A Gröbner-Basis Theory for Divide-and-Conquer Recurrences

Frédéric Chyzak, Philippe Dumas

A Gröbner-Basis Theory for Divide-and-Conquer Recurrences

Frédéric Chyzak
frederic.chyzak@inria.fr
INRIA

Philippe Dumas
philippe.dumas@inria.fr
INRIA

ABSTRACT

We introduce a variety of noncommutative polynomials that represent divide-and-conquer recurrence systems. Our setting involves at the same time variables that behave like words in purely noncommutative algebras and variables governed by commutation rules like in skew polynomial rings. We then develop a Gröbner-basis theory for left ideals of such polynomials. Strikingly, the nature of commutations generally prevents the leading monomial of a polynomial product to be the product of the leading monomials. To overcome the difficulty, we consider a specific monomial ordering, together with a restriction to monic divisors in intermediate steps. After obtaining an analogue of Buchberger's algorithm, we develop a variant of the F_4 algorithm, whose speed we compare.

CCS CONCEPTS

• **Computing methodologies** → **Algebraic algorithms**.

KEYWORDS

Gröbner bases, divide-and-conquer recurrences, skew polynomials

ACM Reference Format:
Frédéric Chyzak and Philippe Dumas. 2020. A Gröbner-Basis Theory for Divide-and-Conquer Recurrences. In *International Symposium on Symbolic and Algebraic Computation (ISSAC '20), July 20–23, 2020, Kalamata, Greece.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

Divide-and-conquer recurrences appear at the interface between mathematics and (theoretical) computer science, namely, in relation to number systems, formal languages, number theory, and complexity theory. For example, the total number u_n of operations $(+, -, \times)$ in Karatsuba's algorithm when multiplying polynomials of degree less than n satisfies $u_0 = 0$, $u_1 = 1$, and the system of recurrences

$$u_{2n+2} - 3u_{n+1} = 8n+4, \quad u_{2n+3} - 2u_{n+2} - u_{n+1} = 8n+8 \quad \text{for } n \geq 0.$$

So far, the literature has focused almost exclusively on finding the asymptotic behavior of some sequence defined by first-order recurrences; see the references in [5, 9]. In the example above, the sequence undoubtedly exists and is defined uniquely. But can we guarantee that any given divide-and-conquer system actually defines a sequence, and this uniquely? This motivates an algorithmic study of suitable left ideals that encode divide-and-conquer systems.

In Section 1, we explain how divide-and-conquer recurrences can be expressed as polynomials of a noncommutative algebra. In Section 2, we develop a Gröbner-basis theory in it, by using a specific monomial ordering that we call *breadth-first ordering*. This leads to a Buchberger algorithm whose correctness we prove by an analogue of the usual criterion on S-polynomials. We then replace the pair-completion approach by a linear-algebraic one in Section 3, and we develop a variant of the algorithm F_4. Timings are briefly presented in Section 4, together with a speed comparison.

We close this introduction with a short review of related works on Gröbner bases, which we hope the reader will keep in mind and contrast to our contribution. The Gröbner-basis theory for commutative polynomial algebras $k[x_1, \ldots, x_n]$ over a given field k is well understood, see textbooks like [2, 4]. The theory has since long been studied in relation to linear algebra [13]. This led to developments like Faugère's algorithm F_4 [6], a big algorithmic speed-up. Extensions to noncommutative contexts range between two extremes. A first line of research is towards free noncommutative algebras $k\langle a_1, \ldots, a_n\rangle$ [16–18] and path algebras $k\Gamma$ [8, 20], replacing commutative monomials by words on letters a_i or by paths on a graph Γ. Noetherianity is typically lost, but algorithms have been given both for one-sided and two-sided ideals. In these contexts, monomials commute with the coefficients from k and the one-sided case is regarded to be simpler than the two-sided. Another line of research concerns k-algebras given by generators and relations, for well-identified forms of relations. Early works in this direction provided algorithms for Weyl algebras, $k\langle x_1, \ldots, x_n, y_1, \ldots, y_n;\ y_j x_i = x_i y_j + \delta_{i,j},\ y_j y_i = y_i y_j,\ x_j x_i = x_i x_j,\ 1 \leq i, j \leq n\rangle$, where $\delta_{i,j}$ is 1 if $i = j$ and 0 otherwise [7], and for enveloping algebras of Lie algebras, that is, given a finite-dimensional Lie algebra $\mathfrak{g}$ with k-basis $(a_1, \ldots, a_n)$, the associative algebra $k\langle x_1, \ldots, x_n;\ x_j x_i = x_i x_j + [a_j, a_i],\ 1 \leq i, j \leq n\rangle$ [1]. These studies focus on one-sided ideals, which is natural for [7] as Weyl algebras have no nontrivial two-sided ideals. They were extended to noncommutative polynomial rings of solvable type $k\langle x_1, \ldots, x_n;\ x_j x_i = c_{i,j} x_i x_j + p_{i,j},\ 1 \leq i, j \leq n\rangle$, where the $c_{i,j}$ are nonzero and the $p_{i,j}$ are polynomials smaller than $x_i x_j$ in a suitable sense [10, 14]. In all such algebras given by generators and relations, again, the monomials commute with the coefficients from k. In contrast, the rings of difference-differential operators over rational-function coefficients [19] can be obtained by tensoring Weyl algebras or similar algebras with the field $F = k(x_1, \ldots, x_n)$: they involve variables y_i that commute with one another but generally not with the coefficients from the field F. A similar situation occurs with Ore algebras [3], which are generalizations to more types of linear functional operators. A generalization of [10] to a sort of solvable polynomials rings whose monomials in the x_i need not commute with the coefficients from k was developed in [11]. All these cases are (left, right, two-sided) Noetherian rings. For an integer $b \geq 2$, the algebra $A := k\langle x, y; yx = x^b y\rangle$ of linear b-Mahler operators with polynomial coefficients directly relates to the algebras of section operators discussed in the present article; see our Conclusion. Its analogue with rational-function coefficients, $k(x) \otimes_{k[x]} A$, is a case of Ore algebras, while the algebra A itself

Preprint, to appear in *ISSAC '20, July 20–23, 2020, Kalamata, Greece*

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

is non-Noetherian. The theory was adapted to A so as to provide computable Gröbner bases for finitely generated one-sided and two-sided ideals [21]. The setting to be introduced in Section 1 inherits from both the free noncommutative algebras $k\langle a_1, \dots, a_n\rangle$ and Ore algebras by considering skew polynomials whose monomials are words that have commutation rules with their coefficients in a field $k(x)$. It is non-Noetherian. A similar situation, but distinct in that the monomials are not just noncommutative words but satisfy commutation rules as well, was introduced in an application to the calculation of symmetries of discrete systems [15]; see the generalization [12].

Acknowledgement. Supported in part by ANR-19-CE40-0018.

1 SKEW POLYNOMIALS

In this work, k is a commutative, computable field. The sequences we have in mind are defined on the set of nonnegative integers $\mathbb{Z}_{\geq 0}$. We also see them as those sequences defined on $\mathbb{Z}$ that have their supports in $\mathbb{Z}_{\geq 0}$. To a sequence $(u_n)_{n\in\mathbb{Z}_{\geq 0}}$, we associate the formal power series $\sum_{n\geq 0} u_n x^n$ in $k[[x]]$. The ring $k[[x]]$ is a subring of the field of formal Laurent series $k((x))$, which proves to be the right algebraic set to think of our sequences.

1.1 Section operators

To formalize the study of divide-and-conquer recurrences we introduce what we call *section operators*. We fix an integer $b \geq 2$, which the reader can think of as the radix of a numeration system. For each integer $0 \leq r < b$, we consider, with the same notation, the operators $S_{b,r}$ that act k-linearly on sequences in $k^{\mathbb{Z}}$, and, respectively, on formal Laurent series in $k((x))$ by

$$S_{b,r} \cdot u_n = u_{bn+r}, \qquad S_{b,r} \cdot \sum_n u_n x^n = \sum_n u_{bn+r} x^n, \tag{1}$$

where, in each case, n ranges in $\mathbb{Z}$.

The operators $S_{b,r}$ given by $0 \leq r < b$ generate a monoid of endomorphisms, which, by extension of the notation, are the $S_{b^\ell,r}$ obtained for all integers $\ell \geq 1$ and $0 \leq r < b^\ell$, and are related by the composition rule

$$S_{b^\ell,r} S_{b^{\ell'},r'} = S_{b^{\ell+\ell'},b^{\ell'}r+r'}. \tag{2}$$

Moreover, for any $\ell \geq 1$ and for each $0 \leq r < b^\ell$, there is a 'Leibniz' formula: for any two formal Laurent series $f(x)$ and $g(x)$,

$$\begin{aligned} S_{b^\ell,r} \cdot (f(x)g(x)) = &\sum_{s=0}^{r} \left(S_{b^\ell,r-s} \cdot f(x)\right) \left(S_{b^\ell,s} \cdot g(x)\right) \\ &+ x \sum_{s=r+1}^{b^\ell-1} \left(S_{b^\ell,r-s+b^\ell} \cdot f(x)\right) \left(S_{b^\ell,s} \cdot g(x)\right). \end{aligned} \tag{3}$$

1.2 Skew polynomials

In order to give a polynomial version of the section operators, we introduce the associative $k(x)$-algebra $k(x)\langle T\rangle$ generated by indeterminates $T_{b,r}$ with $0 \leq r < b$, subject to the product rule

$$T_{b,r} \times f(x) = \sum_{s=0}^{r} \left(S_{b,r-s} \cdot f(x)\right) T_{b,s} + x \sum_{s=r+1}^{b-1} \left(S_{b,r-s+b} \cdot f(x)\right) T_{b,s} \tag{4}$$

for all $f(x) \in k(x)$, which reflects (3) when $\ell = 1$. We refer to the elements of $k(x)\langle T\rangle$ as *skew polynomials*.

As this rule can be used to rewrite its left-hand side into its right-hand side, polynomials from the $k(x)$-algebra can be viewed as having monomials that are noncommutative words in the $T_{b,r}$ and coefficient from $k(x)$, written on the left of monomials.

We can view elements $f(x)$ from $k(x)$ as operators on $k((x))$, by considering their action by multiplication, and each $T_{b,r}$ as an operator on $k((x))$ by endowing it with the action of the section operator $S_{b,r}$. Then, the Leibniz formula (3) provides an expression for $T_{b,r} \cdot f(x) \cdot g(x) = S_{b,r} \cdot (f(x)g(x))$, which, by (4), matches the result $(T_{b,r}f(x)) \cdot g(x)$ of the action of the operator $T_{b,r}f(x)$ on $g(x)$. One checks that this defines a left action of $k(x)\langle T\rangle$ on $k((x))$.

1.3 Exponent notation

In the classical commutative case, computations on ideals use monomial orderings and very basic results about the exponents, which are elements of $\mathbb{Z}^m_{\geq 0}$. This leads us to introduce a parallel notation for the monomials. As exponents, we use words over the alphabet $\mathcal{A}$ of the numeration system with radix b, that is, the alphabet $\mathcal{A} = \{0, 1, \dots, b-1\}$. In other words, we have two notations $T^r = T_{b,r}$ with $0 \leq r < b$ for the generators of $k(x)\langle T\rangle$.

To make an explicit link with section operators, for any word $w \in \mathcal{A}^*$ introduce the integer r whose b-ary expansion is w:

$$r = (w)_b = w_{\ell-1} b^{\ell-1} + \cdots + w_0 b^0.$$

The monomial T^w acts on $k((x))$ as does $S_{b^\ell,r}$, which results from applying Equation (2) iteratively, since the action of a section operator $S_{b^\ell,r}$ on $k((x))$ is the same as the action of T^w, obtained as the composition of the action of $T^{w_{\ell-1}}$ after the action of $w_{\ell-2}, \dots$, after the action of T^{w_0}.

Hence we can extend the notation $T_{b,r}$ with $0 \leq r < b$ into $T_{b^\ell,r}$ with $0 \leq r < b^\ell$ by the relation

$$T^w = T_{b^\ell,r}, \quad \text{with } \ell = |w|,\ r = (w)_b. \tag{5}$$

Upon setting $\ell := |w|$, $\ell' := |w'|$, $r := (w)_b$, and $r' := (w')_b$, Equation (2) thus provides the simple formula $T^w T^{w'} = T^{ww'}$, where the word ww' is the concatenation of w and w'. As a consequence, the monoid of words and the monoid of monomials are clearly isomorphic. Furthermore, Formula (4) generalizes by changing b to b^ℓ in the formula, thus mimicking (3) for general ℓ. Written more loosely and after reindexing by words, the formula becomes

$$T^w \times f(x) = \sum_{|w'|=|w|} g_{w'}(x) T^{w'}, \tag{6}$$

for suitable rational functions $g_{w'}(x)$.

For example, with $b = 2$, both formulas $T^{01}T^{101} = T^{01101}$ and $T_{4,1}T_{8,5} = T_{32,13}$ mean the same; after applying to a sequence u, they become $u_{8(4n+1)+5} = u_{32n+13}$.

The length of words plays a role akin to the degree in the commutative case. This leads us to define the degree in $k(x)\langle T\rangle$ by $\deg 0 = -\infty$ and for a nonzero polynomial by the formula

$$\deg \sum_w c_w T^w = \max\{|w| \mid c_w \neq 0\}. \tag{7}$$

It satisfies the usual property of a degree with respect to the multiplication and the addition.

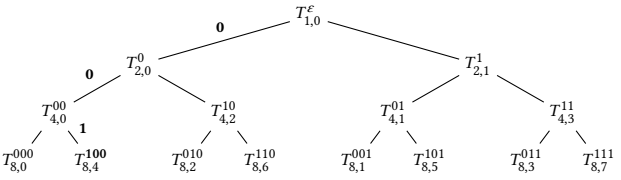

Figure 1: The tree of monomials for $b = 2$. Each node $T^{w}_{b^{\ell},r}$ denotes a monomial in two notations: $T^{w} = T_{b^{\ell},r}$. The path from the root to a monomial follows the word w, read from right to left, that is, from the least significant digits first.

In computational commutative algebra, it is usual to support a piece of reasoning by drawing so-called stairs. A polynomial is seen through its carrier, which is the set of its exponents. Similarly, we view a polynomial in $k(x)\langle T\rangle$ via the set of its exponents, which are words in $\mathcal{A}^*$. Owing to noncommutativity, these are the nodes of a b-ary tree, instead of the nodes of the square lattice (Fig. 1).

2 GRÖBNER BASES

In this section, we develop a theory for ideals of section operators, adapting what can be of the classical commutative theory [2, 4].

2.1 Monomial ordering

As opposed to the ordinary theories of Gröbner bases, our new theory makes use of a single monomial ordering, which is motivated by two constraints.

First, in our applications to divide-and-conquer recurrences, we do not want to produce recurrence formulas like $u_{2n+1} = u_{8n+3} + u_{4n} + u_n$, where the term $u_{b^{\ell}+r}$ on the left-hand side is defined by using some $u_{b^{\ell'}+r'}$ where ℓ' is larger than ℓ. Hence, we need an ordering that refines the degree.

Our second constraint is technical: Property 3 in Proposition 1 below will prove to be crucial to make our theory possible, in particular by the proof of Lemma 3. In practice, this leads to the choice of a single monomial ordering used in what follows.

Ordering monomials T^{w} is equivalent to ordering words w. In the case $b = 2$, our ordering lists the words involved as superscripts in Figure 1 in the order they appear when read by a breadth-first (left to right) traversal of the tree. We call it the *breadth-first ordering*: $\varepsilon < 0 < 1 < 00 < 10 < 01 < 11 < 000 < 100 < 010 < 110 < 001 < 101 < 011 < 111 < \cdots$. It can be defined formally as follows. First, order the alphabet $\mathcal{A}$ according to $0 < 1 < \cdots < b-1$. Next, words over $\mathcal{A}$ are first ranked by length, with ties broken by the lexicographical ordering on words read from right to left. (With our convention for defining $(w)_b$, this means from the least significant digit to the most significant one.) In other words, we define $w < w'$ if $|w| < |w'|$, or else if $|w| = |w'|$ and the two words can be written $w = ujv$ and $w' = u'j'v$ for words u, u', and v, and letters $j < j'$.

Proposition 1. *The breadth-first ordering on the monoid of monomials satisfies the following properties:*

1. *it is total and refines the degree,*
2. *every set of monomials has a smallest element,*
3. *it is* left compatible *with concatenation, that is if $|v| = |v'|$ and $T^{v} < T^{v'}$, then $T^{uv} < T^{u'v'}$ whenever $|u| = |u'|$.*

Proof. The first and third assertions are direct consequences of the definition of the order on words. The second assertion follows from the first and the fact that there exist only finitely many monomials of a given degree. □

2.2 Leading monomials

With a total ordering on monomials at our disposal, we can consider leading monomials and leading coefficients.

Definition 2. *The leading monomial* $\operatorname{lm}(F)$ *of a nonzero skew polynomial F is the largest monomial in F with respect to breadth-first ordering. The leading coefficient* $\operatorname{lc}(F)$ *is the coefficient of the leading monomial* $\operatorname{lm}(F)$.

A key point in the commutative case is the fact that the leading monomial of a product is the product of the leading monomials. Formula (4) induces a breach to this law, which seems to preclude the translation of the commutative case into our noncommutative case. Indeed, when we multiply a term $c(x)T^{v}$ by a monomial T^{u} on the left, we generally obtain all the monomials $T^{u'v}$ with $|u'| = |u|$ and not only the monomial T^{uv}.

Lemma 3. *The leading monomial w.r.t. breadth-first ordering of the product of two nonzero skew polynomials is the product of their leading monomials whenever the right-hand factor is monic.*

Proof. Let F and G be the two skew polynomials to be multiplied, with F nonzero and G monic. Without loss of generality, we can also assume that F is monic, as changing F into $1/\operatorname{lc}(F)\times F$ does not modify the leading monomial of the left factor, and, by associativity, of the product FG. Let T^{u} and T^{v} be the leading monomials of F and G, respectively. Without loss of generality, we can neglect the terms with degree smaller than the degree of those leading monomials, since the ordering refines the degree. So we consider

$$F = T^{u} + \sum_{u'<u} c_{u'}(x)T^{u'}, \qquad G = T^{v} + \sum_{v'<v} d_{v'}(x)T^{v'},$$

where u' and v' are words subject to $|u'| = |u|$, $|v'| = |v|$. Apart from the monomial $T^{u}T^{v} = T^{uv}$, which bears coefficient 1, the product FG includes two types of terms: first, terms $e(x)T^{u''v'}$ with $|u''| = |u|$, $v' < v$; second, terms $c_{u'}(x)T^{u'v}$ with $u' < u$. In the first case, the monomial $T^{u''v'}$ is smaller than T^{uv} since the breadth-first ordering is left compatible. In the second case, $T^{u'v}$ is smaller than T^{uv} because u' is smaller than u and, again, by the left compatibility with concatenation. Hence the leading monomial of the product is the product of the leading monomial. □

For convenience, we augment the monoid of monomials with the element 0, and its ordering so that 0 becomes its minimal element. We then extend the map $\operatorname{lm}(\cdot)$ by giving it the value 0 at the polynomial 0, so that $\operatorname{lm}(0) = 0 < \operatorname{lm}(F)$ for any nonzero skew polynomial F. The total order on the augmented monoid of monomials then induces a preorder on skew polynomials: for any F and G in $k(x)\langle T\rangle$, we say that F is smaller than G, denoted $F < G$, if $\operatorname{lm}(F) < \operatorname{lm}(G)$, and that F is smaller than or equivalent to G, denoted $F \le G$, if $\operatorname{lm}(F) \le \operatorname{lm}(G)$. Observe that the inequality $F < G$ is equivalent to any of the three inequalities obtained by replacing F by $\operatorname{lm}(F)$, G by $\operatorname{lm}(G)$, or both. We will use this equivalence freely.

The property of left compatibility for monomials extends to skew polynomials in the form of the next lemma.

Input: A dividend A, a list of nonzero divisors $(B_1, \dots, B_s)$.
Output: A list of quotients $(Q_1, \dots, Q_s)$ and a remainder R.

1. For i from 1 to s, do $B'_i := \mathrm{lc}(B_i)^{-1} \times B_i$.
2. $R := 0$. For i from 1 to s, do $Q'_i := 0$.
3. While $A \neq 0$ do
   a. if there exists i between 1 and s such that $\mathrm{lm}(B'_i)$ divides $\mathrm{lm}(A)$ on the right, then:
      (1) pick such an i,
      (2) $M := \mathrm{lc}(A)\,\mathrm{lm}(A)\,\mathrm{lm}(B'_i)^{-1}$,
      (3) $Q'_i := Q'_i + M$, $A := A - MB'_i$;
   b. otherwise: $R := R + \mathrm{lc}(A)\,\mathrm{lm}(A)$, $A := A - \mathrm{lc}(A)\,\mathrm{lm}(A)$.
4. For i from 1 to s, do $Q_i := Q'_i \times \mathrm{lc}(B_i)^{-1}$.
5. Return the list $(Q_1, \dots, Q_s)$ and R.

Algorithm 1: Right division algorithm in $k(x)\langle T\rangle$.

Lemma 4. *For any skew polynomials H, K_1 and K_2 from $k(x)\langle T\rangle$, if $H \neq 0$ and $K_1 < K_2$, then $HK_1 < HK_2$.*

Proof. As $\mathrm{lm}(K_1) < \mathrm{lm}(K_2)$, the second of these monomials is nonzero. Therefore, the polynomial K_2 is nonzero, and so is HK_2. Writing $q = \mathrm{lc}(K_2)$, we have $H \times q \neq 0$ and $q^{-1}K_1 < q^{-1}K_2$, so it is sufficient to prove the result for monic K_2. If $K_1 = 0$, then $HK_1 = 0 < HK_2$ and the result is proved. There remains the case $K_1 \neq 0$. For any term $h(x)T^u$ of H and any term $k(x)T^v$ of K_1, by Formula (6) there exist coefficients $g_{u'}(x)$ such that

$$h(x)T^u k(x)T^v = \sum_{u'} g_{u'}(x)T^{u'v},$$

with a sum over those u' satisfying $|u'| = |u|$. The left-compatibility of breadth-first ordering and the strict inequality $T^v \leq \mathrm{lm}(K_1) < \mathrm{lm}(K_2)$ imply $T^{u'v} < T^u\,\mathrm{lm}(K_2)$ for each u'. Therefore,

$$h(x)T^u k(x)T^v \leq \max_{u'} T^{u'v} < T^u\,\mathrm{lm}(K_2) \leq \mathrm{lm}(HK_2),$$

where the last inequality results from the monicity of K_2. Taking a maximum over u and v, we get $HK_1 \leq \max \mathrm{lm}(h(x)T^u k(x)T^v) < HK_2$, thus proving the result. □

2.3 Division

The needed restriction of right quotients to monic skew polynomials is first involved in right division. To work around the difficulty, we write a right division $A = QB+R$ in the form $A = (Q\times c)(c^{-1}\times B)+R$ where c is the leading coefficient of the polynomial B. Of course, we next adjust the computation by changing the quotient $Q' = Q \times c$ into $Q = Q' \times c^{-1}$. This leads to Algorithm 1, which is a simple adaptation to our setting of the usual division algorithm.

Proposition 5. *Given a tuple $(B_1, \dots, B_s)$ of nonzero polynomials in $k(x)\langle T\rangle$, every $A \in k(x)\langle T\rangle$ can be written $A = Q_1B_1+\cdots+Q_sB_s+R$ for polynomials $Q_1, \dots, Q_s, R$ satisfying the following conditions:*

- *the monomials in the remainder R are not divisible by any of the leading monomials of the divisors $B_1, \dots, B_s$;*
- *furthermore, each Q_iB_i satisfies $Q_iB_i \leq A$.*

Such a division is provided by Algorithm 1, whatever choices are made to resolve nondeterminism at Step 3a(1).

Proof. The proof is based on Algorithm 1. Let T^v be the leading monomial of the dividend A at any stage of the computation. If no divisor has a leading monomial that divides T^v, then the term with this monomial is moved from A to the remainder, so that the dividend is made smaller. If there is a divisor

$$B' = T^u + \sum_{u'<u} c_{u'}(x)T^{u'}$$

whose leading monomial T^u divides T^v, then $v = wu$ for some word w. Then we subtract T^wB' from A so that its leading monomial $T^v = T^wT^u$ is killed. According to Lemma 4,

$$T^w \sum_{u'<u} c_{u'}(x)T^{u'} < T^wT^u = T^v,$$

so the next dividend, $A - T^wB'$, is smaller than A.

This process thus produces a strictly decreasing sequence of monomials, given by the $\mathrm{lm}(A)$, which by Proposition 1 must have a lowest element. The process therefore terminates. The correction of the algorithm results from a loop invariant: the value of $A + Q'_1B'_1 + \cdots + Q'_sB'_s + R$ at each entry into the loop body of Step 3 is equal to the initial value of A. As the final value of A is zero, this proves the existence of the division. As the proof above does not depend on the choice of i at Step 3a(1), the final assertion holds. □

Example 6 (Natural ordering)**.** Instead of breadth-first ordering, we could have considered the 'natural' ordering $<_{\mathrm{nat}}$. As breadth-first ordering, it refines degree and is based on lexicographic ordering. But it compares b-ary expansions of integers from the most significant digit to the least significant digit, that is from left to right, contrary to breadth-first ordering which reads from right to left. In other words, given any ℓ and $0 \leq r$, $r' < b^\ell$, natural ordering has $T_{b^\ell,r} <_{\mathrm{nat}} T_{b^\ell,r'}$ if and only if $r < r'$.

Lemma 3 about the leading monomial of a product does not hold true with natural ordering. For example, with $b = 2$, $F = T_{4,2}$, $G = T_{2,1} + \frac{x^3}{1-x^4}T_{2,0}$, the product is $FG = T_{8,5} + \frac{x}{1-x}T_{8,6}$, with leading monomial $T_{8,6}$ for natural ordering, while the product of the leading monomials is $T_{4,2}T_{2,1} = T_{8,5}$.

Moreover, with $<_{\mathrm{nat}}$, it is possible that the division algorithm does not end. For $b = 2$, consider the dividend $F = T_{8,6}$ and the divisors $F_1 = T_{2,1} - \frac{x^3}{1-x^4}T_{2,0}$, $F_2 = T_{4,2} - \frac{1}{1-x^2}T_{4,1}$. The successive dividends are $P_{2k} = \frac{x^k}{(1-x)^{2k}}T_{8,6}$, $P_{2k+1} = \frac{x^k}{(1-x)^{2k+1}}T_{8,5}$, $k \geq 0$. The carrier of the P_k alternates between $T_{8,6}$ and $T_{8,5}$. Note that for breadth-first ordering, the division process ends immediately, because $\mathrm{lm}(F_1) = T_{2,1}$ and $\mathrm{lm}(F_2) = T_{4,1}$, none of which divides $T_{8,6}$.

2.4 Gröbner bases

In contrast with the commutative case, neither Hilbert's basis theorem nor Dickson's lemma is available. As a consequence, in the sequel we restrict to finitely generated left ideals by requesting that ideals be presented by an explicit finite set of generators.

Definition 7. *A Gröbner basis of a left ideal $\mathcal{I}$ in $k(x)\langle T\rangle$ is a finite subset $\mathcal{G}$ of $\mathcal{I}$ whose elements are monic and such that for every F in $\mathcal{I}$, the leading monomial $\mathrm{lm}(F)$ is a left multiple of the leading monomial $\mathrm{lm}(G)$ of some polynomial G in $\mathcal{G}$.*

Proposition 8. *Let $\mathcal{G}$ be a Gröbner basis for a left ideal $\mathcal{I}$. For every polynomial F, there is a unique polynomial R such that $F \equiv R \bmod \mathcal{I}$ and no monomial of R is divisible by a monomial in $\mathrm{lm}(\mathcal{G})$. As a*

consequence, R is the remainder of the division by $\mathcal{G}$ regardless of the chosen division strategy.

Proof. Let us assume that we have $F \equiv R_1 \equiv R_2 \bmod \mathcal{I}$ for distinct R_1 and R_2, both satisfying the condition with regard to $\mathrm{lm}(\mathcal{G})$. Then $R_1 - R_2$ is in $\mathcal{I}$ and nonzero. By the definition of a Gröbner basis, the leading monomial $\mathrm{lm}(R_1 - R_2)$ is divisible by a monomial in $\mathrm{lm}(\mathcal{G})$. But this is impossible since none of the monomials of R_1 and R_2 is divisible by a monomial in $\mathrm{lm}(\mathcal{G})$. We have thus shown the uniqueness of R.

In addition, whatever choices resolve nondeterminism in the division process, division provides us with some polynomial satisfying the two properties, and as a consequence of uniqueness, this polynomial is independent of the choices. □

A crucial ingredient in the theory of Gröbner bases in polynomial rings is the notion of S-polynomials: for two nonzero polynomials F and G, one considers the least common multiple of their leading monomials and forms a combination of F and G that kills this monomial. Owing to noncommutativity, least common multiples of monomials do not always exist in $k(x)\langle T\rangle$ and they are very specific when they do. In $k(x)\langle T\rangle$, a monomial T^u indeed divides another monomial T^v on the right, meaning there exists a quotient $Q \in k(x)\langle T\rangle$ satisfying $T^u = QT^v$ if and only if v is a suffix of u, in which case there exists a monomial w satisfying $u = wv$ and $Q = T^w$. Thus, when two monomials T^u and T^v have a least common multiple, this is necessarily one of the two monomials.

As is usual in theories of Gröbner bases where monomials need not have common multiples, like in the theory for polynomial modules, we define the S-polynomial of two monic polynomials P and Q of $k(x)\langle T\rangle$, with respective leading monomials T^u and T^v, to be 0 when neither T^u divides T^v nor T^v divides T^u, and to be $P - T^wQ$ when $T^u = T^wT^v$ for some w, respectively $Q - T^wP$ when $T^v = T^wT^u$ for some w. Observe that we restrict the definition to monic polynomials, as nonmonic divisors are ill-behaved.

We next obtain a characterization of Gröbner bases in $k(x)\langle T\rangle$ akin to that in commutative polynomial rings, via S-polynomials.

Theorem 9. *A family $\mathcal{G} = (G_i)_{1\le i\le m}$ of monic polynomials is a Gröbner basis of the left ideal $\mathcal{I}$ it generates if and only if, whenever $i \neq j$, there exists a choice resolving nondeterminism in the division of the S-polynomial of G_i and G_j by $\mathcal{G}$ that leads to a zero remainder.*

In relation to the forward implication, notice that by Proposition 8, any resolution of nondeterminism leads to a zero remainder.

Proof. Given a Gröbner basis $\mathcal{G}$, let us consider any two of its elements, H_1 and H_2, and their S-polynomial H. The division of H by $\mathcal{G}$ produces a remainder R in $\mathcal{I}$. If it was nonzero, by the definition of a Gröbner basis, its leading monomial would be a multiple of an element of $\mathcal{G}$, contradicting that R is a remainder. So R is nothing but 0, and more generally so do all S-polynomials.

Conversely, let $\mathcal{G} = (G_i)_{1\le i\le m}$ be a family of monic polynomials whose S-polynomials all admit zero as a remainder upon division by $\mathcal{G}$. Further, let F be a nonzero polynomial in the left ideal $\mathcal{I}$ generated by $\mathcal{G}$, which can be written

$$F = \sum_{i=1}^{m} H_i G_i. \tag{8}$$

In particular, H_i is a nonzero polynomial for at least one i. We set $M_i := \mathrm{lm}(H_iG_i)$ for $1 \le i \le m$, and $M := \max_{1\le i\le m} M_i$. Note the inequalities $0 < \mathrm{lm}(F) \le M$. We will show that if $\mathrm{lm}(F) < M$, then we can change the representation of F so as to reduce M. Postponing the proof, we therefore assume the equality $\mathrm{lm}(F) = M$, implying that M is one of the M_i, and that $\mathrm{lm}(F)$ is right divisible by $\mathrm{lm}(G_i)$. This proves that $\mathcal{G}$ is a Gröbner basis.

When $\mathrm{lm}(F) < M$, we can without loss of generality assume that for some integer s,

$$M = M_1 = \cdots = M_s > M_{s+1} \ge \cdots \ge M_m \ge 0,$$

and that $\mathrm{lm}(G_s) = \min_{1\le i\le s} \mathrm{lm}(G_i)$. Then, for each $\ell < s$, there exists w_ℓ such that $\mathrm{lm}(G_\ell) = T^{w_\ell}\,\mathrm{lm}(G_s)$, so that, by assumption, the S-polynomial $G_\ell - T^{w_\ell}G_s$ admits zero as a remainder upon division by $\mathcal{G}$: there is an exact-division formula

$$G_\ell - T^{w_\ell}G_s = \sum_{i=1}^{m} A_{\ell,i}G_i,$$

where the inequality $A_{\ell,i}G_i < \mathrm{lm}(G_\ell) = \mathrm{lm}(T^{w_\ell}G_s)$ holds for each i. By rewriting the G_ℓ in terms of those new sums into (8), we deduce the new expression

$$F = \sum_{\ell=1}^{s-1} H_\ell \left(T^{w_\ell}G_s + \sum_{i=1}^{m} A_{\ell,i}G_i \right) + \sum_{i=s}^{m} H_iG_i = QG_s + R$$

for

$$Q := H_s + \sum_{\ell=1}^{s-1} H_\ell T^{w_\ell}, \qquad R := \sum_{\ell=1}^{s-1}\sum_{i=1}^{m} H_\ell A_{\ell,i}G_i + \sum_{i=s+1}^{m} H_iG_i.$$

Since $M > 0$, note that H_ℓ is nonzero if $\ell < s$. So, for $\ell < s$ and any i, this and the inequality $A_{\ell,i}G_i < \mathrm{lm}(G_\ell)$ imply by Lemma 4 the strict inequality $H_\ell A_{\ell,i}G_i < \mathrm{lm}(H_\ell G_\ell) = M$. For $i > s$, the inequality $H_iG_i < M$ is strict as well. Adding all terms, this implies $R < M$, then, because $F < M$, also $QG_s = F - R < M$. Up to reordering, this makes $QG_s + R$ a new representation of F, with lowered maximal monomial M. □

2.5 A variant of Buchberger's algorithm

Buchberger's algorithm generalizes with minimal alterations.

Theorem 10. *The noncommutative variant of Buchberger's algorithm provided by Algorithm 2 terminates. Moreover, with the breadth-first ordering, it computes a Gröbner basis for the left ideal generated by the input $(F_i)_{1\le i\le s}$.*

Proof. According to Proposition 5, the calls to the division algorithm in Step 3b(2) return. The set $\mathcal{G} := (G_i)_{i=1,\dots,m}$ can change only in Step 3b(3)iii, if the remainder R is nonzero. In this case, the set of the leading monomials of the elements of $\mathcal{G}$ increases at this point. But all encountered monomials in the algorithm have a degree that is not more than the maximal degree in $\mathcal{F}$, primarily because the S-polynomials H considered at Step 3b(2) have this property. So the set $\mathrm{lm}(\mathcal{G})$ cannot grow indefinitely, proving that Step 3b(3)iii can happen only finitely many times. After that, for each S-polynomial H there exists a division of H by $\mathcal{G}$ with remainder 0, so that the algorithm terminates.

Let $\mathcal{G}_f$ be the value of $\mathcal{G}$ output from the algorithm, and consider a pair (G, G') in it, with G appearing as at a smaller index than G'

Input: A finite list $\mathcal{F} = (F_i)_{i=1,\dots,s}$ of nonzero skew polynomials.
Output: A finite list $\mathcal{G} = (G_i)_{i=1,\dots,m}$ of nonzero skew polynomials.

1. $m := s$. For i from 1 to m, do $G_i := \mathrm{lc}(F_i)^{-1} \times F_i$.
2. $\mathcal{P} := \{(G_i, G_j) \mid 1 \le i < j \le m\}$.
3. While $\mathcal{P} \neq \varnothing$ do:
   a. choose a pair (H_1, H_2) and remove it from $\mathcal{P}$;
   b. if one of the leading monomials of the pair divides the other, say, if $\mathrm{lm}(H_2) = T^w \mathrm{lm}(H_1)$:
      (1) compute the S-polynomial $H = H_2 - T^w H_1$,
      (2) divide H by $(G_i)_{i=1,\dots,m}$,
      (3) if the remainder R is not 0 then
         i. $R := \mathrm{lc}(R)^{-1} \times R$,
         ii. $\mathcal{P} := \mathcal{P} \cup \{(G_i, R) \mid 1 \le i \le m\}$,
         iii. set $m := m + 1$, then $G_m := R$.
4. Return $(G_i)_{i=1,\dots,m}$.

Algorithm 2: A variant of Buchberger's algorithm for the noncommutative algebra $k(x)\langle T\rangle$.

in $\mathcal{G}_f$. As the algorithm never removes any element from $\mathcal{G}$, the pair must have been introduced into $\mathcal{P}$ during the execution, and must have later been dealt with. Let $\mathcal{G}_0$ be the value of $\mathcal{G}$ at the time the pair has been considered, and considering the S-polynomial H of G and G'. A possible choice for the division of H by the final set $\mathcal{G}_f$ is, first, to reuse the exact same division steps that led to R, thus using only elements from $\mathcal{G}_0$, and, second, in case R is nonzero, to add with one division step, dividing by the element R of $\mathcal{G}_f$. In all cases, the division obtains zero as its remainder. Therefore, the output $\mathcal{G}$ is a Gröbner basis, as a consequence of Theorem 9. □

2.6 Reduced Gröbner bases

We continue by exploring properties of Gröbner bases that ensure their uniqueness for a fixed ideal I (and breadth-first ordering). The results and proofs of the present section are very similar to those of the classical commutative case.

Definition 11. *A Gröbner basis $\mathcal{G}$ is minimal, respectively reduced, if, for any two polynomials F and G in $\mathcal{G}$, the leading monomial of F does not divide the leading monomial of G, respectively any monomial of G.*

Proposition 12. *Every Gröbner basis $\mathcal{G}$ of a given left ideal I contains a minimal Gröbner basis for the same ideal I. Furthermore, any two minimal Gröbner bases for I have the same number of elements and the same set of leading monomials.*

Proof. Suppose F and G in $\mathcal{G}$ are such that $\mathrm{lm}(G)$ is a left multiple of $\mathrm{lm}(F)$. By transitivity of right divisibility, $\mathcal{G}' := \mathcal{G} \setminus \{G\}$ is another Gröbner basis. Let $H = G - T^w F$ be the S-polynomial of F and G. The division of H by $\mathcal{G}$ cannot involve the divisor G, as leading monomials exclude this possibility, and it has a zero remainder, because the set $\mathcal{G}$ is a Gröbner basis. So G is in the ideal generated by $\mathcal{G}'$. The latter is also a Gröbner basis for I.

Let $\mathcal{F} = (F_i)_{i=1,\dots,n}$ and $\mathcal{G} = (G_i)_{i=1,\dots,m}$ be two minimal Gröbner bases of I. Because $\mathcal{G}$ is a Gröbner basis, the leading monomial $\mathrm{lm}(F_1)$ is divisible by the leading monomial of some G_i. Without loss of generality, we can reindex the family $\mathcal{G}$ so that $\mathrm{lm}(G_1)$ divides $\mathrm{lm}(F_1)$. But $\mathrm{lm}(G_1)$ is by the same argument divisible by some $\mathrm{lm}(F_i)$, so that $\mathrm{lm}(F_i)$ divides $\mathrm{lm}(F_1)$, hence $i = 1$ as $\mathcal{F}$ is minimal. Consequently, $\mathrm{lm}(G_1)$ divides $\mathrm{lm}(F_1)$ and $\mathrm{lm}(F_1)$ divides $\mathrm{lm}(G_1)$, so that they are equal. We continue with $\mathrm{lm}(F_2)$, which neither divides $\mathrm{lm}(F_1)$ on the right nor is a left multiple of it, because the Gröbner basis $\mathcal{F}$ is minimal. As previously, up to some reindexation, we get G_2 satisfying $\mathrm{lm}(F_2) = \mathrm{lm}(G_2)$. The process continues until one of the lists is finished. If there remains an element in the other, say G_m in $\mathcal{G}$, we obtain a contradiction to minimality: $\mathrm{lm}(G_m)$ would be divisible by some leading monomial $\mathrm{lm}(F_i)$, that it to say by $\mathrm{lm}(G_i)$ with $i < m$. □

Input: A Gröbner basis $\mathcal{F} = (F_i)_{i=1,\dots,m}$ of an ideal of $k(x)\langle T\rangle$.
Output: A reduced Gröbner basis $\mathcal{G} = (G_i)_{i=1,\dots,r}$ of the same ideal.

1. $r := m$.
2. While some $\mathrm{lm}(F_i)$ is a left multiple of some $\mathrm{lm}(F_j)$ with $j \neq i$, set $\mathcal{F} := (F_1, \dots, F_{i-1}, F_{i+1}, \dots, F_r)$ and $r := r - 1$.
3. Set $\mathcal{G} := \mathcal{F}$.
4. For i from 1 to r:
   a. $\mathcal{G}' := (G_1, \dots, G_{i-1}, G_{i+1}, \dots, G_r)$;
   b. set G_i to the remainder R of G_i upon division by $\mathcal{G}'$.
5. Return $\mathcal{G} = (G_i)_{i=1,\dots,r}$.

Algorithm 3: Gröbner-basis reduction algorithm.

Both following propositions show that a reduced Gröbner basis of a left ideal generated by a finite set of skew polynomials exists and is unique. Note that the monomial ordering used is the breadth-first ordering and only this one.

Proposition 13. *A reduced Gröbner basis of a left ideal of $k(x)\langle T\rangle$ is unique.*

Proof. Observe that reduced Gröbner bases are minimal, so that their cardinality is fixed, and so are their set of leading monomials. Let $\mathcal{G} = \{G_1, G_2 \dots G_s\}$ and $\mathcal{G}' = \{G'_1, G'_2 \dots G'_s\}$ be two reduced Gröbner bases of the same left ideal. Without loss of generality, we can assume $\mathrm{lm}(G_i) = \mathrm{lm}(G'_i)$ for each i. Suppose that G_i and G'_i are different for some i. Then, the difference $G_i - G'_i$ is in the ideal, and its leading monomial M appears in at least one of G_i and G'_i, strictly below their leading monomials. If in G_i, M is a left multiple of some $\mathrm{lm}(G_j)$ for $j \neq i$, contradicting that $\mathcal{G}$ is reduced. If in G'_i, a similar argument applies. Therefore, $\mathcal{G} = \mathcal{G}'$. □

Proposition 14. *Algorithm 3 computes a reduced Gröbner basis from a Gröbner basis.*

Proof. Let $\mathcal{I}$ be the input ideal. The first two steps of Algorithm 3 replace $\mathcal{F}$ by some minimal Gröbner basis generating the same ideal by the method implicit in the proof of Proposition 12. Observe that the successive values of $\mathcal{G}$ along the loop at Step 4 are all minimal Gröbner bases of $\mathcal{I}$, with the family $\mathrm{lm}(\mathcal{G})$ kept invariant, as a result of G_i and R sharing the same leading monomial at Step 4b. Additionally, for each i, the remainder R write $\mathrm{lm}(G_i) - Q$ where Q involves no left multiple of any element of $\mathrm{lm}(\mathcal{G}')$, and in fact of any element of $\mathrm{lm}(\mathcal{G})$ as $Q < G_i$. As a result, the final family $\mathcal{G}$ is a reduced Gröbner basis of $\mathcal{I}$. □

Example 15. Let us consider the family of skew polynomials

$$T_{4,3} + \frac{1}{1-2x}T_{4,2} + \frac{x}{1-x^2}T_{2,1} + T_{2,0}, \quad T_{8,3} + \frac{1}{1-x}T_{8,2},$$
$$xT_{8,4} + \frac{2-x}{1-x}T_{4,2} + T_{4,0}, \quad T_{8,1} + T_{8,0} + T_{4,3}, \quad T_{8,1} + T_{8,2} + T_{8,0},$$
$$-\frac{x^3}{1-x^4}T_{8,2} + T_{8,4} + T_{2,1}, \quad T_{8,5} + \frac{x^2}{1-3x}T_{8,4} + T_{2,1}.$$

An instance of execution of Algorithm 2 begins by considering the pair between the first two polynomials, because of the relation $T_{8,3} = T_{2,0}T_{4,3}$. This provides the S-polynomial

$$T_{8,6} - \frac{2x^5 - 8x^4 - 3x^3 - 3x^2 - 9x + 6}{2x^4(2x^2 - 3x + 1)}T_{4,2} + \frac{4x-1}{2-2x}T_{2,0}$$
$$- \frac{4x^4 + 2x^3 + 3x^2 + 3x + 3}{2x^4}T_{4,0} + \frac{3x^6 + 4x^5 - 4x^4 - 3x^2 + 3}{2x^3(x^3 - x^2 - x + 1)}T_{2,1}.$$

The computation results in a Gröbner basis with 14 elements, whose leading monomials are: $T_{8,3}$, $T_{8,5}$, $T_{8,1}$, $T_{8,6}$, $T_{8,2}$, $T_{8,4}$, $T_{8,0}$, $T_{4,3}$, $T_{4,1}$, $T_{4,2}$, $T_{4,0}$, $T_{2,1}$, $T_{2,0}$. One of the polynomials in the basis has rational functions coefficients with degree 31 and numerical coefficients of order 10^{11}. There were 81 pairs dealt with. Among them, 21 gave S-polynomials and 14 of the 21 S-polynomials reduced to 0. After reduction by Algorithm 3, we find the Gröbner basis $\{T_{2,0}, T_{2,1}\}$.

3 THE LINEAR ALGEBRA APPROACH

In this section, we develop an algorithm reminiscent of Faugère's algorithm F_4 [6], but properties of section operators departing from those of commutative polynomials make specific variations needed. First, it results directly from the properties of division and the definition of S-polynomials that our variant of Buchberger's algorithm, Algorithm 2, performs all its calculations on an input $\mathcal{F}$ in the $k(x)$-vector space generated by the monomials $T^w \leq \max \mathrm{lm}(\mathcal{F})$. Second, divisions tend to involve dense polynomials, owing to the relation (6), which is amplified by the exponential growth with d of the number of monomials T^w of degree $|w| = d$.

Consequently, it seems adequate to perform a calculation that is incremental in the way of F_4, but with the unusual property of being confined in a finite-dimensional vector space known beforehand.

Given a finite set of generators of an ideal, we use the basis $\mathcal{B}$ of all monomials that are not larger than an adequate monomial T^u. Then, any polynomial $F \leq T^u$ in $k(x)\langle T\rangle$ can be represented by the row vector $V = \mathrm{mat}_{\mathcal{B}}(F)$, and conversely, any row vector V represents a polynomial $F = \mathrm{poly}_{\mathcal{B}}(V) = \sum_{v \leq u} V_v T^v$. By viewing matrices as families of rows, indexed by integers, a similar bijection is in place between families $\mathcal{F}$ of s polynomials and rectangular matrices M with s rows. We write $M = \mathrm{mat}_{\mathcal{B}}(\mathcal{F})$ and $\mathcal{F} = \mathrm{poly}_{\mathcal{B}}(M)$ accordingly. Furthermore, we extend the notion of leading monomial to vectors through these bijections, that is, we define $\mathrm{lm}(V) := \mathrm{lm}(\mathrm{poly}_{\mathcal{B}}(V))$. In the previous discussion, all (row) vectors and matrices have columns indexed by the words v such that $\varepsilon \leq v \leq u$. For pivoting considerations in linear algebra, we view those columns as sorted according to decreasing v, that is, so to say with u to the left and ε to the right.

As already emphasized in Section 2.4, the notion of S-polynomial is very particular in our context. A pair (H_1, H_2) of polynomials admits a nonzero S-polynomial $H_2 - T^w H_1$ only if $\mathrm{lm}(H_2) = T^w \mathrm{lm}(H_1)$ (up to order). A direct analogue of Faugère's "half pairs" would therefore consist of both polynomials H_2 and $T^w H_1$. But when we get to consider a pair (H_1, H_2) in our variant of the F_4 algorithm, the polynomial H_2 is already in the polynomial list $\mathrm{poly}_{\mathcal{B}}(R)$ of polynomials available as divisors. So, it suffices to add $T^w H_1$ to the list P of new half pairs. This motivates that our definition of HalfPairs intentionally forgets H_2.

Input: A finite family $\mathcal{F}$ of skew polynomials.
Output: A Gröbner basis $\mathcal{G}$ of the left ideal generated by $\mathcal{F}$.

1. Set $\mathcal{B} := (T^v)_{u \geq v \geq \varepsilon}$ for u such that $T^u = \max \mathrm{lm}(\mathcal{F})$.
2. $R := \mathrm{RowEchelon}((\mathrm{mat}_{\mathcal{B}}(\mathrm{lc}(F)^{-1} \times F))_{F \in \mathcal{F}})$.
3. $P := \mathrm{Preproc}(\mathrm{HalfPairs}(R, R), R)$.
4. While $P \neq \varnothing$ do
   a. $R^0 :=$ R augmented by stacking it above P,
   b. $R := \mathrm{RowEchelon}(R^0)$,
   c. $R^+ :=$ the rows V of R such that $\mathrm{lm}(V)$ is not in $\mathrm{lm}(R^0)$,
   d. $P := \mathrm{Preproc}(\mathrm{HalfPairs}(R, R^+), R)$.
5. Return $\mathcal{G} = \mathrm{poly}_{\mathcal{B}}(R)$.

where:

* $\mathrm{HalfPairs}(R^1, R^2)$ returns the rows $\mathrm{mat}_{\mathcal{B}}(T^w \mathrm{poly}_{\mathcal{B}}(V'))$ satisfying $\mathrm{lm}(V) = T^w \mathrm{lm}(V')$ for some word w, some row V in R^1 or R^2, and some row V' in the other one.
* $\mathrm{RowEchelon}(M)$ returns the variant of a row echelon form of M obtained by reducing each row by the rows above it, without interchanging any rows, but removing null rows, and by using leading coefficients of rows as pivots.
* $\mathrm{Preproc}(P, R)$ takes a monomial that appears in $\mathrm{poly}_{\mathcal{B}}(P)$ but not in $\mathrm{lm}(R \cup P)$, and is expressible as a product $T^w \mathrm{lm}(V)$ for some word w, some row V in R, then adds $\mathrm{mat}_{\mathcal{B}}(T^w \mathrm{poly}_{\mathcal{B}}(V))$ to P, and repeats until no such product can be added.

Algorithm 4: A variant of the F_4 algorithm for the noncommutative algebra $k(x)\langle T\rangle$.

Theorem 16. *The variant of the F_4 algorithm provided by Algorithm 4 terminates and returns a Gröbner basis of the left ideal of $k(x)\langle T\rangle$ generated by its input.*

Proof. The successive matrices R at Step 4b generate an increasing family of $k(x)$-vector spaces of rows, all of dimension at most the cardinality of $\mathcal{B}$. The termination of the algorithm is then immediate: as the span of R cannot grow indefinitely, at some point R^+ is empty, forcing P to be empty as well at the next step.

After the initialization of R at Step 2, the left ideal generated by $\mathrm{poly}_{\mathcal{B}}(R)$ is exactly the ideal generated by the input $\mathcal{F}$. Whether it be after Step 3 or Step 4d, $\mathrm{poly}_{\mathcal{B}}(P)$ contains only elements of the ideal generated by $\mathrm{poly}_{\mathcal{B}}(R)$. The ideal generated by $\mathrm{poly}_{\mathcal{B}}(R^0)$ is therefore equal to that generated by $\mathrm{poly}_{\mathcal{B}}(R)$, and this ideal is left unchanged upon changing R^0 to $\mathrm{RowEchelon}(R^0)$. We get that the ideal generated by $\mathrm{poly}_{\mathcal{B}}(R)$ remains unchanged after Steps 4a and 4b. By induction, this ideal, and therefore the ideal generated by the output $\mathcal{G}$, is the ideal generated by the input $\mathcal{F}$.

Next, the construction of R^0 at Step 4a and the definition of RowEchelon are so that the matrix R obtained at Step 4b is equal to the matrix R before, stacked above the matrix R^+ that will be extracted at Step 4c. Thus, any row vector introduced into R by Step 2 or 4b will remain there until the end of the algorithm.

problem	01	35	38	14	39	42	18	15	43
radix	2	2	3	2	3	2	3	2	2
deg/dim	3/14	6/127	4/161	5/63	5/485	4/31	4/161	6/127	5/63
#in/#out	7/2	5/5	5/5	5/5	5/5	24/1	4/4	6/6	48/1
Buchberger	0.29	1.89	2.09	0.46	9.10	4.90	1.64	1.98	69.95
F4	0.26	0.65	0.77	2.76	2.86	5.39	9.68	25.50	77.41
speed-up	1.09	2.91	2.70	0.17	3.18	0.91	0.17	0.08	0.90

Table 1: Selected timings, comparing the speeds of Algorithm 4 (F4) and Algorithm 2 (Buchberger). Our running example (Examples 15 and 17) corresponds to problem 01.

Finally, consider any two polynomials H_1 and H_2 of the output $\mathcal{G}$ satisfying $\operatorname{lm}(H_1) = T^w \operatorname{lm}(H_2)$ for some word w. This w is nonempty since $\operatorname{lm}(\mathcal{G})$ has no repeated no element. If both row vectors $V_1 = \operatorname{mat}_{\mathcal{B}}(H_1)$ and $V_2 = \operatorname{mat}_{\mathcal{B}}(H_2)$ were introduced at Step 2, they are considered at Step 3 to produce the half pair $\operatorname{mat}_{\mathcal{B}}(T^w H_2)$. Otherwise, the most recent of V_1 and V_2 was introduced at Step 4b and both vectors are considered at Step 4d to produce the half pair $\operatorname{mat}_{\mathcal{B}}(T^w H_2)$. In both cases, P is thus nonempty and the calculation continues to Step 4a with both V_1 and V_2 in R. After 4b, they are still in R, and $\operatorname{mat}_{\mathcal{B}}(T^w H_2)$ is a linear combination of the rows of R. As a consequence, the remainder of the division of the S-polynomial $H_1 - T^w H_2$ by $\operatorname{poly}_{\mathcal{B}}(R)$ is zero, and so is the remainder under division by $\mathcal{G}$. By Theorem 9, $\mathcal{G}$ is a Gröbner basis. □

Example 17 (Example 15 continued)**.** The maximum monomial is $T_{8,3} = T^{011}$ so that we use the basis of the T^w with $w = 011, 101, 001, 110, \dots, 0, \varepsilon$, that is $T_{8,3}, T_{8,5}, T_{8,1}, T_{8,6}, \dots, T_{2,0}, T_{1,0}$. The leading monomials of the input polynomials are $T_{4,3}$ and $T_{8,r}$ with $1 \le r \le 5$. The row echelon reduction at Step 2 brings up the monomial $T_{4,2}$. As $T_{4,3}$ divides $T_{8,3}$ and $T_{4,2}$ divides $T_{8,2}$, Step 3 computes two half pairs that are rows with leading monomials $T_{8,3}$ and $T_{8,2}$. Preprocessing adds a row with leading monomial $T_{8,6}$, resulting in P consisting of 3 rows. After stacking R and P at the first execution of the loop, the reduction at Step 4b discovers the monomials $T_{8,0}$ and $T_{4,1}$, hence the matrix R^+ at Step 4c has two rows. Next, Step 4d produces three half pairs with leading monomials $T_{8,6}$, $T_{8,1}$, and $T_{8,5}$, before preprocessing finds no row to be added, resulting in P consisting of only 3 rows. At this point, the matrix contains polynomials with maximal degree 14. It takes 3 executions of the main loop before the computation ends and returns a Gröbner basis with 13 polynomials, whose leading monomials are in fact all the elements of the basis $\mathcal{B}$ except for $T^\varepsilon = T_{1,0}$. The polynomials in intermediate calculations have degrees up to 19 and use integer coefficients up to $\approx 3.7\,10^{19}$.

4 IMPLEMENTATION AND EXPERIMENT

We implemented Algorithms 2 and 4 in Maple and computed reduced Gröbner bases of over 40 ideals. The script and the data are available at https://specfun.inria.fr/chyzak/gbdacr/. The timings obtained (Table 1) do not indicate any clear advantage of F4.

5 CONCLUSION

We have achieved our initial goal of a theory of Gröbner bases for divide-and-conquer systems. To the best of our knowledge, this is the first time such a theory has been developed in a context involving noncommutative words and twisted commutation rules simultaneously. We could overcome the difficulty that the leading monomial of a polynomial product need not be the product of the leading monomials.

As to efficiency, the contribution of the F_4 algorithm is unclear. It needs to be further studied in relation to other ingredients: rejection criteria; an incremental selection strategy of half pairs; modular variants of $k(x)$ compatible with the action of sections operators.

On the other hand, our theory extends to an algorithmic module theory, which we use in applications involving nonhomogeneous recurrence equations and systems. This will be developed elsewhere.

Finally, remark that Example 15 provides a system whose series solutions are all zero, although the ideal does not contain 1. Any series annihilated by the computed Gröbner basis, $\{T^0, T^1\}$, has indeed odd and even parts that are zero, and so is zero. Recovering 1 in the ideal is possible if one extends the algebra with a generator M to represent the Mahler operator, acting on series by $M \cdot f(x) = f(x^b)$. When $b = 2$, the action on series leads to the identity $1 = MT^0 + xMT^1$ to be enforced in the algebra. However, it also leads to $T^0 M = 1$ and $T^1 M = 0$, hence to an algebra with zero divisors. We have not tried to develop a Gröbner-basis theory for it.

REFERENCES

[1] J. Apel and W. Lassner. An extension of Buchberger's algorithm and calculations in enveloping fields of Lie algebras. *J. Symbolic Comput.*, 6(2-3):361–370, 1988.
[2] T. Becker and V. Weispfenning. *Gröbner Bases.* Springer, 1993.
[3] F. Chyzak and B. Salvy. Non-commutative elimination in Ore algebras proves multivariate identities. *Journal of Symbolic Computation*, 26(2):187–227, 1998.
[4] D. A. Cox, J. Little, and D. O'Shea. *Ideals, Varieties, and Algorithms.* Springer, 2015.
[5] P. Dumas. Asymptotic expansions for linear homogeneous divide-and-conquer recurrences. *Theoretical Computer Science*, 548:25–53, Sept. 2014.
[6] J.-C. Faugère. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra*, 139(1–3):61–88, June 1999.
[7] A. Galligo. Some algorithmic questions on ideals of differential operators. In *Eurocal'85, Vol. 2 (Linz, 1985)*, volume 204 of *Lecture Notes in Comput. Sci.*, pages 413–421. Springer, 1985.
[8] E. L. Green. An introduction to noncommutative Gröbner bases. In K. G. Fischer, P. Loustaunau, J. Shapiro, E. L. Green, and D. Farkas, editors, *Computational Algebra*, volume 151 of *Lecture Notes in Pure and Appl. Math.*, pages 167–190. Dekker, 1994. Proc. of the 5th Mid-Atlantic Algebra Conference (1993).
[9] H.-K. Hwang, S. Janson, and T.-H. Tsai. Exact and asymptotic solutions of a divide-and-conquer recurrence dividing at half. *ACM Transactions on Algorithms*, 13(4):1–43, Dec. 2017.
[10] A. Kandri-Rody and V. Weispfenning. Noncommutative Gröbner bases in algebras of solvable type. *Journal of Symbolic Computation*, 9(1):1–26, 1990.
[11] H. Kredel. *Solvable Polynomial Rings.* Reihe Mathematik. Shaker, Germany, 1993.
[12] R. La Scala and V. Levandovskyy. Skew polynomial rings, Gröbner bases and the letterplace embedding of the free associative algebra. *J. Symbolic Comput.*, 48:110–131, 2013.
[13] D. Lazard. Gröbner bases, Gaussian elimination and resolution of systems of algebraic equations. In *Lecture Notes in Comput. Sci.*, pages 146–156. 1983.
[14] V. Levandovskyy and H. Schönemann. Plural: a computer algebra system for noncommutative polynomial algebras. In *Proceedings of ISSAC'03 (Philadelphia, USA).* ACM Press, 2003.
[15] E. L. Mansfield and A. Szanto. Elimination theory for differential difference polynomials. In *Proceedings of ISSAC'03 (Philadelphia, USA).* ACM Press, 2003.
[16] F. Mora. Groebner bases for non-commutative polynomial rings. In *Algebraic Algorithms and Error-Correcting Codes*, pages 353–362. Springer, 1986.
[17] T. Mora. Standard bases and non-noetherianity: Non-commutative polynomial rings. In T. Beth and M. Clausen, editors, *Applicable Algebra, Error-Correcting Codes, Combinatorics and Computer Algebra*, pages 98–109. Springer, 1988.
[18] T. Mora. Groebner bases in non-commutative algebras. In *Proceedings of ISSAC'89 (Portland, USA)*, pages 150–161. Springer, 1989.
[19] N. Takayama. Gröbner basis and the problem of contiguous relations. *Japan Journal of Applied Mathematics*, 6(1):147–160, 1989.
[20] V. A. Ufnarovskiĭ. On the use of graphs for computing a basis, growth and Hilbert series of associative algebras. *Math. Sb.*, 68(2):417–428, 1991.
[21] V. Weispfenning. Finite Gröbner bases in non-Noetherian skew polynomial rings. In *Proceedings of ISSAC'92 (Berkeley, USA).* ACM Press, 1992.