 0$ requires $O(r(n - r)\mathrm{M}(\mu) + r\mu\mathrm{M}(\mu)\log\mu)$ or $O\tilde{~}(r(n-r)\mu + r\mu^2)$ operations in K when $r \le n$.*

Proof. Start by computing $\sigma^i(b_j)$ for $0 \le i < \mu$ and $0 \le j \le r$. This requires $O(r\mu\mathrm{M}(\mu)\log\mu)$ operations in K. At stage i, computing $f^{(i)} - h^{(i)}g$ requires r operations in F. There are at most $n - r$ stages requiring a total of $O(r(n-r))$ operations in F or $O(r(n-r)\mathrm{M}(\mu))$ operations in K. $\square$

Using the above division algorithm, modular equivalence can be meaningfully defined: Given $f_1, f_2, g \in \mathsf{F}[x;\sigma]$, we write $f_1 \equiv f_2 \bmod g$ if and only if there exists a $Q \in \mathsf{F}[x;\sigma]$ such that $f_1 - f_2 = Qg$. It is left as an exercise to the reader that "equivalence modulo h" is indeed an equivalence relation in $\mathsf{F}[x;\sigma]$.

Ore (1933) proved the main structure theorem on complete factorizations in $\mathsf{F}[x;\sigma]$, a somewhat simplified version of which is stated below (this can also be proven as a consequence of the Jordan-Holder theorem — see Jacobson (1943)).

Theorem 1.3. (Ore) *If $f \in \mathsf{F}[x;\sigma]$ factors completely as*

$$\begin{aligned} f &= f_1 f_2 \cdots f_k \\ &= g_1 g_2 \cdots g_t, \end{aligned}$$

where $f_1, \ldots, f_k, \ldots, g_1, \ldots, g_t \in \mathsf{F}[x;\sigma]$ are irreducible, then $k = t$ and there exists a permutation φ of $\{1, \ldots, k\}$ such that for $1 \le i \le k$, $\deg f_i = \deg g_{\varphi(i)}$.

2. Common Multiples and Divisors

From the existence of a right division algorithm in $\mathsf{F}[x;\sigma]$ follows the existence of a right Euclidean scheme in the usual way (see van der Waerden (1970), pp. 55). This implies the existence of *greatest common right divisors* and *least common left multiples* (defined below), the non-commutative analogues of greatest common divisors and least common multiples in a commutative Euclidean domain. It also gives a fast algorithm for computing these.

The Greatest Common Right Divisor (GCRD) of f_1 and f_2, denoted $\operatorname{gcrd}(f_1, f_2)$, is the unique monic polynomial $w \in \mathsf{F}[x;\sigma]$ of highest degree such that there exist $u_1, u_2 \in \mathsf{F}[x;\sigma]$ with $f_1 = u_1 w$ and $f_2 = u_2 w$. It's existence and uniqueness is easily derived from the algorithm presented below, and is demonstrated by Ore (1933). In the usual polynomial ring $\mathsf{F}[x] = \mathsf{F}[x; id]$ we have $\operatorname{gcrd}(f_1, f_2) = \gcd(f_1, f_2)$, the usual greatest common divisor of $f_1, f_2 \in \mathsf{F}[x]$.

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

The set $\mathsf{F}[x;\sigma]f \cap \mathsf{F}[x;\sigma]g$ is also a left ideal, consisting of all polynomials in $\mathsf{F}[x;\sigma]$ which are left multiples of both f and g. Since this left ideal is principal, it is generated by a unique monic $h = \operatorname{lclm}(f, g) \in \mathsf{F}[x;\sigma]$, the Least Common Left Multiple (LCLM) of f and g. The LCLM h is the unique monic polynomial in $\mathsf{F}[x;\sigma]$ of lowest degree such that there exist $u_1, u_2 \in \mathsf{F}[x;\sigma]$ with $h = u_1 f$ and $h = u_2 g$. In $\mathsf{F}[x] = \mathsf{F}[x; id]$ the LCLM is simply the usual least common multiple in $\mathsf{F}[x]$.

Assume $f_1, f_2 \in \mathsf{F}[x;\sigma] \setminus \{0\}$ with $\delta_1 := \deg f_1$, $\delta_2 := \deg f_2$ and $\delta_1 \ge \delta_2$. We can compute an extended Euclidean scheme in $\mathsf{F}[x;\sigma]$ much as we can in $\mathsf{F}[x]$. For $3 \le i \le k+1$, let $f_i, q_i \in \mathsf{F}[x;\sigma]$ be the quotient and remainder of f_{i-2} divided by f_{i-1},

$$f_i = f_{i-2} - q_i f_{i-1}, \quad \delta_i := \deg(f_i), \quad \delta_{i-1} > \delta_i \text{ for all } i \text{ with } 3 \le i \le k, \quad f_{k+1} = 0.$$

Analogous to the commutative case we have $f_k = \operatorname{gcrd}(f_1, f_2)$. Furthermore, let $s_i, t_i \in$

$\mathsf{F}[x;\sigma]$ be the multipliers in the extended Euclidean scheme, i.e.,

$$\begin{aligned} s_1 &:= 1; \quad s_2 := 0; \quad s_i := s_{i-2} - q_i s_{i-1}; \\ t_1 &:= 0; \quad t_2 := 1; \quad t_i := t_{i-2} - q_i t_{i-1}; \end{aligned}$$

$$s_i f_1 + t_i f_2 = f_i,$$

for all i with $3 \le i \le k+1$. It follows by an easy induction on i that for all $3 \le i \le k+1$ $\deg(s_i) = \delta_2 - \delta_{i-1}$ and $\deg(t_i) = \delta_1 - \delta_{i-1}$.

To obtain the LCLM, note that $s_{k+1}f_1 + t_{k+1}f_2 = f_{k+1} = 0$, hence $v = s_{k+1}f_1 = -t_{k+1}f_2$ is a common multiple of f_1 and f_2. We see that $\deg v = (\delta_2 - \delta_k) + \delta_1 = \deg f_1 + \deg f_2 - \deg \operatorname{gcrd}(f_1, f_2)$, which Ore (1933) shows to be the degree of the LCLM. It must therefore be the case that $v = \operatorname{lclm}(f_1, f_2)$.

A similar presentation of the extended Euclidean scheme (and computation of GCRD's and LCLM's) in skew-polynomial rings may be found in Bronstein & Petkovšek (1994), Section 1.

Lemma 2.1. *If* $f_1, f_2 \in \mathsf{F}[x;\sigma]$ *with* $n = \deg f_1 \ge \deg f_2$*, then we can compute* $\operatorname{gcrd}(f_1, f_2)$ *and* $\operatorname{lclm}(f_1, f_2)$ *with* $O(n^2\mathrm{M}(\mu)\mu\log\mu)$ *or* $O\tilde{~}(n^2\mu^2)$ *operations in* K.

Proof. For $3 \le i \le k+1$ we can compute f_i with $O((\delta_{i-2}-\delta_{i-1})\mathrm{M}(\mu)+\delta_{i-1}\mathrm{M}(\mu)\mu\log\mu)$ operations in K. The cost to compute all the f_i's is therefore

$$\begin{aligned} \sum_{3\le i\le k} (\delta_{i-2} - \delta_{i-1}) \cdot \delta_{i-1}\mathrm{M}(\mu) + \delta_{i-1}\mathrm{M}(\mu)\mu\log\mu \\ \le \sum_{3\le i\le k} (\delta_{i-2}^2 - \delta_{i-1}^2)\mathrm{M}(\mu) + \sum_{3\le i\le k} \delta_{i-1}\mathrm{M}(\mu)\mu\log\mu \\ \le \delta_1^2\mathrm{M}(\mu) + \delta_2^2\mathrm{M}(\mu)\mu\log\mu. \end{aligned}$$

We can compute all s_i, for $1 \le i \le k+1$, with

$$\begin{aligned} \sum_{3\le i\le k} (\delta_2 - \delta_{i-3})\mu + (\delta_{i-2}-\delta_{i-1})(\delta_2 - \delta_{i-2})\mathrm{M}(\mu) + (\delta_2 - \delta_{i-2})\mu\mathrm{M}(\mu)\log\mu \\ \le \delta_2^2\mu + \delta_2^2\mu\mathrm{M}(\mu)\log\mu + \delta_1\delta_2\mathrm{M}(\mu). \end{aligned}$$

All t_i's, for $1 \le i \le k+1$ can be computed with similar cost. Therefore, in total, our algorithm requires $O(n^2\mu\mathrm{M}(\mu)\log\mu)$ operations in K. □

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

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

for $1 \le i \le l$, 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 $\mathsf{F}[x]$, though the two notions are equivalent in a commutative domain. An *LCLM-decomposition* of $f \in \mathsf{F}[x;\sigma]$ is a list $(f_1, \ldots, f_l) \in \mathsf{F}[x;\sigma]^l$ of mutually co-prime polynomials such that $f = \operatorname{lclm}(f_1, \ldots, f_l)$; f is *LCLM-indecomposable* if it admits no non-trivial LCLM-decompositions. If $f_1, \ldots, f_l$ are also all irreducible if $\mathsf{F}[x;\sigma]$, then f is said to be *completely irreducible* (see Ore (1933) – he refers to "LCLM-indecomposable" polynomials as simply "indecomposable" polynomials). The following result of Ore (1933) captures the uniqueness of polynomial decompositions in any skew-polynomial ring.

THEOREM 2.2. (ORE, 1933) *Let* $f \in \mathsf{F}[x;\sigma]$ *be monic such that* $f = \mathrm{lclm}(f_1,\ldots,f_l)$, *where* $f_1,\ldots,f_l \in \mathsf{F}[x;\sigma]$ *are LCLM-indecomposable and mutually co-prime.*

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

(ii) If, for $1 \le i \le l$, $f_i = f_{i,1}f_{i,2}\cdots f_{i,s_i}$, *where each* $f_{i,j} \in \mathsf{F}[x;\sigma]$ *is irreducible for* $1 \le j \le s_i$, *and* $f = h_1h_2\cdots h_k$, *where* $h_1,\ldots,h_k \in \mathsf{F}[x;\sigma]$ *are irreducible, then there exists a bijection* φ *from* $\{1,\ldots,k\}$ *to* $\{(i,j) \mid 1 \le i \le l,\ 1 \le j \le s_i\}$ *such that* $\deg h_e = \deg f_{\varphi(e)}$ *for* $1 \le e \le k$.

3. Finding Complete Factorizations

To completely factor any non-constant $f \in \mathsf{F}[x;\sigma]$, we construct a small finite associative algebra $\mathfrak{A}$ over K with the property that each non-zero zero divisor in $\mathfrak{A}$ yields a non-zero factorization of f. An associative algebra $\mathfrak{A}$ over K is a K-vector space with a product $\times\colon \mathfrak{A} \to \mathfrak{A}$ such that $\mathfrak{A}$ is a ring under $+$ and $\times$ (we write ab for $a \times b$ for $a, b \in \mathfrak{A}$). A candidate for $\mathfrak{A}$ is the quotient $\mathsf{F}[x;\sigma]/\mathsf{F}[x;\sigma]f$, but it is in general only a $\mathsf{F}[x;\sigma]$-module, and not an algebra. It is only an algebra when $\mathsf{F}[x;\sigma]f$ is a two-sided ideal in $\mathsf{F}[x;\sigma]$. To regain some of the desirable structure of finite algebras, we follow Cohn (1985), Section 0.7, and introduce the concept of an eigenring. For notational brevity, let $\mathfrak{S} = \mathsf{F}[x;\sigma]$ throughout this section. Define $I(\mathfrak{S}f) = \{u \in \mathfrak{S} \mid fu \equiv 0 \bmod f\}$, the *idealizer* of $\mathfrak{S}f$. The set $I(\mathfrak{S}f)$ is the largest subalgebra of $\mathfrak{S}$ in which $\mathfrak{S}f$ is a two-sided ideal. The *eigenring* $E(\mathfrak{S}f)$ of $\mathfrak{S}f$ is defined as the quotient $E(\mathfrak{S}f) = I(\mathfrak{S}f)/\mathfrak{S}f$, a finite K-algebra since $\mathfrak{S}$ is an K-algebra and $\mathfrak{S}f$ a two-sided ideal in $I(\mathfrak{S}f)$. If $\deg f = n$, the eigenring $E(\mathfrak{S}f)$ is isomorphic to the K-algebra

$$\mathfrak{A} = \{u \in I(\mathfrak{S}f) \mid \deg u < n\} = \{u \in \mathfrak{S} \mid fu \equiv 0 \bmod f \text{ and } \deg u < n\},$$

under addition in $\mathfrak{S}$ and multiplication in $\mathfrak{S}$ reduced modulo f (i.e., each element in $E(\mathfrak{S}f)$ is represented by its unique residue modulo f). The key facts about $E(\mathfrak{S}f)$, which we shall prove in the sequel, are that it is a field if and only if f is irreducible, and that non-zero zero divisors in $E(\mathfrak{S}f)$ allow us to compute non-zero factors of f efficiently.

To prove the desired properties of the eigenring we need to characterize the centre $\mathfrak{C}$ of $\mathfrak{S}$, and the two-sided ideals in $\mathfrak{S}$. McDonald (1974), pages 24-25, shows $\mathfrak{C} = \mathsf{K}[x^\mu;\sigma] \subseteq \mathfrak{S}$, the polynomials in x^μ with coefficients in K. This follows since the subset of $\mathfrak{S}$ commuting with Θ_F is $\mathsf{F}[x^\mu]$, while the subset of $\mathfrak{S}$ commuting with x is $\mathsf{K}[x]$. The elements Θ_F and x generate $\mathfrak{S}$ as a K-algebra, whence $\mathsf{K}[x^\mu] = \mathsf{F}[x^\mu] \cap \mathsf{K}[x]$ is the centre of $\mathfrak{S}$. Letting $y = x^\mu$, we identify $\mathfrak{C} = \mathsf{K}[y]$, the usual ring of polynomials over K in the indeterminate y, so in particular, $\mathfrak{C}$ is a commutative unique factorization domain. The degree (in x) of any element in $\mathfrak{C}$ will always be a multiple of μ. Clearly, any $\hat{f} \in \mathsf{K}[y]$ generates a two-sided ideal $\mathfrak{S}\hat{f}$. In fact, the two-sided ideals in $\mathfrak{S}$ are exactly those of the form $\mathfrak{S}(\hat{f}x^s)$ for some $\hat{f} \in \mathsf{K}[y]$ and $s \in \mathbb{N}$. The maximal (non-zero) two-sided ideals in $\mathfrak{S}$ are $\mathfrak{S}x$, and $\mathfrak{S}\hat{u}$, where $\hat{u} \in \mathsf{K}[y] \setminus \{y\}$ is irreducible as a polynomial in y. An important characteristic of the left ideal $\mathfrak{S}f$ is the largest two sided ideal $\mathfrak{o}$ it contains, called the *bound* for $\mathfrak{S}f$ (see Jacobson (1943), Chapter 3, Sections 5 and 6). Closely related to the bound for $\mathfrak{S}f$

is the *minimal central left multiple* $\hat{f} \in \mathsf{K}[y]$ of f, the polynomial in $\mathsf{K}[y]$ of minimal degree which is a left multiple of f. Such a polynomial always exists (we show how to construct it efficiently in Lemma 4.2), and if $\mathrm{gcrd}(f, x) = 1$ then $\mathfrak{o} = \mathfrak{S}\hat{f}$. More generally, if $f = \mathrm{lclm}(f_0, x^s)$ for some $s \geq 0$ and some $f_0 \in \mathfrak{S}$ co-prime with x and with minimal central left multiple $\hat{f}_0 \in \mathsf{K}[y]$, then $\mathfrak{o} = \mathfrak{S} \cdot \hat{f}_0 x^s$.

We recall some basic facts about associative algebras before we proceed. An algebra $\mathfrak{A}$ is *simple* if its only two-sided ideals are $\{0\}$ and $\mathfrak{A}$, and is *semi-simple* if it is a direct sum of simple algebras. Next, we summarize some well known facts about finite simple algebras (see for example Lang (1984), Chapter 17).

FACT 3.1. *Let $\mathfrak{A}$ be a finite, simple algebra of dimension d over K, and let L be a left ideal in $\mathfrak{A}$.*

(i) $\mathfrak{A}$ is isomorphic to $\mathsf{E}^{m\times m}$, where $m \geq 1$, E is the centre of $\mathfrak{A}$ and a finite extension field of degree r over K, and $d = m^2 r$.

(ii) There exist minimal left ideals $L_1, \ldots, L_m \subseteq \mathfrak{A}$ and $l \leq m$ such that $L = L_1 \oplus \cdots \oplus L_l$ and $\mathfrak{A} = L_1 \oplus \cdots \oplus L_m$. Furthermore, each minimal left ideal has dimension rm as an K-vector space.

(iii) There exist maximal left ideals $M_1, \ldots, M_m \subseteq \mathfrak{A}$ and $k \leq m$ such that $L = M_1 \cap \cdots \cap M_k$, $M_1 \cap \cdots \cap M_m = \{0\}$, and $M_i + (M_1 \cap \cdots \cap M_{i-1} \cap M_{i+1} \cap \cdots \cap M_m) = \mathfrak{A}$ for $1 \leq i \leq m$. Furthermore, each maximal left ideal has dimension $rm^2 - rm$ as a K-vector space.

A K-algebra of particular interest is $\mathfrak{A} = \mathfrak{S}/\mathfrak{S}\hat{f}$, where $\hat{f} \in \mathsf{K}[y] \setminus \{y\}$ is irreducible as a polynomial in y. Since $\mathfrak{S}\hat{f}$ is a maximal two-sided ideal in $\mathfrak{S}$, $\mathfrak{A}$ is a simple algebra. From $\mathfrak{S}$, $\mathfrak{A}$ inherits the property of being a left principal ideal ring. Suppose $g_1 + \mathfrak{S}\hat{f}$ and $g_2 + \mathfrak{S}\hat{f}$ are in some left ideal $J \subseteq \mathfrak{A}$, where $g_1, g_2 \in \mathfrak{S}$. Then there exist $h_1, h_2 \in \mathfrak{S}$ such that $h_1 g_1 + h_2 g_2 = \mathrm{gcrd}(g_1, g_2)$ and

$$(h_1 + \mathfrak{S}\hat{f})(g_1 + \mathfrak{S}\hat{f}) + (h_2 + \mathfrak{S}\hat{f})(g_2 + \mathfrak{S}\hat{f}) = \mathrm{gcrd}(g_1, g_2) + \mathfrak{S}\hat{f} \in J.$$

Thus, left ideals are closed under GCRD's (of their pre-images in $\mathfrak{S}$) and each left ideal J in $\mathfrak{A}$ is generated by some unique $g + \mathfrak{S}\hat{f}$, where $g \in \mathfrak{S}$ is monic of minimal degree. Since $\mathrm{gcrd}(g, \hat{f}) + \mathfrak{S}\hat{f} \in J$ and g has minimal degree, g is a right factor of $\hat{f}$. We call such a g the *minimal modular generator* of J. The following lemma relates left ideals in $\mathfrak{A}$ with the left ideals in $\mathfrak{S}$ generated by their minimal modular generators.

LEMMA 3.2. *Let $J_1, J_2 \subseteq \mathfrak{A}$ be non-zero left ideals in $\mathfrak{A}$, with respective minimal modular generators $g_1, g_2 \in \mathfrak{S}$.*

(i) The left ideal $J_3 = J_1 \cap J_2$ in $\mathfrak{A}$ has minimal modular generator $g_3 = \mathrm{lclm}(g_1, g_2)$ if $J_3 \neq \{0\}$. Otherwise $J_3 = \{0\}$ and $\hat{f} = \mathrm{lclm}(g_1, g_2)$.

(ii) The left ideal $J_4 = J_1 + J_2$ in $\mathfrak{A}$ has minimal modular generator $g_4 = \mathrm{gcrd}(g_1, g_2)$.

PROOF. To prove (i) we note that $\mathrm{lclm}(g_1, g_2) + \mathfrak{S}\hat{f} \in J_3$, so we must show that $\mathrm{lclm}(g_1, g_2)$ is the minimal modular generator of J_3. Suppose $h + \mathfrak{S}\hat{f} \in J_3$ for some $h \in \mathfrak{S}$. Then $h \equiv w_1 g_1 \equiv w_2 g_2 \bmod \hat{f}$ for some $w_1, w_2 \in \mathfrak{S}$. It follows that since both g_1 and g_2 are right factors of $\hat{f}$, they are also both right factors of h as well. Thus $h \equiv 0 \bmod \mathrm{lclm}(g_1, g_2)$, so the pre-image in $\mathfrak{S}$ of every element in J_3 is in $\mathfrak{S}\,\mathrm{lclm}(g_1, g_2)$. If

$\operatorname{lclm}(g_1, g_2) \neq \hat{f}$ then $\operatorname{lclm}(g_1, g_2)$ is the minimal modular generator of J_3. If $\operatorname{lclm}(g_1, g_2) = \hat{f}$ then $J_3 = \{0\}$.

To prove (ii), we note

$$J_1 + J_2 = (\mathfrak{S}g_1 \bmod \mathfrak{S}\hat{f}) + (\mathfrak{S}g_2 \bmod \mathfrak{S}\hat{f}) = (\mathfrak{S}g_1 + \mathfrak{S}g_2) \bmod \hat{f} = \mathfrak{S}u \bmod \mathfrak{S}\hat{f},$$

where $u = \operatorname{gcrd}(g_1, g_2)$. Thus $u + \mathfrak{S}\hat{f}$ generates J_4 and $\hat{f} \equiv 0 \bmod u$ since both g_1 and g_2 are right factors of $\hat{f}$. For any $h \in \mathfrak{S}$ such that $h + \mathfrak{S}\hat{f} \in J_4$, $h \equiv Qu \bmod \hat{f}$ for some $Q \in \mathfrak{S}$, and since u is a right factor of $\hat{f}$ and Qu, u is a right factor of h. It follows that u is the polynomial in $\mathfrak{S}$ of smallest degree such that $u + \mathfrak{S}$ generates J_4, that is, u is the minimal modular generator of J_4. $\square$

The next theorem characterizes the LCLM-decompositions of those $f \in \mathfrak{S}$ whose minimal central left multiples are irreducible as polynomials in y.

Theorem 3.3. *For $f \in \mathfrak{S}$, the eigenring $E(\mathfrak{S}f)$ is a (finite) field if and only if f is irreducible in $\mathfrak{S}$.*

Proof. If f is irreducible McDonald (1974), Exercise 2.24, shows $E(\mathfrak{S}f)$ is a finite field.

We now show that if f is reducible then $E(\mathfrak{S}f)$ possesses zero divisors. If f is reducible and LCLM-decomposable, then $f = \operatorname{lclm}(f_1, f_2)$, where $f_1, f_2 \in \mathfrak{S} \setminus \mathsf{F}$ and $g_1 f_1 + g_2 f_2 = 1$ for some $g_1, g_2 \in \mathfrak{S}$. Note that if $h \equiv 0 \bmod f_1$ and $h \equiv 0 \bmod f_2$ for any $h \in \mathfrak{S}$, then $h \equiv 0 \bmod f$. We now construct a pair of non-zero zero divisors in $E(\mathfrak{S}f)$. Let $h_1 = g_1 f_1$ and $h_2 = g_2 f_2$, neither of which are equivalent to zero modulo f. Then

$$fh_1 = f(1 - g_2 f_2) \equiv 0 \bmod f_2 \quad \text{and} \quad fh_1 = f g_1 f_1 \equiv 0 \bmod f_1,$$

so $fh_1 \equiv 0 \bmod f$. Similarly $fh_2 \equiv 0 \bmod f$, so $h_1, h_2 \in I(\mathfrak{S}f)$. Moreover, $h_1 h_2 = h_1 - h_1^2 = h_2 - h_2^2$, which is equivalent to zero modulo both f_1 and f_2, and hence modulo f. Thus $(h_1 + \mathfrak{S}f)(h_2 + \mathfrak{S}f) \equiv 0 \bmod f$ and $h_1 + \mathfrak{S}f$ and $h_2 + \mathfrak{S}f$ are non-zero zero divisors in $E(\mathfrak{S}f)$.

If f is reducible but indecomposable then Jacobson (1943), Theorem 3.13, shows $\hat{f} = \hat{g}^e \in \mathsf{K}[y]$ is the minimal central left multiple of f, where $\hat{g} \in \mathsf{K}[y]$ is irreducible as a polynomial in y and $e \geq 1$. If $\hat{g} = y$ then $f = x^d$ for some $d \geq 2$, and $\mathfrak{S}f$ is a two-sided ideal in $\mathfrak{S}$. Thus $E(\mathfrak{S}/\mathfrak{S}f) = \mathfrak{S}/\mathfrak{S}f$ and $x + \mathfrak{S}f$ is a zero divisor in $E(\mathfrak{S}/\mathfrak{S}f)$. Now assume that $\hat{g} \neq y$. The set $f + \mathfrak{S}\hat{f}$ generates a left ideal L in $\mathfrak{A} = \mathfrak{S}/\mathfrak{S}\hat{f}$. We now show that $e > 1$ by contradiction. Suppose that $e = 1$ so that $\mathfrak{A}$ is simple. Then by Fact 3.1, there exist maximal left ideals $M_1, \ldots, M_m \subseteq \mathfrak{A}$ such that $M_1 \cap \cdots \cap M_k = L$ and $M_i + (M_1 \cap \cdots \cap M_{i-1} \cap M_{i+1} \cap \cdots \cap M_m) = \mathfrak{A}$. Since f is reducible we know L is not maximal so $k \geq 2$. Each maximal left ideal M_i has an irreducible minimal modular generator $h_i \in \mathfrak{S}$, for $1 \leq i \leq k$. By Lemma 3.2, $f = \operatorname{lclm}(h_1, \ldots, h_k)$. Moreover, since $M_i + (M_1 \cap \cdots \cap M_{i-1} \cap M_{i+1} \cap \cdots \cap M_m) = \mathfrak{A}$ for $1 \leq i \leq m$, we know $M_i + (M_1 \cap \cdots \cap M_{i-1} \cap M_{i+1} \cap \cdots \cap M_k) = \mathfrak{A}$ for $1 \leq i \leq k$, and by Lemma 3.2 it follows that $\operatorname{gcrd}(h_i, \operatorname{lclm}(h_1, \ldots, h_{i-1}, h_{i+1}, \ldots, h_k)) = 1$ for $1 \leq i \leq k$. Thus, $h_1, \ldots, h_k$ are pairwise co-prime. In particular, since $k \geq 2$, f is decomposable, which is a contradiction. Assume then that $e \geq 2$. Note that $\hat{g} \in I(\mathfrak{S}f)$ and $\hat{g} \not\equiv 0 \bmod f$, so the image $\hat{g} + \mathfrak{S}f \in E(\mathfrak{S}f)$ of $\hat{g}$ in $E(\mathfrak{S}f)$ is non-zero. Since $\hat{g}^e \equiv 0 \bmod f$, we see that $(\hat{g} + \mathfrak{S}f)(\hat{g}^{e-1} + \mathfrak{S}f) \equiv 0 \bmod \mathfrak{S}f$ and $E(\mathfrak{S}f)$ is not a field. $\square$

Next we show that left zero divisors in $\mathfrak{A} \cong E(\mathfrak{S}f)$ allow us to split f.

THEOREM 3.4. *For* $f \in \mathfrak{S}$, *if* $u, v \in \mathfrak{A} \setminus \{0\}$ *with* $uv \equiv 0 \bmod f$, *then* $\operatorname{gcrd}(f, u) \neq 1$.

PROOF. Suppose $\operatorname{gcrd}(f, u) = 1$. There exist $s, t \in \mathfrak{S}$ such that $sf + tu = 1$ and $sfv + tuv = v$. But $fv \equiv 0 \bmod f$ and $uv \equiv 0 \bmod f$ so $v \equiv 0 \bmod f$, a contradiction. □

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

Algorithm: `Complete-Factorization`
Input: $f \in \mathsf{F}[x;\sigma]$ of degree n;
Output: $f_1, \ldots, f_k \in \mathsf{F}[x;\sigma]$ irreducible, with $f = f_1 \cdots f_k$.
(1) Compute a basis for $\mathfrak{A}$ (above) as a K-algebra;
(2) If $\mathfrak{A}$ is a field Then Return f;
Else
(3) Find a non-zero left zero divisor $u \in \mathfrak{A}$;
(4) Compute $h = \operatorname{gcrd}(f, u)$ and $g \in \mathsf{F}[x;\sigma]$ with $f = gh$;
(5) Recursively factor $g = g_1 \cdots g_r$ and $h = h_1 \cdots h_s$
with $g_1, \ldots, g_r, h_1, \ldots, h_s \in \mathsf{F}[x;\sigma]$ irreducible;
(6) Return $g_1, \ldots, g_r, h_1, \ldots h_s$;
End.

The polynomial $f \in \mathsf{F}[x;\sigma]$ is irreducible if and only if $\mathfrak{A}$ is a field, and the algorithm halts correctly in this case. If $f \in \mathfrak{S}$ is reducible then Theorem 3.3 implies $\mathfrak{A}$ is not a field, and therefore possesses non-zero zero divisors (Wedderburn's Theorem implies every finite algebra, whose only zero divisor is zero, is a field). By Theorem 3.4 any left zero divisor has a non-zero GCRD with f, yielding a proper factorization in step 4. The algorithm recurses on g and h, each of which has degree less than n. Since there is no recursion when f is irreducible, the procedure `Complete-Factorization` will be called at most n times, each time on a polynomial of degree at most n.

The number of operations in K required by each step is now determined:

Step 1. A basis for $\mathfrak{A}$ can be found as follows. Let $W \subseteq \mathsf{F}[x;\sigma]$ be the set of all $g \in \mathsf{F}[x;\sigma]$ with $\deg g < n$. As a K-vector space W is isomorphic to $\mathsf{F}[x;\sigma]/\mathsf{F}[x;\sigma]f$, with basis

$$\{\Theta_{\mathsf{F}}^i x^j \mid 0 \leq i < \mu,\ 0 \leq j < n\},$$

and dimension $n\mu$. Multiplication on the left by f induces an K-linear map $T\colon W \to W$: if $u \in W$ then $T(u) = v \equiv fu \bmod f$, for some $v \in W$. The elements of $\mathfrak{A}$ are exactly those elements in the null space of T, a basis which is found by constructing a matrix for T (an $n\mu \times n\mu$ matrix over K) and then using linear algebra techniques to compute a basis for the null space. This matrix is computed by evaluating T at each of the basis elements of W, i.e., finding $f\Theta_{\mathsf{F}}^i x^j \bmod f$ for $0 \leq i < \mu$ and $0 \leq j < n$, requiring a total of $O(n^3\mu \mathrm{M}(\mu) + n^2\mu^2\mathrm{M}(\mu)\log\mu)$ operations in K. The linear algebra to find a basis for the null space of T, and hence for $\mathfrak{A}$, requires $O(\mathrm{MM}(n\mu))$ additional operations in K.

Steps 2–3. We have not yet shown how to determine if $\mathfrak{A}$ is a field, and if it is not, produce a non-zero zero divisor in $\mathfrak{A}$. In Rónyai (1987) it is shown that this problem is reducible, with $(n\mu\log q)^{O(1)}$ operations in K, to factoring polynomials in

$\mathbb{F}_p[x]$ of degree $(n\omega)^{O(1)}$ (recall $[\mathsf{F} : \mathbb{F}_p] = \omega$). A faster Las Vegas type probabilistic algorithm for this problem is presented in Section 5, and requires $\mathrm{O}(n\mu\chi + \mathrm{MM}(n\mu) + \mathrm{M}(n\mu)\log(n\mu)\log q)$ operations in K, where χ operations in K are required to multiply two elements of $\mathfrak{A}$. A multiplication in $\mathfrak{A}$ can be done with $O(n^2\mathrm{M}(\mu) + n\mu\mathrm{M}(\mu)\log\mu)$ operations in K, so we can determine if $\mathfrak{A}$ is a field, and if not, find a zero divisor in $\mathfrak{A}$, with $O(n^3\mu\mathrm{M}(\mu) + n^2\mu^2\mathrm{M}(\mu)\log\mu + \mathrm{MM}(n\mu) + \mathrm{M}(n\mu)\log(n\mu)\log q)$ or $O\tilde{}(n^3\mu^2 + n^2\mu^3 + \mathrm{MM}(n\mu) + n\mu\log q)$ operations in K.

Step 4. The polynomials g and h can be computed with $O(n^2\mathrm{M}(\mu)\mu\log\mu)$ operations in K by Lemma 2.1.

As noted above, there are at most n recursive calls, each on a polynomial of degree less than n. This yields the following theorem:

THEOREM 3.5. *Let $f \in \mathsf{F}[x;\sigma]$ have degree n. The algorithm* `Complete-Factorization` *correctly finds a complete factorization of f in $\mathsf{F}[x;\sigma]$, and proves:*

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

(ii) the complete factorization problem is solvable by a Las Vegas type algorithm with $O(n^4\mu\mathrm{M}(\mu) + n^3\mu^2\mathrm{M}(\mu)\log\mu + n\mathrm{MM}(n\mu) + n\mathrm{M}(n\mu)\log(n\mu)\log q)$ or $O\tilde{}(n^4\mu^2 + n^3\mu^3 + n\mathrm{MM}(n\mu) + n^2\mu\log q)$ operations in K.

4. Bi-Factorization With Two-Sided Ideals

Finding the minimal central left multiple $\hat{f} \in \mathsf{F}[y]$ of an $f \in \mathsf{F}[x;\sigma]$ provides the key to bi-factorization. The following theorem demonstrates how the factorization over $\mathsf{K}[y]$ of $\hat{f}$ yields a partial factorization of f. Once again, we let $\mathfrak{S} = \mathsf{F}[x;\sigma]$ throughout this section.

THEOREM 4.1. *Let $f \in \mathsf{F}[x;\sigma]$ and $\hat{f} \in \mathsf{K}[y] \setminus \{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. If $\hat{f} = \hat{f}_1 \cdots \hat{f}_l$ for pairwise co-prime $\hat{f}_1, \ldots, \hat{f}_l \in \mathsf{K}[y]$, then $f = \mathrm{lclm}(h_1, \ldots, h_l)$, where $h_i = \mathrm{gcrd}(\hat{f}_i, f)$ for $1 \le i \le l$, and $h_1, \ldots, h_l$ are pairwise co-prime.*

PROOF. From the definitions of GCRD and LCLM in Section 2, this theorem can be restated in terms of ideals: $\mathfrak{S}f = \mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l$ and

$$\mathfrak{L}_i + (\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_{i-1} \cap \mathfrak{L}_{i+1} \cap \cdots \cap \mathfrak{L}_l) = \mathfrak{S},$$

for $1 \le i \le l$, where $\mathfrak{L}_i = \mathfrak{S}f + \mathfrak{S}\hat{f}_i = \mathfrak{S}h_i$.

We start by showing that $\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l = \mathfrak{S}f$. For any $u \in \mathfrak{S}f$, we know $u \equiv 0 \bmod f$ and hence $u \equiv 0 \bmod h_i$ and $u \in \mathfrak{L}_i$ for $1 \le i \le l$. Thus $\mathfrak{S}f \subseteq \mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l$. To show $\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l \subseteq \mathfrak{S}f$ assume $u \in \mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l$. Thus $u = v_i f + w_i \hat{f}_i$ for some $v_i, w_i \in \mathfrak{S}$, and $u \equiv v_i f \bmod \hat{f}_i$, for $1 \le i \le l$. We know that $\mathfrak{S}/\mathfrak{S}\hat{f}$ is isomorphic as a ring to $\mathfrak{S}/\mathfrak{S}\hat{f}_1 \oplus \cdots \oplus \mathfrak{S}/\mathfrak{S}\hat{f}_l$. By the Chinese remainder theorem, since u is a left multiple of f modulo each $\hat{f}_i$, u is a left multiple of f modulo $\hat{f}$, i.e., $u \equiv vf \bmod \hat{f}$ for some $v \in \mathfrak{S}$. From this and the fact that $\hat{f} \equiv 0 \bmod f$ we see $u \equiv 0 \bmod f$, and therefore that $u \in \mathfrak{S}f$ and $\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_l = \mathfrak{S}f$.

To show that

$$\mathfrak{L}_i + (\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_{i-1} \cap \mathfrak{L}_{i+1} \cap \cdots \cap \mathfrak{L}_l) = \mathfrak{S},$$

for $1 \le i \le l$, we observe that $\mathfrak{S}\hat{f} = \mathfrak{S}\hat{f}_1 \cap \cdots \cap \mathfrak{S}\hat{f}_l$, where

$$\mathfrak{S}\hat{f}_i + (\mathfrak{S}\hat{f}_1 \cap \cdots \cap \mathfrak{S}\hat{f}_{i-1} \cap \mathfrak{S}\hat{f}_{i+1} \cap \cdots \cap \mathfrak{S}\hat{f}_l) = \mathfrak{S}.$$

This follows since $\mathsf{K}[y]$ is a unique factorization domain. Thus, for $1 \le i \le l$, there exists $u_i \in \mathfrak{S}\hat{f}_i$ and $v_i \in \mathfrak{S}\hat{f}_1 \cap \cdots \cap \mathfrak{S}\hat{f}_l$ such that $u_i + v_i = 1$. Since $\mathfrak{L}_i \supseteq \mathfrak{S}\hat{f}_i$ for $1 \le i \le l$, we know $u_i \in \mathfrak{L}_i$ and $v_i \in \mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_{i-1} \cap \mathfrak{L}_{i+1} \cap \cdots \cap \mathfrak{L}_l$, so

$$1 \in \mathfrak{L}_i + (\mathfrak{L}_1 \cap \cdots \cap \mathfrak{L}_{i-1} \cap \mathfrak{L}_{i+1} \cap \cdots \cap \mathfrak{L}_l) = \mathfrak{S}$$

for $1 \le i \le l$. □

The above theorem is used to get a partial decomposition of f by factoring its minimal central left multiple $\hat{f} \in \mathsf{K}[y]$, as a polynomial in y, into pairwise co-prime polynomials in $\mathsf{K}[y]$ and then taking GCRD's between f and each of these factors. We now address the question of finding $\hat{f}$.

LEMMA 4.2. *Given $f \in \mathsf{F}[x;\sigma]$ of degree n, we can find the minimal central left multiple of f with $O(n^3\mu\mathrm{M}(\mu) + n^2\mu^2\mathrm{M}(\mu)\log\mu + \mathrm{MM}(n\mu))$ or $O\tilde{}(n^3\mu^2 + n^2\mu^3 + \mathrm{MM}(n\mu))$ operations in K.*

PROOF. First, compute the sequence $x^{i\mu} = Q_i f + R_i$ for $0 \le i \le n\mu$, where $Q_i, R_i \in \mathsf{F}[x;\sigma]$ and $\deg R_i < \deg f = n$. The set of all polynomials in $\mathsf{F}[x;\sigma]$ of degree less than n forms a K-vector space of dimension $n\mu$, where each coefficient in F is expanded with respect to the given basis of F/K. For $1 \le i \le n\mu$, if

$$R_i = \sum_{0\le j<n}\sum_{0\le l<\mu} w_{jl}\Theta_{\mathsf{F}}^l x^j, \quad \text{let } \bar{R}_i = (w_{0,0}, w_{0,1}, \ldots, w_{n-1,\mu-1})^t \in \mathsf{K}^{n\mu\times 1}.$$

Since there are $n\mu + 1$ polynomials $R_0, \ldots, R_{n\mu}$, there exists a minimal $t \le n\mu$ and $\alpha_0, \ldots, \alpha_t \in \mathsf{K}$, not all zero, such that $\sum_{0\le i\le t}\alpha_i R_i = 0$ and hence that $\sum_{0\le i\le t}\alpha_i \bar{R}_i = 0$. The minimal central left multiple $\hat{f}$ of f is then $\hat{f} = \alpha_t^{-1}\sum_{0\le i\le t}\alpha_i x^{\mu i}$.

Let B be the $n\mu\times(n\mu+1)$ matrix over K whose ith column is $\bar{R}_{i-1}$. Since R_t is linearly dependent (over K) on $R_0, \ldots, R_{t-1}$, and $R_{t+i} \equiv x^\mu R_{t+i-1} \bmod f$, it follows that R_{t+i} is also linearly dependent (over K) on $R_0, \ldots, R_{t-1}$, for $i \ge 0$. Thus $t = \mathrm{rank}B$, and if

$$v = (\alpha_0, \ldots, \alpha_t, 0, \ldots, 0) \in \mathsf{K}^{n\mu},$$

then $Bv = 0$. Conversely, any non-zero $v \in \mathsf{K}^{n\mu}$ of the form

$$v = (\beta_0, \ldots, \beta_t, 0, \ldots, 0) \in \mathsf{K}^{n\mu},$$

and in the null space of B, yields a scalar multiple $\sum_{0\le i\le t}\beta_i x^i$ of the minimal central left multiple of $\hat{f}$. Hence we can now solve for the minimal central left multiple of $\hat{f}$ with linear algebra over K.

To determine the cost of this algorithm, start by computing $X_i \equiv x^i \bmod f$ with $\deg X_i < \deg f$, for $0 \le i \le n + \mu - 1$; this can be accomplished with $O(n^2\mu\mathrm{M}(\mu) + n\mu^2\mathrm{M}(\mu)\log\mu)$ operations in K. Now for any

$$g = \sum_{0\le i<n} b_i x^i$$

with $b_0, \ldots, b_{n-1} \in \mathsf{F}$, we know

$$x^\mu g = \sum_{0 \le i < n} b_i x^{i+\mu} \equiv \sum_{0 \le i < n} b_i X_{i+\mu} \bmod f.$$

Using the fact that $R_i \equiv x^\mu R_{i-1} \bmod f$ for $i > 0$, we can compute R_i from R_{i-1} as an F-linear combination of $X_\mu, \ldots, X_{\mu+n-1}$, with $O(n^2\mu)$ operations in F. Finding $R_0, \ldots, R_{n\mu}$ then takes $O(n^3\mu\mathrm{M}(\mu))$ operations in K, and the linear algebra to compute $\hat{f}$ from $\bar{R}_0, \ldots, \bar{R}_{n\mu}$ requires an additional $O(\mathrm{MM}(n\mu))$ operations in K. $\square$

The next lemma characterizes the LCLM-decompositions of those $f \in \mathsf{F}[x;\sigma]$ whose minimal central left multiples are irreducible as polynomials in y.

Theorem 4.3. *Let $f \in \mathsf{F}[x;\sigma]$ and $\hat{f} \in \mathsf{K}[y]$ the minimal central left multiple of f with $\deg \hat{f} = n\mu$. If $\hat{f}$ is irreducible as a polynomial in y and $f = gh$ for some irreducible $h \in \mathsf{F}[x;\sigma]$, then $\deg h = n$.*

Proof. If $\hat{f} = y = x^\mu$ then $n = 1$. The only irreducible right factor of $\hat{f}$ in this case is x, which has degree 1.

Assume then that $\hat{f} \neq y$. The quotient $A = \mathfrak{S}/\mathfrak{S}\hat{f}$ is a simple algebra (since $\mathfrak{S}\hat{f}$ is a maximal left ideal in $\mathfrak{S}$) of dimension $n\mu^2$ over K. By Fact 3.1, for some $m \geq 1$, A is isomorphic to the ring of all $m \times m$ matrices over the centre E of A, where E is an extension field of K. If $[\mathsf{E} : \mathsf{K}] = r$, then $n\mu^2 = rm^2$.

The centre E of A is simply the image of $\mathsf{K}[y]$ in A. To see this, let $g \in \mathfrak{S}$ and $\bar{g}$ its image in A. If $g \in \mathsf{K}[y]$ then $\bar{g}$ is certainly in E. Conversely, if $\bar{g} \in \mathsf{E}$, then we may assume $\deg g < \deg_x \hat{f}$, i.e., we choose the polynomial of least degree in $\mathfrak{S}$ which is equivalent to $\bar{g}$ modulo $\hat{f}$. Now $g\Theta_{\mathsf{F}} - \Theta_{\mathsf{F}} g \equiv 0 \bmod \hat{f}$ and $gx - xg \equiv 0 \bmod \hat{f}$, since g is in the centre of A. The degrees of $g\Theta_{\mathsf{F}}$ and $\Theta_{\mathsf{F}} g$ are both less than $\deg_x \hat{f}$, so $g\Theta_{\mathsf{F}} - \Theta_{\mathsf{F}} g = 0$, which is only true if $g \in \mathsf{F}[x^\mu]$. Assume now that $\mu \geq 2$ (if $\mu = 1$ then $\mathfrak{S} = \mathsf{F}[x]$ and the theorem is trivially true). Since $g \in \mathsf{F}[x^\mu]$ it has degree less than $n\mu - 1$ and both gx and xg have degrees less than $n\mu$, whence $gx - xg = 0$. The elements x and Θ_{F} generate $\mathfrak{S}$ as a K-algebra, and since g commutes with both of them, g must be in the centre of $\mathfrak{S}$. Therefore the centre E of A is the image of $\mathsf{K}[y]$ in A, and has degree n over K. It follows that $r = [\mathsf{E} : \mathsf{K}] = n$, $m = \mu$, and $A \cong \mathsf{E}^{\mu \times \mu}$.

Maximal left ideals in A are exactly those whose minimal modular generators are irreducible in $\mathfrak{S}$. In particular, the left ideal generated by $h + \mathfrak{S}$ is maximal. By Fact 3.1, each maximal left ideal in A has dimension $n\mu^2 - n\mu$ as a K-vector space. Since the left ideal in A generated by $h + \mathfrak{S}$ is equal to the set of left multiples of h of degree less than $n\mu$, reduced modulo $\hat{f}$, it has dimension $n\mu^2 - n\deg h$ as a K-vector space. Thus $\deg h = n$. $\square$

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

Theorem 4.4. *Let $f \in \mathsf{F}[x;\sigma]$ and $\hat{f} \in \mathsf{K}[y]\backslash\{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. Furthermore, suppose $\hat{f} = \hat{g}^e$ for some $\hat{g} \in \mathsf{K}[y]\backslash\{0\}$ and $e \geq 1$, where $\hat{g}$ is irreducible as a polynomial in $\mathsf{K}[y]$, and $\deg_x \hat{g} = d\mu$. Then for all complete factorizations $f = f_1 \cdots f_l$, with $f_1, \ldots, f_l \in \mathsf{F}[x;\sigma]$ irreducible in $\mathsf{F}[x;\sigma]$, we have $\deg f_i = d$.*

PROOF. Suppose $f = f_1 \cdots f_k$, where $f_1, \ldots, f_k \in \mathsf{F}[x;\sigma]$ are irreducible. We proceed by induction on k. If $k = 1$, then Jacobson (1943), Chapter 12, Theorem 13 shows $e = 1$, and by Theorem 4.3, $\deg f_1 = d$. Assume that the theorem is true for complete factorizations with fewer than k irreducible factors. The minimal central left multiple of f_k must be irreducible as a polynomial in y and must divide $\hat{g}^e$, whence $\hat{g} \equiv 0 \bmod f_k$. By Theorem 4.3, $\deg f_k = d$. Moreover, by Jacobson (1943), Chapter 12, Theorem 12, $\hat{g}^e \equiv 0 \bmod f_1 \cdots f_{k-1}$, so by induction then $\deg f_1 = \cdots = \deg f_{k-1} = d$. $\square$

COROLLARY 4.5. *Let $f \in \mathsf{F}[x;\sigma]$ and $\hat{f} \in \mathsf{K}[y] \setminus \{0\}$ be such that $\hat{f} \equiv 0 \bmod f$. Furthermore, suppose $\hat{f} = \hat{g}_1^{e_1} \hat{g}_2^{e_2} \cdots \hat{g}_l^{e_l}$ where $e_1, \ldots, e_l \geq 1$ and $\hat{g}_1, \ldots, \hat{g}_l \in \mathsf{K}[y]$ are distinct and irreducible as polynomials in $\mathsf{K}[y]$, all with the same degree $d\mu$ in x. Then for any complete factorization $f = f_1 \cdots f_k$, with $f_1, \ldots, f_k \in \mathsf{F}[x;\sigma]$ irreducible, we have $\deg f_i = d$.*

PROOF. By Theorem 4.1 we know $f = \mathrm{lclm}(h_1, \ldots, h_l)$, where $h_i = \mathrm{gcrd}(\hat{f}_i^{e_i}, f)$ for $1 \leq i \leq l$. Since $\hat{f}_i^{e_i} \equiv 0 \bmod h_i$ for $1 \leq i \leq l$, we know by Theorem 4.3 that every complete factorization

$$h_i = h_{i,1} h_{i,2} \cdots h_{i,s_i},$$

where each $h_{i,j} \in \mathfrak{S}$ is irreducible, is such that $\deg h_{i,j} = d$ for $1 \leq j \leq s_i$ and $1 \leq i \leq l$. Theorem 2.2 implies that if

$$f = f_1 f_2 \cdots f_k,$$

where $f_1, \ldots, f_k \in \mathfrak{S}$ are irreducible, then $\deg f_i = d$ for $1 \leq i \leq k$. $\square$

Corollary 4.5 yields an efficient reduction from the bi-factorization problem to the complete factorization problem.

Algorithm: `Bi-Factorization`

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

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

(1) Compute the minimal central left multiple $\hat{f} \in \mathsf{K}[y]$ 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 $\hat{f}_i \in \mathsf{K}[y]$ is such that if $\hat{g} \in \mathsf{K}[y]$ divides $\hat{f}_i$, and $\hat{g}$ is irreducible as a polynomial in $\mathsf{K}[y]$, then $\deg_y \hat{g} = i$, for $1 \leq i \leq n$ (some f_i's may have degree zero).

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

(4) Factor each h_i completely in $\mathsf{F}[x;\sigma]$ 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 $d_1, \ldots, d_n \in \mathbb{N}$ 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 $g, h \in \mathsf{F}[x;\sigma]$, where $h = \mathrm{lclm}(\bar{h}_1, \bar{h}_2, \ldots, \bar{h}_n)$, 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 $g \in \mathsf{F}[x;\sigma]$ is such that $f = gh$;

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

End.

Any pair $g, h \in \mathsf{F}[x;\sigma]$ 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 4.1, $f = \mathrm{lclm}(h_1, \ldots, h_n)$, where $h_i = \mathrm{gcrd}(\hat{f}_i, f)$ as computed in step 3. By Corollary 4.5, all complete factorizations of $h_i = h_{i,1}h_{i,2}\cdots h_{i,e_i}$ into irreducible $h_{i,j} \in \mathsf{F}[x;\sigma]$, are such that $\deg h_{i,j} = i$ for $1 \le j \le e_i$ and $1 \le i \le n$. If $h = \mathrm{lclm}(\bar{h}_1, \ldots, \bar{h}_n)$, then Theorem 2.2 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 $1 \le i \le n$.

If $f = uv$ for some $u, v \in \mathsf{F}[x;\sigma]$ and $\deg v = s$, this algorithm finds some right factor h of f of degree s. Suppose $v = v_1v_2\cdots v_t$, with $v_1, \ldots, v_t \in \mathsf{F}[x;\sigma]$ irreducible. If exactly d_i of the factors $v_1, \ldots, v_t$ have degree i for $1 \le i \le n$, then $d_i \le e_i$ by Theorem 2.2. Hence $h = \mathrm{lclm}(\bar{h}_1, \ldots, \bar{h}_n)$, computed in step 5, has degree s.

The the number of operations required by the algorithm `Bi-factorization` is now determined.

Step 1. Computing the minimal central left multiple of f requires $O(n^3\mu\mathrm{M}(\mu)+\mathrm{MM}(n\mu))$ operations in K by Lemma 4.2.

Step 2. Distinct degree factorization can be computed with $O(\mathrm{M}(n\mu)\cdot\mathrm{M}(\sqrt{n\mu})\cdot\sqrt{n\mu}\log n + \mathrm{M}(n\mu)\log q)$ operations in K, using the algorithm of von zur Gathen & Shoup (1992).

Step 3. Use the fact that if $u \equiv v \bmod f$ then $\mathrm{gcrd}(f, u) = \mathrm{gcrd}(f, v)$, for any $u, v \in \mathsf{F}[x;\sigma]$. During the computation of the minimal central left multiple $\hat{f}$ in step 1 we found $R_j \equiv x^{\mu j} \bmod f$, where $\deg R_j < \deg f$ for $1 \le j \le n\mu$. Each polynomial $\hat{f}_i \bmod f$ is then just a linear combination of the R_j's, so we can compute $\hat{f}_1 \bmod f$, $\hat{f}_2 \bmod f, \ldots, \hat{f}_n \bmod f$ within the time required for step 1. The required GCRD's can now be computed with $O(n^3\mathrm{M}(\mu)\mu\log\mu)$ operations in K.

Step 4. Completely factoring n polynomials in $\mathsf{F}[x;\sigma]$ of total degree n is accomplished by the algorithm `Complete-Factorization` of Section 3. The cost of completely factoring n polynomials in $\mathsf{F}[x;\sigma]$ of total degree n is at most the cost of completely factoring a single polynomial in $\mathsf{F}[x;\sigma]$ of degree n. By Theorem 3.5 this can be done with a (Las Vegas) probabilistic algorithm requiring $O(n^4\mu\mathrm{M}(\mu) + n^3\mu^2\mathrm{M}(\mu)\log\mu + n\mathrm{MM}(n\mu) + n\mathrm{M}(n\mu)\log(n\mu)\log q)$ operations in K. By the same theorem, this problem is deterministically reducible, with $(n\mu\log p)^{O(1)}$ operations in K, to the problem of factoring univariate polynomials in $\mathbb{F}_p[x]$ of degree $(n\mu)^{O(1)}$.

Step 5. Determining if $d_1, \ldots, d_n$ exist, and finding them if they do, while not performed with K-operations (and hence not really "counted" in our model of computation), can be accomplished efficiently with a simple dynamic programming algorithm. The LCLM can be performed with $O(n^2\mathrm{M}(\mu)\mu\log\mu)$ operations in K.

THEOREM 4.6. *Let $f \in \mathsf{F}[x;\sigma]$ have degree n and $s < n$. The algorithm* `bi-factorization` *above correctly solves the problem of determining if there exist $g, h \in \mathsf{F}[x;\sigma]$ 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 K, to the problem of factoring polynomials in $\mathbb{F}_p[x]$ of degree $(n\omega)^{O(1)}$, and is solvable with a deterministic algorithm requiring $(n\omega p)^{O(1)}$ operations in K;

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

5. A Fast Algorithm for Finding Zero Divisors

Let $\mathfrak{A}$ be any finite dimensional associative algebra (with identity) of dimension ν over a finite field $\mathsf{K} \cong \mathbb{F}_q$, where q is a power of a prime p. $\mathfrak{A}$ is described computationally as a K-vector space with a basis $\mathcal{B} = \{w_1, \ldots, w_\nu\} \subseteq \mathfrak{A}$. A representation of $1 \in \mathfrak{A}$ is assumed to be supplied. Addition in $\mathfrak{A}$ is component-wise and a "black box" algorithm for multiplication in $\mathfrak{A}$, which requires χ operations in K, is assumed to be provided.

Our algorithm is based on finding and factoring the minimal polynomial of a randomly selected element $a \in \mathfrak{A}$, then evaluating one of these factors at a. Recall that the *minimal polynomial* $\min_{\mathsf{K}}(b) \in \mathsf{K}[x]$ of $b \in \mathfrak{A}$ is the monic polynomial $f \in \mathsf{K}[x]$ of minimal degree such that $f(b) = 0$. It does not depend on how $\mathfrak{A}$ is represented as an extension of K, and has degree at most $\dim \mathfrak{A} = \nu$.

Algorithm: `FindZeroDivisor`
Input: an algebra $\mathfrak{A}$ of dimension ν over K (see above);
Output: $b_1, b_2 \in \mathfrak{A} \setminus \{0\}$ with $b_1 b_2 = 0$, or a report that $\mathfrak{A}$ is a field, or failure;
(1) Choose random $a_1, a_2 \in \mathfrak{A}$;
  For $b \in \{a_1,\, a_2,\, a_1a_2 - a_2a_1\} \setminus \{0\}$ Do
(2)   Compute $f = \min_{\mathsf{K}}(b) \in \mathsf{K}[x]$;
(3)   Factor f in $\mathsf{K}[x]$;
    If f is reducible with $f = gh$ for $g, h \in \mathsf{K}[x] \setminus \{0\}$
(4)   Return $g(b)$, $h(b)$;
    Else if $\deg f = \nu$ (and f is irreducible)
(5)   Return "$\mathfrak{A}$ is a field (and has no zero divisors)";
  End For;
(6) Return "Failure";
End.

To see that the algorithm is correct, examine two cases: when $\mathfrak{A}$ has non-trivial zero divisors, and when $\mathfrak{A}$ is a (finite) 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 (commutative) field. If $\mathfrak{A}$ is not a field, let $b \in \mathfrak{A}$ have a reducible minimal polynomial $f \in \mathsf{K}[x]$ (we shall show that there are many such elements). Factoring $f = gh$, for some $g, h \in \mathsf{K}[x] \setminus \mathsf{K}$, 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 $b \in \mathfrak{A}$ has a minimal polynomial $f \in \mathsf{K}[x]$ which is irreducible of degree ν, then $\mathfrak{A} = \mathsf{K}[b]$ and $\mathsf{K}[b]$ is isomorphic to the finite field $\mathsf{K}[x]/(f) \cong \mathsf{K}_{q^\nu}$ under the isomorphism mapping b to $x \bmod f$.

While determining the complexity of this algorithm, assume failure probability $\varrho < 1$. In Theorems 5.2 and 5.3 and 5.9 below, we show that in fact $\varrho \le 8/9$. Computing f in step 2 can be accomplished by first computing the sequence $1, b, b^2, \ldots, b^\nu \in \mathfrak{A}$, requiring $O(\nu\chi)$ operations in K. Using linear algebra f can then be found with $O(\mathrm{MM}(\nu))$ additional operations in K. Factoring f can be done using the Las Vegas type probabilistic algorithm of Berlekamp (1970), with $O(\mathrm{MM}(\nu) + \mathrm{M}(\nu)\log\nu\log q)$ operations in K. Evaluating $g(b)$ and $h(b)$ in step 4 can be done with $O(\nu^2)$ operations in K, using the powers of b computed in step 2. We have shown the following.

THEOREM 5.1. *Let $\mathfrak{A}$ be an algebra with dimension ν over $\mathsf{K} = \mathbb{F}_q$. The algorithm* `FindZeroDivisor` *requires $O(\nu\chi + \mathrm{MM}(\nu) + \mathrm{M}(\nu)\log\nu\log q)$ operations in K to determine whether $\mathfrak{A}$ is a field extension of K, or to produce $b_1, b_2 \in \mathfrak{A} \setminus \{0\}$ with $b_1b_2 = 0$. It requires as additional input (to the description of $\mathfrak{A}$) two randomly selected elements of $\mathfrak{A}$ and fails with probability $\varrho < 1$.*

The proof that the probability of failure ϱ satisfies $\varrho \le 8/9$ for any algebra $\mathfrak{A}$ is quite involved, the hardest case being when $\mathfrak{A}$ has a non-trivial zero divisor. In the course of the proofs that follow, we will need both upper and lower bounds for the number of irreducible polynomials of a fixed degree over a finite field. Let $\Delta \in \mathbb{N}$ be a prime power and $\mathcal{I}_\Delta(n) \subseteq \mathbb{F}_\Delta[x]$ the set of monic irreducible polynomials in $\mathbb{F}_\Delta[x]$ of degree n, and $N_\Delta(n) = \#\mathcal{I}_\Delta(n)$. By Lidl & Niederreiter (1983), Exercises 3.27 and 3.28,

$$\frac{\Delta^n}{n} - \frac{\Delta}{\Delta - 1} \cdot \frac{\Delta^{n/2} - 1}{n} \le N_\Delta(n) \le \frac{\Delta^n - \Delta}{n}. \tag{5.1}$$

First, consider the case when $\mathfrak{A}$ is a field extension of K.

THEOREM 5.2. *Let $\mathfrak{A}$ be field of dimension ν over K. The algorithm* `FindZeroDivisor` *with input $\mathfrak{A}$ reports that $\mathfrak{A}$ is a field with probability at least $1/4$ and reports "failure" with probability at most $3/4$.*

PROOF. An element $b \in \mathbb{N}$ always has a minimal polynomial of degree dividing ν, since $\mathsf{K}[b]$ is a subfield of $\mathfrak{A}$, and $[\mathsf{K}[b] : \mathsf{K}]$ divides $[\mathfrak{A} : \mathsf{K}]$. The elements $b \in \mathfrak{A}$ such that $\deg \min_{\mathsf{K}}(b) = \nu$ are exactly those that satisfy an irreducible polynomial in $\mathsf{K}[x]$ of degree ν, and ν distinct elements of $\mathfrak{A}$ satisfy each such polynomial. Applying (5.1),

$$N_q(\nu) \ge \frac{q^\nu}{\nu} - \frac{q}{q-1} \cdot \frac{q^{\nu/2} - 1}{\nu} = \frac{q^\nu}{2\nu} \cdot \left(2 - 2 \cdot \frac{q^{\nu/2} - 1}{(q-1)q^{\nu-1}}\right) \ge \frac{q^\nu}{2\nu},$$

where it is easily verified that $(q^{\nu/2} - 1)/((q-1)q^{\nu-1}) \le 1/2$. The number of $b \in \mathfrak{A}$ satisfying irreducible polynomials in $\mathsf{K}[x]$ of degree ν is $\nu N_q(\nu) \ge \nu \cdot q^\nu/(2\nu) = q^\nu/2$. Thus, the number of $b \in \mathfrak{A}$ with $\deg \min_{\mathsf{K}}(b) < \nu$ must be less than $q^\nu/2$, so $\varrho \le 1/4$ since we choose two elements $b \in \mathfrak{A}$ independently and test each of them. $\square$

Now let $\mathfrak{A}$ be an algebra with at least one non-trivial zero divisor, i.e., $\mathfrak{A}$ is not a local algebra. We call an element $b \in \mathfrak{A}$ *reducible* over K if its minimal polynomial in $\mathsf{K}[x]$ is reducible, and *irreducible* over K otherwise. Define

$$\Lambda(\mathfrak{A}) = \#\{b \in \mathfrak{A} \mid b \text{ is irreducible over } \mathsf{K}\},$$

The failure probability of `FindZeroDivisor` is at most $(\Lambda(\mathfrak{A})/q^\nu)^2$, ignoring for now the possibility that $b = b_1b_2 - b_2b_1$ yields a zero divisor of $\mathfrak{A}$ (this is only used when $\mathfrak{A}$ is a local algebra over $\mathfrak{A}$, that is, when $\mathfrak{A}/\mathrm{rad}(\mathfrak{A})$ is a (finite) field, where $\mathrm{rad}(\mathfrak{A})$ is the Jacobson radical of $\mathfrak{A}$ – see Theorem 5.9 below).

THEOREM 5.3. *Let $\mathfrak{A}$ be an algebra of dimension ν over K which is not local and possess a non-trivial zero divisor. The probability ϱ that the algorithm* `FindZeroDivisor` *fails to find a non-trivial zero divisor in $\mathfrak{A}$ is at most $8/9$.*

We begin by proving two theorems dealing with significant special cases. In Theorem 5.7 we bound ϱ for $\mathfrak{A}$ simple, and in Theorem 5.8 for $\mathfrak{A}$ semi-simple.

For now, assume that $\mathfrak{A}$ is simple and therefore isomorphic to a full matrix algebra $\mathsf{E}^{r\times r}$ of all $r \times r$ matrices over some algebraic extension field $\mathsf{E} \supseteq \mathsf{K}$ (see Lang 1984, Chapter 17). We set $\mu = [\mathsf{E} : \mathsf{K}]$, so $\mathsf{E} \cong \mathbb{F}_{q^\mu}$ and $\nu = \mu r^2$.

Let $a \in \mathfrak{A}$ and $B \in \mathsf{E}^{r\times r}$ its image in $\mathsf{E}^{r\times r}$. The minimal polynomial in $\mathsf{K}[x]$ of $a \in \mathfrak{A}$ is the monic $f \in \mathsf{K}[x] \setminus \{0\}$ of minimal degree such that $f(B) = 0$. This minimal polynomial is intimately related to the minimal polynomial $g \in \mathsf{E}[x]$ of the matrix B: f is the monic polynomial of smallest degree in $\mathsf{K}[x] \setminus \{0\}$ such that $g \mid f$. We write $f = \min_{\mathsf{K}}(B) \in \mathsf{K}[x]$ and $g = \min_{\mathsf{E}}(B) \in \mathsf{E}[x]$. Theorem 5.7 is proved by showing that for at most a constant fraction of matrices $B \in \mathsf{E}^{r\times r}$ that $\min_{\mathsf{K}}(B)$ is irreducible in $\mathsf{K}[x]$. First, we show that every matrix in $\mathsf{E}^{r\times r}$ similar to a companion matrix is similar to at least $q^{\mu r^2-\mu r}(1-q^{-\mu})^{q^\mu/(q^\mu-1)}$ distinct matrices. This implies that a substantial fraction of all matrices in $\mathsf{E}^{r\times r}$ are similar to companion matrices in $\mathsf{E}^{r\times r}$. We then show that for most matrices similar to companion matrices that their minimal polynomials in $\mathsf{K}[x]$ are products of at least two distinct irreducible factors.

LEMMA 5.4. *Let $r \geq 2$ and $B \in \mathsf{E}^{r\times r}$ be such that $g = \min_{\mathsf{E}}(B)$ and $\deg g = r$. Then B is similar to at least $q^{\mu r^2-\mu r} \cdot (1 - q^{-\mu})^{q^\mu/(q^\mu-1)}$ distinct matrices in $\mathsf{E}^{r\times r}$, exactly one of which is a companion matrix.*

PROOF. Matrices in $\mathsf{E}^{r\times r}$ whose minimal polynomials in $\mathsf{E}[x]$ have degree r are exactly those similar to the companion matrix of their minimal polynomial. Since the minimal polynomial is the only invariant factor if it has degree r, it completely characterizes the similarity class. Since no two distinct companion matrices are similar, we know that B is similar to exactly one companion matrix.

It is well known (see, for example, Hodges 1958) that the number of matrices similar to a given matrix $B \in \mathsf{E}^{r\times r}$ is the total number $\mathcal{L}(\mathsf{E}, r)$ of non-singular matrices in $\mathsf{E}^{r\times r}$ divided by the number of non-singular matrices in $\mathsf{E}^{r\times r}$ which commute with B. In the case of a $B \in \mathsf{E}^{r\times r}$ with $\deg \min_{\mathsf{E}}(B) = r$, it is shown by Gantmacher (1990), Section 8.2, that the only matrices commuting with B are in $\mathsf{E}[B]$, whence there are $q^{\mu r}$ of them. From (Dickson 1901, Part II, Chapter 1), we have

$$\mathcal{L}(\mathsf{E}, r) = \prod_{0\leq i<r} (q^{\mu r} - q^{\mu i}) = q^{\mu r^2} \prod_{1\leq i\leq r} (1 - 1/q^{\mu i}).$$

We bound $\prod_{1\leq i\leq r}(1 - q^{-\mu i})$ from below by considering its logarithm

$$\begin{aligned}\log \prod_{1\leq i\leq r} (1 - q^{-\mu i}) &= -\sum_{1\leq i\leq r}\sum_{j\geq 1} \frac{1}{jq^{\mu ij}} = -\sum_{j\geq 1} \frac{1}{j} \cdot \frac{1}{q^{\mu j}-1}\left(1 - \frac{1}{q^{\mu jr}}\right) \\ &\geq -\sum_{j\geq 1} \frac{1}{jq^{\mu j}} \frac{q^{\mu j}}{q^{\mu j}-1} \geq \log\left(1 - \frac{1}{q^\mu}\right)^{q^\mu/(q^\mu-1)}.\end{aligned}$$

□

LEMMA 5.5. *The number of monic $g \in \mathsf{E}[x]$ of degree $r \geq 2$ such that the $f \in \mathsf{K}[x] \setminus \{0\}$ of smallest degree with $g \mid f$ is irreducible in $\mathsf{K}[x]$, is less than $(3/4) \cdot q^{\mu r}$.*

PROOF. The proof is broken into two parts: when $r \geq 3$ and when $r = 2$.

Assume $r \geq 3$. We prove that the number of monic $g \in \mathsf{E}[x]$ of degree $r \geq 3$ such that $g = g_1 g_2$ where g_1 is monic, irreducible and $r/2 < \deg g_1 < r$ is greater than $q^{\mu r}/4$. Since any $g \in \mathsf{E}[x]$ has at most one such factor g_1, f is reducible for such g (f has roots in two distinct extension fields of E). The exact number of such g is

$$\begin{aligned}
q^{\mu r - \mu l} \sum_{r/2<l<r} N_{q^\mu}(l) &> q^{\mu r} \sum_{r/2<l<r} \left(\frac{1}{l} - \frac{q^\mu}{q^\mu - 1} \cdot \frac{1}{l \cdot q^{\mu l/2}} \right) \\
&\geq q^{\mu r} \cdot \left(\sum_{r/2<l<r} \frac{1}{l} - \frac{2}{r} \cdot \frac{q^\mu}{q^\mu - 1} \cdot \sum_{r/2<l<r} \frac{1}{q^{\mu l/2}} \right) \\
&\geq q^{\mu r} \cdot \left(\log(2) + \log\left(\frac{r-1}{r}\right) + \frac{1}{r-1} - \frac{2}{r} \right. \\
&\qquad \left. - \frac{2}{r} \cdot \frac{q^\mu}{q^\mu - 1} \cdot \frac{q^{\mu/2}}{q^{\mu/2} - 1} \cdot \frac{q^{\mu r/4 - \mu/2} - 1}{q^{\mu r/2}} \right) \\
&\geq q^{\mu r}/4,
\end{aligned}$$

except possibly when $3 \leq r \leq 14$. We use the Euler summation formula and (5.1). The lemma is easily verified when $3 \leq r \leq 14$ by explicitly expanding $q^{\mu r - \mu l} \sum_{r/2<l<r} N_{q^\mu}(l)$ as a polynomial in q^μ.

When $r = 2$ the above approach does not work since no such factors g_1 exist in the desired degree range. In this case, g is either irreducible in $\mathsf{E}[x]$ or it factors into two linear factors. When g is irreducible then f is irreducible and there are $N_{q^\mu}(2)$ such g. Suppose g factors, and each of these factors divides an irreducible $f \in \mathsf{K}[x]$. It must be the case that f factors completely in $\mathsf{E}[x]$, thus $s = \deg f$ divides μ. Moreover, for each irreducible $f \in \mathsf{K}[x]$ of degree s, there are $\binom{s}{2}$ distinct ways of choosing 2 factors of f in $\mathsf{E}[x]$ to form a $g \in \mathsf{E}[x]$ with this (irreducible) minimal degree multiple $f \in \mathsf{K}[x]$. Thus, the total number of $g \in \mathsf{E}[x]$ of degree 2, such that the minimal degree $f \in \mathsf{K}[x]$ with $g \mid f$ is irreducible in $\mathsf{K}[x]$ is

$$\begin{aligned}
N_{q^\mu}(2) + \sum_{s \mid \mu} N_q(s) \cdot \frac{s(s-1)}{2} &\leq \frac{q^{2\mu}}{2} + \sum_{s \mid \mu} \frac{q^s(s-1)}{2} \\
= \frac{q^{2\mu}}{2} + \frac{q^\mu(\mu-1)}{2} + \sum_{\substack{s \mid \mu \\ 1 \leq s \leq \mu/2}} \frac{q^s(s-1)}{2} &\leq \frac{q^{2\mu}}{2} + \frac{q^\mu(\mu-1)}{2} + \sum_{1 \leq s \leq \mu/2} \frac{q^s(s-1)}{2} \\
= \frac{q^{2\mu}}{2} + \frac{q^\mu(\mu-1)}{2} + \frac{q}{(q-1)^2} \cdot \frac{q^{\mu/2+1}\mu - 2q^{\mu/2+1} - q^{\mu/2}\mu + 2q}{4} &\leq \frac{3q^{2\mu}}{4},
\end{aligned}$$

using (5.1) and elementary calculus. □

We combine Lemmas 5.4 and 5.5 to count the number of matrices in $\mathsf{E}^{r \times r}$ whose minimal polynomials in $\mathsf{K}[x]$ are irreducible in $\mathsf{K}[x]$.

Theorem 5.6. *The number of matrices in $\mathsf{E}^{r \times r}$ whose minimal polynomials in $\mathsf{K}[x]$ are reducible is greater than $q^{\mu r^2}/16$.*

Proof. By Lemma 5.5 there are at most $3q^{\mu r}/4$ polynomials $g \in \mathsf{E}[x]$ of degree r such that the $f \in \mathsf{K}[x] \setminus \{0\}$ of least degree with $g \mid f$ is irreducible in $\mathsf{K}[x]$. Hence there are greater than $q^{\mu r}/4$ such that f is reducible. A matrix in $\mathsf{E}^{r\times r}$ similar to a companion matrix of such an g will have a reducible minimal polynomial in $\mathsf{K}[x]$. By Lemma 5.4 a companion matrix $C_g \in \mathsf{E}^{r\times r}$ of such a g is similar to at least $q^{\mu r^2 - \mu r}(1-1/q^\mu)^{q^\mu/(q^\mu-1)}$ distinct matrices in $\mathsf{E}^{n\times n}$, of which C_g is the only companion matrix. Thus there are at least

$$\frac{q^{\mu r^2}}{4} \cdot \left(1 - \frac{1}{q^\mu}\right)^{q^\mu/(q^\mu-1)} \geq 1/16$$

matrices in $\mathsf{E}^{r\times r}$ with reducible minimal polynomials in $\mathsf{K}[x]$. $\square$

Theorem 5.7. *If $\mathfrak{A}$ is a simple algebra of dimension ν over K which is not a field. Then $\Lambda(\mathfrak{A}) \leq (15/16) \cdot q^\nu$.*

Proof. The number of irreducible $a \in \mathfrak{A}$ is equal to the number of $B \in \mathsf{E}^{r\times r}$ with $\min_{\mathsf{K}}(B) \in \mathsf{K}[x]$ irreducible in $\mathsf{K}[x]$. By Theorem 5.6 this is at most $15/16 \cdot q^{\mu r^2}$. $\square$

Theorem 5.8. *If $\mathfrak{A}$ is a semi-simple algebra of dimension ν over K which is not local, then $\Lambda(\mathfrak{A}) \leq 15/16 \cdot q^\nu$.*

Proof. If $\mathfrak{A}$ is simple then Theorem 5.7 implies that this theorem is true, so assume $\mathfrak{A}$ is not simple. The Wedderburn Decomposition Theorem (see Pierce 1982, Section 3.5) yields a decomposition of $\mathfrak{A}$ as

$$\mathfrak{A} \cong \mathfrak{A}_1 \oplus \mathfrak{A}_2 \oplus \ldots \oplus \mathfrak{A}_k,$$

where $\mathfrak{A}_i$ is a simple algebra of dimension $\nu_i \geq 1$ for $1 \leq i \leq k$. This is also a decomposition of $\mathfrak{A}$ as an K-vector space, so $\nu = \nu_1 + \nu_2 + \cdots + \nu_k$, with $k \geq 2$ since $\mathfrak{A}$ is not simple. Each simple component $\mathfrak{A}_i$ is isomorphic to $\mathsf{E}_i^{r_i\times r_i}$, where E_i is a finite extension field of K. Under this isomorphism, each $b \in \mathfrak{A}$ has an image $(b_1, \ldots, b_k) \in \mathfrak{A}_1 \oplus \mathfrak{A}_2 \oplus \ldots \oplus \mathfrak{A}_k$, and $\min_{\mathsf{K}}(b) = \operatorname{lcm}(\min_{\mathsf{K}}(b_1), \ldots, \min_{\mathsf{K}}(b_k))$. It is clear that $\Lambda(\mathfrak{A}) \leq \Lambda(\mathfrak{A}_1)\Lambda(\mathfrak{A}_2)\cdots\Lambda(\mathfrak{A}_k)$, since the minimal polynomial of an element of $\mathfrak{A}$ is a power of an irreducible only if the minimal polynomial of each of its components in $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$ is a power of an irreducible (in fact, each component must be a power of the *same* irreducible, which is not reflected in this inequality). We consider two cases in this proof:

Case (i): $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$ are all fields. Here $r_1 = r_2 = \ldots = r_k = 1$, and the minimal polynomials of all elements in $\mathfrak{A}$ will be squarefree. Hence, we need only consider the case when the minimal polynomial is irreducible. We consider $\mathfrak{A}_1 \cong \mathbb{F}_{q^{\nu_1}}$ and $\mathfrak{A}_2 \cong \mathbb{F}_{q^{\nu_2}}$, and show that at most half the elements of $\mathfrak{A}_1 \oplus \mathfrak{A}_2$ have an irreducible minimal polynomial.

Start by determining the number of elements of $\mathfrak{A}_1 \oplus \mathfrak{A}_2$ annihilated by a single irreducible polynomial in $\mathsf{K}[x]$. If $b \in \mathfrak{A}$ and $\min_{\mathsf{K}}(b)$ is irreducible over K, then $\min_{\mathsf{K}}(b_1) = \min_{\mathsf{K}}(b_2)$, since the minimal polynomial of a field element is always irreducible. Thus $\mathsf{K}(b_1)$ and $\mathsf{K}(b_2)$ are isomorphic as fields, and $d = \deg \min_{\mathsf{K}}(b)$ divides both ν_1 and ν_2 since $\mathsf{K}(b_1) \subseteq \mathfrak{A}_1$ and $\mathsf{K}(b_2) \subseteq \mathfrak{A}_2$. In particular $d \mid \gcd(\nu_1, \nu_2)$. Any one irreducible $f \in \mathsf{K}[x]$ of degree d has d^2 roots in $\mathfrak{A}_1 \oplus \mathfrak{A}_2$.
We now count the number of elements $a \in \mathfrak{A}_1 \oplus \mathfrak{A}_2$ for which there exists a monic irreducible $f \in \mathsf{K}[x]$ with $f(a) = 0$. If $d \in \mathbb{N}$ divides $\gcd(\nu_1, \nu_2)$, exactly $d^2 N_q(d)$

elements in $\mathfrak{A}_1 \oplus \mathfrak{A}_2$ are annihilated by polynomials in $\mathcal{I}_{\mathsf{K}}(d)$. On the other hand, if $d \nmid \gcd(\nu_1, \nu_2)$ then no element of $\mathfrak{A}_1 \oplus \mathfrak{A}_2$ is annihilated by an irreducible polynomial of degree d in $\mathsf{K}[x]$. Making use of the fact that $\sum_{d \mid t} dN_q(d) = q^t$ for any $t \geq 1$ (see Lidl & Niederreiter 1983 Corollary 3.21),

$$\begin{aligned}\Lambda(\mathfrak{A}_1 \oplus \mathfrak{A}_2) = &\sum_{d \mid \gcd(\nu_1,\nu_2)} d^2 N_q(d) \leq \min(\nu_1, \nu_2) \sum_{d \mid \min(\nu_1,\nu_2)} dN_q(d) \\ &\leq \min(\nu_1,\nu_2) q^{\min(\nu_1,\nu_2)} \leq 15/16 \cdot q^{2\min(\nu_1,\nu_2)} \leq 15/16 \cdot q^{\nu_1+\nu_2},\end{aligned}$$

using the easily proven fact that $z \leq 15/16 \cdot q^z$ for all $z \geq 1$. Finally,

$$\begin{aligned}\Lambda(\mathfrak{A}) &\leq \Lambda(\mathfrak{A}_1 \oplus \mathfrak{A}_2)\Lambda(\mathfrak{A}_3 \oplus \ldots \oplus \mathfrak{A}_k) \leq \Lambda(\mathfrak{A}_1 \oplus \mathfrak{A}_2) q^{\nu-\nu_1-\nu_2} \\ &\leq 15/16 \cdot (q^{\nu_1+\nu_2}) \cdot q^{\nu-\nu_1-\nu_2} = 15/16 \cdot q^{\nu}.\end{aligned}$$

Case (ii): At least one of $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$ is not a field. Without loss of generality we can assume that $\mathfrak{A}_1$ is not a field. Thus $\mathfrak{A}_1 \cong \mathsf{E}_1^{r_1 \times r_1}$, where $r_1 \geq 2$, and E_1 is an extension field of K. By Theorem 5.7, $\Lambda(\mathfrak{A}_1) \leq 15/16 \cdot q^{\nu_1}$, and

$$\Lambda(\mathfrak{A}) \leq \Lambda(\mathfrak{A}_1)\Lambda(\mathfrak{A}_2)\cdots\Lambda(\mathfrak{A}_k) \leq \Lambda(\mathfrak{A}_1) q^{\nu-\nu_1} \leq 15/16 \cdot q^{\nu} \cdot q^{\nu-\nu_1} = 15/16 \cdot q^{\nu}.$$

□

The proof of Theorem 5.3 is completed by showing its validity when $\mathfrak{A}$ is not semi-simple or local.

Proof. [of Theorem 5.3] If $\mathfrak{A}$ is semi-simple, Theorem 5.8 implies that the theorem is true, so assume $\mathfrak{A}$ is not semi-simple, i.e., $\mathrm{rad}\mathfrak{A} \neq \{0\}$.

The Jacobson radical of $\mathfrak{A}$ is a nilpotent subalgebra of $\mathfrak{A}$; that is, for all $c \in \mathrm{rad}\mathfrak{A}$ the minimal polynomial of c in $\mathsf{K}[x]$ is a power of x. By the Wedderburn-Malcev Principal Theorem (see McDonald 1974, Theorem 8.28) $\mathfrak{A} = \mathfrak{S} + \mathrm{rad}\mathfrak{A}$, where $\mathfrak{S} \cong \mathfrak{A}/\mathrm{rad}\mathfrak{A}$ is a semi-simple subalgebra of $\mathfrak{A}$, and $\mathfrak{S} \cap \mathrm{rad}\mathfrak{A} = \{0\}$. Thus, every $a \in \mathfrak{A}$ can be written uniquely as $a = b + c$, where $b \in \mathfrak{S}$ and $c \in \mathrm{rad}\mathfrak{A}$. If a has minimal polynomial $f \in \mathsf{K}[x]$, observe that

$$0 = f(a) = f(b+c) = f(b) + [f(b+c) - f(b)],$$

and since every term in the expansion of $(b+c)^i - b^i$ for $i \geq 1$ contains a positive power of c, so too does every term in the expansion of $f(b+c) - f(b)$. Since $\mathrm{rad}\mathfrak{A}$ is an ideal in $\mathfrak{A}$ and $c \in \mathrm{rad}\mathfrak{A}$, $f(b+c) - f(b) \in \mathrm{rad}\mathfrak{A}$. Thus $f(b) = 0$.

If $\mathfrak{S}$ has dimension $\tau > 0$ over K, and is not a field, then the number of elements of $\mathfrak{S}$ whose minimal polynomial over K is irreducible is at most $15/16 \cdot q^{\tau}$ by Theorem 5.8. The minimal polynomial of an $a \in \mathfrak{A}$ is irreducible in $\mathsf{K}[x]$ only if the minimal polynomial of its component in $\mathfrak{S}$ is as well, whence

$$\Lambda(\mathfrak{A}) \leq 15/16 \cdot q^{\tau} \cdot q^{\nu-\tau} = 15/16 \cdot q^{\nu}.$$

The probability ϱ of failure is then at most $(15/16)^2 < 8/9$. □

Now consider the case when $\mathfrak{A}$ is a local algebra over K, of dimension ν.

THEOREM 5.9. *Let $\mathfrak{A}$ be a local algebra of dimension ν over K which is not a field. If $\mathfrak{A}$ is commutative then the probability ϱ that the algorithm* `FindZeroDivisor` *fails is less than* $1/4$*. If $\mathfrak{A}$ is non-commutative then* $\varrho \leq 3/4$.

PROOF. First we prove $\Lambda(\mathfrak{A}) \leq q^\nu/2$ when $\mathfrak{A}$ is commutative. The Wedderburn-Malcev Principal Theorem gives a decomposition $\mathfrak{A} = \mathfrak{S} + \mathrm{rad}(\mathfrak{A})$, where $\mathfrak{S}$ is a subalgebra of $\mathfrak{A}$ isomorphic to $\mathfrak{A}/\mathrm{rad}(\mathfrak{A})$, and $\mathfrak{S} \cap \mathrm{rad}(\mathfrak{A}) = \{0\}$. Since $\mathfrak{A}$ is local, $\mathfrak{S}$ is a finite field of dimension τ over K, say $\mathfrak{S} \cong \mathbb{F}_{q^\tau}$ for some $\tau \geq 1$. The algebra $\mathfrak{A}$ possesses a non-trivial zero divisor, so it is not a field and $\mathrm{rad}(\mathfrak{A}) \neq \{0\}$. Let $k > 1$ be the nullity of $\mathrm{rad}(\mathfrak{A})$, the smallest integer k such that $\mathrm{rad}(\mathfrak{A})^k = \{0\}$.

For any $a \in \mathfrak{A}$, suppose $f = \min_{\mathsf{K}}(a) \in \mathsf{K}[x]$ is irreducible of degree n. We prove this implies $a \in \mathfrak{S}$. Consider the subalgebra $\mathsf{K}[a] \subsetneq \mathfrak{A}$. The minimal polynomial of a is irreducible so $\mathsf{K}[a] \cong \mathbb{F}_{q^n}$ is a field extension of K, and by Fermat's Little Theorem $a^{q^n} = a$. Let

$$t = \min_{t_0 > 0} \left\{ q^{t_0} \geq k \text{ and } t_0 \equiv 0 \bmod n \right\},$$

and consider the K-linear ring morphism $\phi : \mathfrak{A} \to \mathfrak{A}$ defined by $\phi(z) = z^{q^t}$ for any $z \in \mathfrak{A}$. Since $t \equiv 0 \bmod n$, we know $\phi(a) = a^{q^t} = a^{q^n} = a$. Suppose $a = b + c$ with $b \in \mathfrak{S}$ and $c \in \mathrm{rad}(\mathfrak{A})$. Then

$$a = \phi(a) = \phi(b + c) = \phi(b) + \phi(c) = b^{q^t} + c^{q^t} = b^{q^t},$$

since q^t is greater than the nullity of $\mathrm{rad}(\mathfrak{A})$ so $c^{q^t} = 0$. This implies $a \in \mathfrak{S}$ since $a = b^{q^t}$ and $\mathfrak{S}$ is a subalgebra of $\mathfrak{A}$. Thus, the only elements of $\mathfrak{A}$ with irreducible polynomials are those in $\mathfrak{S}$. Since $\mathfrak{S} \subsetneq \mathfrak{A}$ it follows that $\Lambda(\mathfrak{A}) \leq q^{\nu-1} \leq q^\nu/2$.

When $\mathfrak{A}$ is non-commutative we note $b = a_1a_2 - a_2a_1 \in \mathrm{rad}(\mathfrak{A})$ since $\mathfrak{A}/\mathrm{rad}(\mathfrak{A})$ is a (commutative) field. Hence the minimal polynomial of b is x^i for some $i \geq 1$, and $i \geq 2$ if $b \neq 0$. We must show that at least $q^{2\nu}/4$ pairs $a_1, a_2 \in \mathfrak{A}$ satisfy $a_1a_2 - a_2a_1 \neq 0$. The centre $\mathfrak{C}$ of $\mathfrak{A}$ is a subalgebra of $\mathfrak{A}$ with at most $q^{\nu-1}$ elements since $\mathfrak{A}$ is non-commutative. For every $a_1 \in \mathfrak{A} \setminus \mathfrak{C}$, the nullspace of the linear map $\varphi : \mathfrak{A} \to \mathfrak{A}$ defined by $\varphi(x) = a_1x - xa_1$ has at most $q^{\nu-1}$ elements since $a_1 \neq \mathfrak{C}$. Thus there are at least $(q^\nu - q^{\nu-1})^2$ pairs $a_1, a_2 \in \mathfrak{A}$ with $a_1a_2 - a_2a_1 \neq 0$, and $(q^\nu - q^{\nu-1})^2/q^{2\nu} \geq 1/4$, so the probability ϱ of the algorithm failing is at most $3/4$. □

For any algebra $\mathfrak{A}$, the failure probability ϱ of the algorithm `FindZeroDivisor` is bounded by $\varrho \leq 8/9$, using Theorem 5.2 when $\mathfrak{A}$ is a field and Theorems 5.3 and 5.9 when it is not. This yields the following corollary to Theorem 5.1.

COROLLARY 5.10. (TO THEOREM 5.1) *Let $\mathfrak{A}$ be an algebra of dimension ν over $\mathsf{K} = \mathbb{F}_q$. The algorithm* `FindZeroDivisor` *requires $O(\nu\chi + \mathrm{MM}(\nu) + \mathrm{M}(\nu)\log\nu\log q)$ operations in K to determine whether $\mathfrak{A}$ is a field extension of K, or to produce $b_1, b_2 \in \mathfrak{A} \setminus \{0\}$ with $b_1b_2 = 0$ (where χ is the number of operations in K required for a single multiplication in $\mathfrak{A}$), or to fail with probability at most* $8/9$.

6. Application to the Functional Decomposition of Polynomials

The problem of functionally decomposing polynomials has received considerable attention recently, and there exist a number of classes of polynomials for which no polynomial-time solution has been found. We consider such a class — the linearized or additive polynomials — and show that it is isomorphic (in a computationally trivial way) to a skew-polynomial ring. This allows us to employ our algorithms for complete and bi-factorization in skew-polynomial rings to obtain very fast algorithms for the functional decomposition of linearized polynomials.

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

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

Now consider the skew-polynomial ring $\mathsf{F}[x;\psi]$, where $\psi(a)=a^p$ for any $a\in\mathsf{F}$ and $xa=a^px$ for all $a\in\mathsf{F}$. This skew-polynomial ring is isomorphic to the ring $\mathbb{A}_\mathsf{F}$ under the map $\Phi\colon\mathbb{A}_\mathsf{F}\to\mathsf{F}[x;\psi]$, which acts as the identity on F and sends λ^{p^i} to x^i for $i\ge 0$ (see McDonald (1974), Theorem 2.13). Note that if $f\in\mathbb{A}_\mathsf{F}$, then $\deg\Phi(f)=\log_p(\deg f)$, so this isomorphism removes some of the "sparseness" of linearized polynomials. Computationally, Φ just maps between two interpretations of the input, and is free of charge.

The functional decomposition problem for general polynomials in $\mathsf{F}[\lambda]$ comes in two flavours analogous to our complete factorization and bi-factorization problems for $\mathsf{F}[x;\sigma]$. Given a polynomial $f\in\mathsf{F}[\lambda]$ of degree N, the (functional) complete decomposition problem asks for functionally indecomposable $f_1,\ldots,f_k\in\mathsf{F}[\lambda]$ such that $f=f_1\circ\cdots\circ f_k$ (any $h\in\mathsf{F}[\lambda]\setminus\mathsf{F}$ is functionally indecomposable if all its bi-decompositions contain a linear composition factor). When $p\nmid\mathrm{N}$, the so-called "tame" case for complete decomposition, fast deterministic algorithms for complete decomposition are presented in von zur Gathen *et al.* (1987). When $p\,|\,\mathrm{N}$, the "wild" case, an algorithm of Zippel (1991) apparently solves the complete decomposition problem in time $(\deg f)^{O(1)}$, although the exact running time is not calculated. All polynomials $f\in\mathbb{A}_\mathsf{F}$ 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 $f\in\mathsf{F}[\lambda]$ and $\mathrm{S}\in\mathbb{N}$, the (functional) bi-decomposition problem asks if there exist $g,h\in\mathsf{F}[\lambda]$ such that $f=g\circ h$ and $\deg h=\mathrm{S}$, and if so, find such g, h. The tame case, when $p\nmid(\mathrm{N}/\mathrm{S})$, is solved efficiently in von zur Gathen *et al.* (1987). When $p\,|\,(\mathrm{N}/\mathrm{S})$, 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 (1990b). All non-trivial bi-decompositions of linearized polynomials fall into the wild case, since, if $f\in\mathbb{A}_\mathsf{F}$ and $f=g\circ h$ for $g,h\in\mathsf{F}[\lambda]$, then Dorey & Whaples (1974) show that $\deg g=p^r$ for some $r\in\mathbb{N}$.

When $f\in\mathbb{A}_\mathsf{F}$, we can solve both the bi-decomposition and complete decomposition problems using our algorithms for complete factorization and bi-factorization in $\mathsf{F}[x;\psi]$. The key observation is that we need only consider decompositions of $f\in\mathbb{A}_\mathsf{F}$ into linearized polynomials: Dorey & Whaples (1974) show that if $f=\bar f_1\circ\cdots\circ\bar f_k$ for any $\bar f_1,\ldots,\bar f_k\in\mathsf{F}[\lambda]$, then there exist $f_1,\ldots,f_k\in\mathbb{A}_\mathsf{F}$ 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 $f\in\mathbb{A}_\mathsf{F}$ of degree p^n can be found by finding a complete factorization of $\Phi(f)$ in $\mathsf{F}[x;\psi]$. Similarly, the bi-decomposition problem on input $f\in\mathbb{A}_\mathsf{F}$ of degree p^n and $\mathrm{S}\in\mathbb{N}$, is equivalent to the bi-factorization problem in $\mathsf{F}[x;\psi]$ on inputs $\Phi(f)\in\mathsf{F}[x;\psi]$ and $\log_p\mathrm{S}$.

THEOREM 6.1. *Let $f\in\mathbb{A}_\mathsf{F}$ have degree $\mathrm{N}=p^n$ and $\mathrm{S}=p^s<\mathrm{N}$, where $\mathsf{F}=\mathbb{F}_{p^\omega}$ for some prime $p\in\mathbb{N}$ and $\omega\ge 1$. We can produce a complete decomposition of f in $\mathbb{A}_\mathsf{F}$, and determine if there exist $g,h\in\mathbb{A}_\mathsf{F}$ such that $\deg h=\mathrm{S}$ and $f=g\circ h$, and if so, find such*

g, h, with a deterministic algorithm requiring $(n\omega p)^{O(1)}$ operations in $\mathbb{F}_p$ or a probabilistic algorithm requiring $O(n^4\omega \mathrm{M}(\omega) + n^3\omega^2\mathrm{M}(\omega)\log\omega + n\mathrm{MM}(n\omega) + n\mathrm{M}(n\omega)\log(n\omega)\log p)$ operations in $\mathbb{F}_p$.

Acknowledgement

The author would like to thank the anonymous referees for their improvements to Lemma 2.1 and their many other helpful comments.

References

E. R. Berlekamp. Factoring polynomials over large finite fields. *Math. Comp.* **24**, pp. 713–735, 1970.
M. Bronstein and M. Petkovšek. On Ore rings, linear operators and factorisation. *Programmirovanie* **20**, pp. 27–45, 1994.
M. Bronstein and M. Petkovšek. An introduction to pseudo-linear algebra. *Theoretical Computer Science* **157**, pp. 3–33, 1996.
D. Cantor and E. Kaltofen. Fast multiplication of polynomials over arbitrary algebras. *Acta Informatica* **28**, pp. 693–701, 1991.
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.
L. E. Dickson. *Linear Groups with an Exposition of the Galois Field Theory.* Teubner, Leipzig, 1901; Dover, New York, 1958 (Leipzig), 1901. Dover, New York, 1958.
F. Dorey and G. Whaples. Prime and composite polynomials. *J. Algebra* **28**, pp. 88–101, 1974.
F. R. Gantmacher. *The Theory of Matrices, Vol. I.* Chelsea Publishing Co. (New York NY), 1990.
J. von zur Gathen. Functional decomposition of polynomials: the tame case. *J. Symb. Comp.* **9**, pp. 281–299, 1990a.
J. von zur Gathen. Functional decomposition of polynomials: the wild case. *J. Symb. Comp.* **10**, pp. 437–452, 1990b.
J. von zur Gathen and V. Shoup. Computing Frobenius maps and factoring polynomials. *Computational Complexity* **2**, pp. 187–224, 1992.
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.
M. Giesbrecht. Factoring in skew-polynomial rings. In *Proc. LATIN'92*, pp. 191–203, Sao Paulo, Brasil, 1992.
D. Yu. Grigoriev. Complexity of factoring and calculating GCD of linear differential operators. *J. Symb. Comp.* **10**, pp. 7–37, 1990.
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.
D. Kozen and S. Landau. Polynomial decomposition algorithms. *J. Symb. Comp.* **7**, pp. 445–456, 1989.
S. Lang. *Algebra.* Addison-Wesley (Reading MA), 1984.
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 Algebras.* 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.
A. Schönhage. Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2. *Acta Informatica* **7**, pp. 395–398, 1977.
A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing* **7**, pp. 281–292, 1971.
M. F. Singer. Testing reducibility of linear differential operators: A group theoretic perspective. *Applicable Algebra in Engineering, Communication and Computing* **7**(2), pp. 77–104, 1996.
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.