Computing Border Bases

Achim Kehrein and Martin Kreuzer

Fachbereich Mathematik
Universität Dortmund
D-44221 Dortmund, Germany

Abstract

This paper presents several algorithms that compute border bases of a zero-dimensional ideal. The first relates to the FGLM algorithm as it uses a linear basis transformation. In particular, it is able to compute border bases that do not contain a reduced Gröbner basis. The second algorithm is based on a generic algorithm by Bernard Mourrain originally designed for computing an ideal basis that need not be a border basis. Our fully detailed algorithm computes a border basis of a zero-dimensional ideal from a given set of generators. To obtain concrete instructions we appeal to a degree-compatible term ordering σ and hence compute a border basis that contains the reduced σ-Gröbner basis. We show an example in which this computation actually has advantages over Buchberger's algorithm. Moreover, we formulate and prove two optimizations of the Border Basis Algorithm which reduce the dimensions of the linear algebra subproblems.

Key words: border basis, Gröbner basis, FGLM algorithm, Buchberger's algorithm

1 Introduction

Border bases play a key role in numerical polynomial algebra because they behave numerically better than Gröbner bases (see Stetter's book [11]). Auzinger and Stetter [1], Möller [9], and Mourrain [10] successfully used border bases to solve zero-dimensional polynomial systems of equations. In previous papers (see [7] and [6]) Robbiano and the authors laid a foundation for the algebraic theory of border bases. Here we address the question of how to compute a border basis of a zero-dimensional polynomial ideal I from a given set of generators.

Email addresses: Achim.Kehrein@mathematik.uni-dortmund.de (Achim Kehrein), Martin.Kreuzer@uni-dortmund.de (Martin Kreuzer).

The most straightforward idea is to compute a Gröbner basis and then to perform a base change in the spirit of the well-known FGLM technique (see [3]). To compare this algorithm to other approaches, we spell it out in Section 2. This easy method is able to compute the border basis of I with respect to any order ideal for which a border basis exists. Its drawback is that it involves a Gröbner basis computation which can be quite time consuming.

Then we move on to a technique inspired by Faugère's algorithms F_4 and F_5 (see [4] and [5]). A general framework for this technique was given by Mourrain in [10]. The idea is to enlarge the given set of generators by repeatedly applying all possible linear syzygies while simultaneously keeping the computation in a finite-dimensional vector subspace of the polynomial ring. This computational *universe* is enlarged only when necessary. Mourrain's generic algorithm applies this idea in a more general setting than we do: his sets of monomials *connected to 1* are not necessarily order ideals, and the generators it produces are not necessarily a border basis. The price for this generality is that it is difficult to make the choices involved in the algorithm explicit, and that it is uncertain whether the resulting type of generating sets satisfies the wonderful characterizations of border bases explained in [6]. Moreover, from the numerical point of view, according to [11, Example 2.21], "the computational use of such bases [not involving order ideals] is awkward".

In Section 4 we formulate and prove an algorithm which yields a concrete construction of an order ideal $\mathcal{O}$ and an $\mathcal{O}$-border basis of I. We buy its explicit nature at the expense of introducing a term ordering. Thus the resulting border basis will contain a reduced Gröbner basis and we have lost some of the flexibility border bases offer. On the other hand, in many cases our algorithm behaves better than Buchberger's algorithm does. We owe this to the simplicity of the Buchberger Criterion for border bases (see Proposition 4): in contrast to Gröbner bases, only *neighboring pairs* need to be considered and thus the degree of the border adapted S-polynomials stays comparatively small. The entire computation takes place in a degree bounded part of the polynomial ring and the algorithm never has to reduce S-polynomials of a degree much larger than the maximal degree of a border term. Another advantage is that the computation requires only K-linear reductions. And if a border basis with respect to some other order ideal is required, we can still follow the algorithm with the FGLM technique explained before.

Finally, in Section 5, we present improved versions of our border basis algorithm. These versions minimize the enlargements of the computational universe. Thereby we keep the size of the necessary linear algebra operations as small as possible. The resulting algorithm has been implemented in CoCoA and performs well in concrete examples.

2 Definitions and Basic Algorithm

In the following we adopt the notation from [7]. So, we work in the polynomial ring $P = K[x_1, \ldots, x_n]$ over a field K. The monoid of terms is $\mathbb{T}^n := \{x_1^{\alpha_1} \cdots x_n^{\alpha_n} \mid \alpha_1, \ldots, \alpha_n \in \mathbb{N}\}$ and, for every $d \in \mathbb{N}$, we let $\mathbb{T}^n_{\leq d}$ denote the set of terms with total degree at most d.

Definition 1 A set of terms $\mathcal{O} \subseteq \mathbb{T}^n$ is called an **order ideal** if it is closed under divisors, i.e. if $t \in \mathcal{O}$ and $t' \mid t$ imply $t' \in \mathcal{O}$. The **border** of a non-empty order ideal $\mathcal{O}$ is the set of terms $\partial\mathcal{O} = \{x_i\, t \mid 1 \leq i \leq n,\ t \in \mathcal{O}\} \setminus \mathcal{O}$; for the empty order ideal, we define $\partial\emptyset := \{1\}$.

The concept of an order ideal appears under many different names in the literature. We use "order ideal" in agreement with [2] and [12]. In the present paper, order ideals and consequently their borders consist of only finitely many terms. Unlass stated otherwise, we write $\mathcal{O} = \{t_1 \ldots, t_\mu\}$ and $\partial\mathcal{O} = \{b_1 \ldots, b_\nu\}$. (For $\mu = 0$, this includes the case $\mathcal{O} = \emptyset$.) Moreover, we shall reserve calligraphic symbols for finite sets of polynomials.

Definition 2 Let $\mathcal{O} = \{t_1, \ldots, t_\mu\}$ be an order ideal with border $\partial\mathcal{O} = \{b_1, \ldots, b_\nu\}$. Let $\mathcal{G} = \{g_1 \ldots, g_\nu\} \subset P$ be a set of polynomials and let $I \subseteq P$ be an ideal.

(1) The set $\mathcal{G}$ is an $\mathcal{O}$-**border prebasis** if the polynomials have the form

$$g_j = b_j - \sum_{i=1}^{\mu} \alpha_{ij} t_i \qquad \text{for} \quad 1 \leq j \leq \nu \quad \text{and} \quad \alpha_{ij} \in K.$$

(2) An $\mathcal{O}$-border prebasis $\mathcal{G}$ is an $\mathcal{O}$-**border basis** of I if $\mathcal{G}$ generates I and if $P = I \oplus \langle\mathcal{O}\rangle_K$ as vector spaces. If there exists an $\mathcal{O}$-border basis of I, we say that the order ideal $\mathcal{O}$ **supports** a border basis of I.

Four remarks are in order. First, it is sometimes convenient to express the direct sum condition "$P = I \oplus \langle\mathcal{O}\rangle_K$" in the the equivalent form "the residue classes of the terms in $\mathcal{O}$ constitute a vector basis of P/I." Second, the condition $\langle\mathcal{G}\rangle_P = I$ follows already from the mere inclusion $\mathcal{G} \subseteq I$ in combination with the direct sum condition. This is shown in [7] and is analoguous to Gröbner basis theory [8, Proposition 2.4.3a]. Third, since the order ideal is stipulated to consist of only finitely many, namely μ terms, the border basis definition implies $\dim_K(P/I) = \mu$. Thus the ideal I is necessarily zero-dimensional. Fourth, for each order ideal there is at most one border basis $\{g_1, \ldots, g_\nu\}$ of I because of the unique decomposition $b_j = g_j \oplus \sum_{i=1}^{\mu} \alpha_{ij} t_i$ for each $1 \leq j \leq \nu$.

In the termination proof of the Border Basis Algorithm below we require the following notions and result.

Definition 3 Let $\mathcal{G} = \{g_1, \ldots, g_\nu\}$ be an $\mathcal{O}$-border prebasis as in Definition 2. Two prebasis polynomials g_k, g_ℓ are **neighbors** if their border terms are related according to $x_i b_k = x_j b_\ell$ or $x_i b_k = b_\ell$ for some indeterminates x_i, x_j. Then, the corresponding **S-polynomials** are

$$S(g_k, g_\ell) := x_i g_k - x_j g_\ell \qquad \text{and} \qquad S(g_k, g_\ell) := x_i g_k - g_\ell$$

respectively.

Proposition 4 (Buchberger Criterion for Border Bases)
An $\mathcal{O}$-border prebasis $\mathcal{G} = \{g_1 \ldots, g_\nu\}$ is an $\mathcal{O}$-border basis of an ideal I if and only if $\mathcal{G} \subset I$ and, for each pair of neighboring prebasis polynomials, there are constant coefficients $c_j \in K$ such that

$$S(g_k, g_\ell) = c_1 g_1 + \ldots + c_\nu g_\nu.$$

The proof and a detailed discussion of these notions are included in [6].

It is helpful to understand the relationship between border bases and reduced Gröbner bases. So let us take a closer look at it. The Gröbner basis approach uses a term ordering σ and the set of leading terms $\mathrm{LT}_\sigma\{I\} = \{\mathrm{LT}_\sigma(f) \mid f \in I \setminus \{0\}\}$. The complementary set of terms $\mathcal{O}_\sigma\{I\} := \mathbb{T}^n \setminus \mathrm{LT}_\sigma\{I\}$ is an order ideal. Hence order ideals appear naturally in Gröbner basis theory.

Let $\mathcal{O}_\sigma\{I\} = \{s_1, \ldots, s_\mu\}$. By Macaulay's Basis Theorem the residue classes $\{\bar{s}_1, \ldots, \bar{s}_\mu\}$ form a vector basis of P/I; equivalently, we have the direct sum decomposition $P = I \oplus \langle \mathcal{O}_\sigma\{I\} \rangle_K$. Now let $\{\ell_1, \ldots, \ell_\lambda\} \subset \mathrm{LT}_\sigma\{I\}$ be the set of leading terms that are minimal in the sense that $\ell_l \in \mathrm{LT}_\sigma\{I\}$ while all proper divisors are in $\mathcal{O}_\sigma\{I\} = \mathbb{T}^n \setminus \mathrm{LT}_\sigma\{I\}$. The direct sum decomposition provides for each ℓ_l a unique polynomial h_l and unique coefficients $\beta_{il} \in K$ with $\ell_l = h_l + \sum_i \beta_{il} s_i$ and this yields the reduced σ-Gröbner basis $\{h_1, \ldots, h_\lambda\}$ of I. In the same way we obtain the $\mathcal{O}_\sigma\{I\}$-border basis $\{g_1, \ldots, g_\nu\}$ from the decompositions $b_j = g_j + \sum_i \alpha_{ij} s_i$ for each border term in $\{b_1, \ldots, b_\nu\} = \partial(\mathcal{O}_\sigma\{I\})$. Due to the minimal property, we have $\ell_1, \ldots, \ell_\lambda \in \partial(\mathcal{O}_\sigma\{I\})$, and the uniqueness of the decompositions implies $\{h_1, \ldots, h_\lambda\} \subseteq \{g_1 \ldots, g_\nu\}$. In this sense the $\mathcal{O}_\sigma\{I\}$-border basis of an ideal extends its reduced σ-Gröbner basis.

These observations motivate the following straightforward algorithm.

Proposition 5 (Basis Transformation Algorithm)
Let $I \subseteq P$ be a zero-dimensional ideal and $\mathcal{O} = \{t_1, \ldots, t_\mu\}$ an order ideal.

The following algorithm checks whether $\mathcal{O}$ supports a border basis of I and, in the affirmative, computes the $\mathcal{O}$-border basis $\{g_1, \ldots, g_\nu\}$ of I.

(T1) Choose a term ordering σ and compute $\mathcal{O}_\sigma\{I\} := \mathbb{T}^n \setminus \mathrm{LT}_\sigma\{I\}$.
(T2) If $|\mathcal{O}_\sigma\{I\}| \neq \mu$ then return "$\mathcal{O}$ has the wrong cardinality to support a border basis of I" and stop.
(T3) Let $\mathcal{O}_\sigma\{I\} = \{s_1, \ldots, s_\mu\}$. For $1 \leq m \leq \mu$, compute the coefficients $\tau_{im} \in K$ of the normal form $\mathrm{NF}_{\sigma,I}(t_m) = \sum_{i=1}^{\mu} \tau_{im} s_i$. Let $\mathcal{T}$ be the matrix $(\tau_{im})_{1 \leq i,m \leq \mu}$.
(T4) If $\det \mathcal{T} = 0$ then return "$\mathcal{O}$ has the wrong form to support a border basis of I" and stop.
(T5) Let $\partial\mathcal{O} = \{b_1, \ldots, b_\nu\}$. For $1 \leq j \leq \nu$, compute the coefficients $\beta_{ij} \in K$ of $\mathrm{NF}_{\sigma,I}(b_j) = \sum_{i=1}^{\mu} \beta_{ij} s_i$. Let $\mathcal{B}$ be the matrix $(\beta_{ij})_{1 \leq i \leq \mu, 1 \leq j \leq \nu}$.
(T6) Compute $(\alpha_{ij}) = \mathcal{T}^{-1}\mathcal{B}$. Return $g_j := b_j - \sum_{i=1}^{\mu} \alpha_{ij} t_i$ for $1 \leq j \leq \nu$.

Proof. By Macaulay's basis theorem, $|\mathcal{O}_\sigma\{I\}| = \dim_K P/I$. Step (T2) checks whether $\mathcal{O}$ has the correct number of terms to represent a vector basis of P/I. Step (T3) calculates the expansions of the vectors $\bar{t}_i$ with respect to the vector basis $\bar{\mathcal{O}}_\sigma(I) = \{\bar{s}_1, \ldots, \bar{s}_\mu\}$ of P/I:

$$\bar{t}_i = \tau_{1i}\bar{s}_1 + \ldots + \tau_{\mu i}\bar{s}_\mu \qquad \text{for} \quad 1 \leq i \leq \mu.$$

In matrix notation, $(\,\bar{t}_1 \;\; \ldots \;\; \bar{t}_\mu\,) = (\,\bar{s}_1 \;\; \ldots \;\; \bar{s}_\mu\,)\,\mathcal{T}$. Therefore, $\bar{\mathcal{O}}$ is a vector basis of P/I, if and only if $\mathcal{T}$ is invertible. Finally, the computation

$$\bar{b}_j = (\,\bar{s}_1 \;\; \ldots \;\; \bar{s}_\mu\,) \begin{pmatrix} \beta_{1j} \\ \vdots \\ \beta_{\mu j} \end{pmatrix} = (\,\bar{t}_1 \;\; \ldots \;\; \bar{t}_\mu\,)\,\mathcal{T}^{-1} \begin{pmatrix} \beta_{1j} \\ \vdots \\ \beta_{\mu j} \end{pmatrix}$$

in P/I implies that steps (T5) and (T6) produce the correct result. □

The following example serves several purposes. First, it provides a particular instance of the preceding algorithm. Secondly, it demonstrates that not every order ideal of the correct cardinality supports a border basis. Thirdly, it presents a border basis that is not an $\mathcal{O}_\sigma\{I\}$-border basis. In other words, there are border bases that are not extensions of reduced Gröbner bases.

Example 6 This example appeared in a different context in [7]. Let I be the vanishing ideal of the five points $(-1,1)$, $(1,1)$, $(0,0)$, $(1,0)$, and $(0,-1)$ in $\mathbb{A}^2(\mathbb{Q})$. Let $\sigma :=$ `DegRevLex`. Then $\mathcal{O}_\sigma\{I\} = \{1,\ x,\ y,\ xy,\ y^2\}$ and hence $\dim_{\mathbb{Q}} \mathbb{Q}[x,y]/I = 5$. The bivariate set of terms $\mathbb{T}^2$ contains seven order ideals with five elements, namely $\mathcal{O}_1 = \{1, x, x^2, x^3, x^4\}$, $\mathcal{O}_2 = \{1, x, x^2, x^3, y\}$, $\mathcal{O}_3 = \{1, x, y, y^2, y^3\}$, $\mathcal{O}_4 = \{1, y, y^2, y^3, y^4\}$ $\mathcal{O}_5 = \{1, x, x^2, y, xy\}$, $\mathcal{O}_6 = \{1, x, y, xy, y^2\}$, and $\mathcal{O}_7 = \{1, y, y^2, x, x^2\}$. Applying the Basis Transformation Algorithm to all seven order ideals respectively, we obtain three classes of results:

The order ideals $\mathcal{O}_1$, $\mathcal{O}_2$, $\mathcal{O}_3$, and $\mathcal{O}_4$ have the wrong form to support a border basis of I and, accordingly, the algorithm terminates in step (T4). (This can also be seen as follows: the vanishing ideal I contains $x^3 - x$ and $y^3 - y$, so $\langle \mathcal{O}_i \rangle_K \cap I \neq 0$ for $i = 1, 2, 3, 4$.)

Next, $\mathcal{O}_5$ and $\mathcal{O}_6$ support border bases of I. The former consists of $h_1 := x^3 - x$, $h_2 := x^2y - x^2 - xy + x$, $h_3 := y^2 - 2x^2 - 2xy + 2x + y$, and $h_4 := xy^2 - xy$; the latter comprises $k_1 := x^2 + xy - \frac{1}{2}y^2 - x - \frac{1}{2}y$, $k_2 := xy^2 - xy$, $k_3 := y^3 - y$, and $k_4 := x^2y - \frac{1}{2}y^2 - \frac{1}{2}y$. Both order ideals are of the form $\mathcal{O}_\sigma\{I\}$, for instance, with respect to lexicographical term orderings with $y < x$ and $x < y$ respectively. The first three polynomials constitute the reduced Gröbner basis respectively.

Finally, the calculation for $\mathcal{O}_7$ produces the border basis $g_1 := xy - \frac{1}{2}y - \frac{1}{2}y^2 - x + x^2$, $g_2 := xy^2 - \frac{1}{2}y - \frac{1}{2}y^2 - x + x^2$, $g_3 := x^3 - x$, $g_4 := y^3 - y$, and $g_5 := x^2y - \frac{1}{2}y - \frac{1}{2}y^2$ Note that this border basis consists of five polynomials in contrast to the two previous examples. Most importantly, this order ideal cannot be the complement of any $\mathrm{LT}_\sigma\{I\}$: the leading term of g_1 with respect to an arbitrary term ordering is x^2 or y^2; in either case, $\mathrm{LT}_\sigma\{I\} \cap \mathcal{O}_7 \neq \emptyset$.

The Basis Transformation Algorithm is similar to the well-known FGLM Algorithm [3]: both compute an ideal basis from a known Gröbner basis via a vector basis transformation. However, there is a fundamental difference. While the FGLM Algorithm uses a σ-Gröbner basis of I to compute $\mathcal{O}_\tau\{I\}$ term by term with respect to some new term ordering τ, the Basis Transformation Algorithm requires the complete order ideal $\mathcal{O}$ as input.

The Basis Transformation Algorithm is unsatisfactory since it significantly uses Gröbner basis calculations. In section 4 we present an algorithm that uses linear algebra instead. It is an adaptation of the algorithm by Mourrain described in the next section.

3 Mourrain's Generic Algorithm

Mourrain [10] proposed a generic algorithm for computing a more general concept than that of a border basis. For the reader's convenience, we list briefly the parts of his work that are pertinent to our approach. We slightly rephrase some material to prepare the ground for our adaptation.

Definition 7 Let V and W be vector subspaces of P and let $v_0 \in V$.

(1) Let $W^+ := W + x_1W + \ldots + x_nW$.

(2) Inductively define the vector subspaces

$$V_0 := \langle v_0 \rangle_K \quad \text{and} \quad V_{k+1} := V_k^+ \cap V \quad \text{for } k = 0, 1, 2, \ldots$$

Then V is **connected to** v_0 if the ascending chain $V_0 \subseteq V_1 \subseteq V_2 \subseteq \ldots$ converges to V in the sense $\bigcup_{k \geq 0} V_k = V$.

When V is generated by a set of terms we also say that the set of terms is connected to v_0. The set of terms $\{1, x, xy\}$ is connected to 1, but not an order ideal. In Mourrain's algorithm, sets of terms connected to 1 play the role that order ideals play in border basis theory. For easy reference, we introduce the following name.

Definition 8 Let $I \subseteq P$ be an ideal and $B \subseteq P$ be a vector subspace. We call B a **Mourrain base space** for I if it is connected to 1 and if $P = B \oplus I$ as vector spaces.

Proposition 9 (Mourrain's Generic Algorithm)
Let $\mathcal{F} := \{f_1, \ldots, f_s\} \subset P$ be a set of polynomials that generates a zero-dimensional ideal $I = \langle \mathcal{F} \rangle_P$. The following algorithm computes a Mourrain base space B for I.

(M1) Determine a finite-dimensional vector subspace $L \subseteq P$ that is connected to 1 and contains $f_1, \ldots, f_s$.
(M2) Let K_0 be the vector subspace $\langle f_1 \ldots, f_s \rangle_K$. Let $\ell = 0$.
(M3) Compute $K_{\ell+1} := K_\ell^+ \cap L$.
(M4) If $K_{\ell+1} \neq K_\ell$ then increase ℓ by 1 and go to step M3.
(M5) Compute a vector subspace B connected to 1 such that $L = B \oplus K_\ell$.
(M6) If $B^+ \not\subseteq L$ then replace L with L^+ and go to step 3. Otherwise return B and stop.

Mourrain [10] remarks on step M5 , "since L is connected to 1, the vector space B connected to 1 and supplementary to K_* [$= K_\ell$ in step M5] can be computed incrementally, starting from 1 (in the case where $1 \notin K_*$)." However, he does not specify how this computation should be effected. We are unaware of an example whose result is a set of terms connected to 1 but not an order ideal. The problem of making step M5 explicit is the starting point for the development of the border basis algorithm in the next section.

4 The Border Basis Algorithm

The K-linear span $\langle \mathcal{O} \rangle_K$ of an order ideal $\mathcal{O}$ is connected to 1. Thus Mourrain's framework relates to the border basis setting. Indeed, in this section we present an adaptation of Mourrain's algorithm to border bases. To do so, we serve the algorithm in easily digestible pieces. We begin with distilling a fundamental concept from Mourrain's generic algorithm.

Definition 10 Let $F \subseteq L$ be vector subspaces of P. We think of L as the **universe** in which our calculations are taking place. Define inductively the vector subspaces

$$F_0 := F \quad \text{and} \quad F_{k+1} := F_k^+ \cap L \quad \text{for } k = 0, 1, 2, \ldots$$

The union $F_L := \bigcup_{k \geq 0} F_k$ of the ascending chain $F_0 \subseteq F_1 \subseteq F_2 \subseteq \ldots$ is called the **L-stable span of** F. In particular, if $\mathcal{F} := \{f_1, \ldots, f_s\}$ is a set of polynomials spanning $F = \langle \mathcal{F} \rangle_K$, then we also write $\mathcal{F}_L$ instead of F_L.

This definition is motivated by the special case $L = P$ in which the P-stable span equals the ideal generated by $\mathcal{F}$, i.e. $F_P = \langle f_1, \ldots, f_r \rangle_P$. But of course, to keep things computable, we prefer finite-dimensional universes L. Let us collect some basic properties of stable spans that will come in handy later.

Lemma 11 Let $F \subseteq G \subseteq U \subseteq L$ be vector subspaces of P. Then the following relations hold.

$$F \subseteq F_L\,, \quad F_L = (F_L)_L\,, \quad F_L \subseteq G_L\,, \quad F_U \subseteq F_L\,, \quad \text{and} \quad F_L = (F_U)_L.$$

Proof. The first two relations are immediate consequences of the stable span's definition. Next, let $F_0 := F$ and $F_{k+1} := F_k^+ \cap L$ for $k \in \mathbb{N}$. Analogously define G_k for $k \in \mathbb{N}$. From $F \subseteq G$ and $F_{k+1} = F_k^+ \cap L \subseteq G_k^+ \cap L = G_{k+1}$ we deduce inductively $F_k \subseteq G_k$ for all k. Hence $F_L = \bigcup_k F_k \subseteq \bigcup_k G_k = G_L$ which proves the third relation. The similar proof of the fourth is skipped.

To derive the last relation, apply the third relation to $F \subseteq F_U$ to obtain $F_L \subseteq (F_U)_L$. The converse inclusion follows from forming the L-stable spans of both sides of $F_U \subseteq F_L$. □

The last property shows that the L-stable span can be computed from the U-stable span instead of from scratch.

Our next goal is to describe explicitly how to compute the stable span for the particular universe $L = \langle \mathbb{T}^n_{\leq d} \rangle_K$. This includes the repeatedly occuring

subproblem of computing a basis extension for vector spaces $\langle \mathcal{V} \rangle_K \subseteq \langle \mathcal{V} \cup \mathcal{G} \rangle_K$ where $\mathcal{V}$ is a basis of the smaller space and $\mathcal{G}$ comprises the additional generators. For this computation we use Gaussian elimination in the following form.

Lemma 12 Let σ be a term ordering and $\mathcal{V} = \{v_1, \ldots, v_r\} \subset P \setminus \{0\}$ a finite set of polynomials with pairwise different leading terms and leading coefficients equal to 1. Let $\mathcal{G} = \{g_1, \ldots, g_s\} \subset P$ be a finite set of polynomials. The following algorithm computes a finite set $\mathcal{W} \subset P$ with leading coefficients equal to 1 and such that $\mathcal{V} \cup \mathcal{W}$ has pairwise different leading terms and $\langle \mathcal{V} \cup \mathcal{W} \rangle_K = \langle \mathcal{V} \cup \mathcal{G} \rangle_K$. (The set $\mathcal{V}$ or $\mathcal{W}$ may be empty.)

(1) Let $\mathcal{H} := \mathcal{G}$ and $\varrho := 0$.
(2) If $\mathcal{H} = \emptyset$ then return $\mathcal{W} := \{v_{r+1} \ldots, v_{r+\varrho}\}$ and stop.
(3) Choose $f \in \mathcal{H}$ and remove it from $\mathcal{H}$. Let $i := 1$.
(4) If $f = 0$ or $i > r + \varrho$ then go to step 7.
(5) If $\mathrm{LT}_\sigma(f) = \mathrm{LT}_\sigma(v_i)$ then replace f with $f - \mathrm{LC}_\sigma(f) \cdot v_i$, reset $i := 1$ and go to step 4.
(6) Increase i by 1, and go to step 4.
(7) If $f \neq 0$ then increase ϱ by 1, and put $v_{r+\varrho} := f / \mathrm{LC}_\sigma(f)$. Continue with step 2.

Proof. The algorithm maintains the following invariant: The leading terms of the polynomials $v_1, \ldots, v_{r+\varrho}$ are pairwise different and

$$\langle \{v_1, \ldots, v_{r+\varrho}\} \cup \{f\} \cup \mathcal{H} \rangle_K = \langle \mathcal{V} \cup \mathcal{G} \rangle_K. \tag{1}$$

In the beginning, when f is still undefined, interpret $\{f\}$ as the empty set.

The loop of steps 4–6 is finite, since in each iteration either i increases or it is reset along with a reduction of the leading term of f. The latter can happen only finitely many times, since σ is a well ordering. Hence after finitely many iterations i is not reset anymore and, eventually, surpasses the unchanged upper bound $r + \varrho$. The reduction in step 5 does not alter the span in equation (1) and, when the loop terminates, either $f = 0$ or $\mathrm{LT}_\sigma(f) \notin \{\mathrm{LT}_\sigma(v_1) \ldots, \mathrm{LT}_\sigma(v_{r+\varrho})\}$. Also, the loop of steps 2–7 terminates: the set $\mathcal{H}$ is initialized as the finite set $\mathcal{G}$ and then each iteration removes one element from $\mathcal{H}$ while no elements are added. Thus, the whole algorithm terminates.

At termination, $\mathcal{H} = \emptyset$ and $\{f\} \subseteq \{0, v_{r+\varrho}\}$, so the invariant verifies the algorithm's correctness. □

The reason for using vector bases with pairwise different leading terms is

the following: if $\mathcal{V} = \{v_1, \ldots, v_r\}$ is a basis of a vector subspace $V \subseteq P$ with pairwise different leading terms then the set $\mathrm{LT}_\sigma\{V\} := \{\mathrm{LT}_\sigma(v) \mid v \in V \backslash \{0\}\}$ of leading terms of elements of V equals the set of leading terms of its basis $\mathrm{LT}_\sigma\{\mathcal{V}\} := \{\mathrm{LT}_\sigma(v_1), \ldots, \mathrm{LT}_\sigma(v_r)\}$.

Next, we must make the $^+$-operation on a vector subspace $\langle \mathcal{F} \rangle_K \subseteq P$ more explicit. So we abuse notation and also define a $^+$-operation on a set of polynomials $\mathcal{F} := \{f_1, \ldots, f_r\}$ by letting $\mathcal{F}^+ := \mathcal{F} \cup x_1\mathcal{F} \cup \ldots \cup x_n\mathcal{F}$. Thus we have $\langle \mathcal{F} \rangle_K^+ = \langle \mathcal{F}^+ \rangle_K$.

Proposition 13 (Computing a Stable Span)
Let $\mathcal{F} := \{f_1, \ldots, f_r\} \subset P$ and $L := \langle \mathbb{T}^n_{\leq d} \rangle_K$ with $f_1, \ldots, f_r \in L$, i.e. such that $d \geq \max\{\deg f_i \mid 1 \leq i \leq r\}$. Let σ be a degree-compatible term ordering. The following algorithm computes a vector basis $\mathcal{V}$ of the stable span $\mathcal{F}_L$. Moreover, the basis vectors have pairwise different leading terms.

(1) Compute a vector basis $\mathcal{V}$ of $\langle \mathcal{F} \rangle_K$ with pairwise different leading terms. (Apply the lemma to $\mathcal{V} = \emptyset$ and $\mathcal{G} := \mathcal{F}$.)
(2) Compute a basis extension $\mathcal{W}' := \{v'_{r+1} \ldots, v'_{r+\varrho'}\}$ for $\langle \mathcal{V} \rangle_K \subseteq \langle \mathcal{V}^+ \rangle_K$ so that the elements of $\mathcal{V} \cup \mathcal{W}'$ have pairwise different leading terms. (Apply the lemma to $\mathcal{V}$ and $\mathcal{G} := \mathcal{V}^+ \setminus \mathcal{V}$.)
(3) Let $\mathcal{W} := \{v_{r+1}, \ldots, v_{r+\varrho}\} := \{v \in \mathcal{W}' \mid \deg(v) \leq d\}$.
(4) If $\varrho > 0$ then replace $\mathcal{V}$ with $\mathcal{V} \cup \mathcal{W}$, increase r by ϱ, and go to step 2.
(5) Return $\mathcal{V}$.

Proof. Steps 2–4 maintain the following loop invariant. Each iteration of the loop of steps 2–4 starts with a finite set $\mathcal{V}$ with pairwise different leading terms and computes a finite set $\mathcal{W}$ such that $\mathcal{V} \cup \mathcal{W}$ has pairwise different leading terms and

$$\langle \mathcal{V} \rangle_K \subseteq \langle \mathcal{V} \cup \mathcal{W} \rangle_K = \langle \mathcal{V}^+ \rangle_K \cap L \subseteq L.$$

In particular, $\mathcal{V} \cup \mathcal{W}$ is a vector basis of $\langle \mathcal{V}^+ \rangle_K \cap L$.

By Lemma 12, step 1 computes a finite set $\mathcal{V}$ with pairwise different leading terms. So the loop invariant is correctly initialized. By the same lemma, step 2 determines a vector basis extension $\mathcal{W}'$ such that $\mathcal{V} \cup \mathcal{W}'$ is a vector basis of $\langle \mathcal{V} \rangle_K^+$ with pairwise different leading terms. Then step 4 intersects this subspace with L by discarding the polynomials of degree larger than d; here we use the degree-compatibility of $L = \langle \mathbb{T}^n_{\leq d} \rangle_K$ and of σ.

Another iteration is called in step 4 if and only if a non-empty basis extension $\mathcal{W}'$ has been computed. Since r increases by a positive ϱ with each new iteration while the upper bound $r < \dim_K L$ stays constant, this loop terminates. After termination the loop invariant becomes $\langle \mathcal{V} \rangle_K = \langle \mathcal{V} \rangle_K^+ \cap L$ which proves correctness. □

The stable span $\mathcal{F}_L$ is contained in L as well as in the ideal generated by $\mathcal{F}$, i.e. $\mathcal{F}_L \subseteq L \cap \langle \mathcal{F} \rangle_P$. The following example shows that this inclusion can be strict and that, in insufficiently large universes, this approximation depends on the set of generators $\mathcal{F}$ and not only on the generated ideal $\langle \mathcal{F} \rangle_P$.

Example 14 Let $\mathcal{F} := \{f_1, f_2, f_3\}$ with $f_1 := x^2y^2 + 1$, $f_2 := x^4$, and $f_3 := y^4$. Also, let $\mathcal{H} := \{1\}$. The sets $\mathcal{F}$ and $\mathcal{H}$ generate the same trivial ideal $\langle 1 \rangle_P$ because $1 = f_2 \cdot f_3 - f_1^2 + 2f_1$.

(1) Let $L := \langle \mathbb{T}^n_{\leq 4} \rangle_K$. Then $\mathcal{F}_L = \langle \mathcal{F} \rangle_K$ with $\dim_K \mathcal{F}_L = 3$, while $\mathcal{H}_L = L$ with $\dim_K \mathcal{H}_L = 10$.
(2) Let $L := \langle \mathbb{T}^n_{\leq 5} \rangle_K$. Then $\mathcal{F}_L = L = \mathcal{H}_L$.

The above computation of the stable span also includes information about an order ideal that is a candidate for supporting a border basis.

Proposition 15 *Let $\mathcal{F} := \{f_1, \ldots, f_s\} \subset P$ and let $L = \langle \mathbb{T}^n_{\leq d} \rangle_K$ such that $\mathcal{F} \subseteq L$. Then there exists an order ideal $\mathcal{O}$ such that*

$$L = \mathcal{F}_L \oplus \langle \mathcal{O} \rangle_K.$$

Namely, if σ is a degree-compatible term ordering and $\mathcal{V} := \{v_1, \ldots, v_r\}$ a vector basis of $\mathcal{F}_L$ with pairwise different leading terms then $\mathrm{LT}_\sigma\{\mathcal{F}_L\}$ is closed under multiples in L, i.e. if $t \in \mathbb{T}^n_{\leq d}$ and $\mathrm{LT}_\sigma(v) \mid t$ for some $v \in \mathcal{F}_L$ then $t = \mathrm{LT}_\sigma(w)$ for some $w \in \mathcal{F}_L$. Dually this states that $\mathcal{O} = \mathbb{T}^n_{\leq d} \setminus \{\mathrm{LT}_\sigma(v_1), \ldots, \mathrm{LT}_\sigma(v_r)\}$ is an order ideal. It also satisfies the above direct sum decomposition.

Proof. The definition of $\mathcal{O}$ and Steinitz's exchange lemma yield

$$L = \langle \mathrm{LT}_\sigma(v_1), \ldots \mathrm{LT}_\sigma(v_r) \rangle_K \oplus \langle \mathcal{O} \rangle_K = \langle v_1, \ldots v_r \rangle_K \oplus \langle \mathcal{O} \rangle_K = \mathcal{F}_L \oplus \langle \mathcal{O} \rangle_K.$$

To prove the order ideal property of $\mathcal{O}$, we show dually that for each term $t \in \mathbb{T}^n \setminus \mathcal{O}$ and each indeterminate x_i the product $x_i\, t$ is in $\mathbb{T}^n \setminus \mathcal{O}$. Since $\mathcal{O} \subseteq \mathbb{T}^n_{\leq d}$, we only have to consider the case $x_i\, t \in \mathbb{T}^n_{\leq d}$. As $t \notin \mathcal{O}$, there is a basis element $v \in \mathcal{V}$ such that $t = \mathrm{LT}_\sigma(v)$. We have $x_i v \in \mathcal{V}^+ \subseteq \mathcal{F}_L^+$ and, by case consideration, $\mathrm{LT}_\sigma(x_i v) = x_i t \in \mathbb{T}^n_{\leq d}$. Since σ is degree-compatible, $x_i v \in \langle \mathbb{T}^n_{\leq d} \rangle_K = L$. Thus, $x_i v \in \mathcal{F}_L^+ \cap L = \mathcal{F}_L$, and therefore $\mathrm{LT}_\sigma(x_i v) \in \mathrm{LT}_\sigma\{\mathcal{F}_L\} = \mathrm{LT}_\sigma\{\mathcal{V}\}$ which shows $x_i t \in \mathbb{T}^n_{\leq d} \setminus \mathcal{O}$. □

The next proposition presents the statement that will serve as stop criterion in the Border Basis Algorithm below. It checks whether the candidate order ideal actually supports a border basis. Note how the special case $L = P$ and $\tilde{I} = I$ resembles the definition of a border basis.

Proposition 16 *Let L be a vector subspace of P. Let $\tilde{I}$ be a vector subspace of a zero-dimensional ideal $I \subseteq P$ such that $\tilde{I}^+ \cap L = \tilde{I}$ and $\langle \tilde{I} \rangle_P = I$ (In a sense, $\tilde{I}$ is an L-stable approximation of I). Let $\mathcal{O} = \{t_1, \ldots, t_\mu\}$ be an order ideal such that*

$$L = \tilde{I} \oplus \langle \mathcal{O} \rangle_K.$$

If $\partial\mathcal{O} \subseteq L$ then $\mathcal{O}$ supports a border basis of I.

Proof. For each border term $b_j \in \partial\mathcal{O} \subseteq L$ the direct sum decomposition defines a polynomial $g_j \in \tilde{I}$ according to

$$b_j = g_j + \sum_{i=1}^{m} \alpha_{ij} t_i \in \tilde{I} \oplus \langle \mathcal{O} \rangle_K.$$

By construction, $G := \{g_1, \ldots, g_\nu\}$ is an $\mathcal{O}$-border prebasis.

Consider two neighboring prebasis polynomials g_k, g_ℓ. The support of their S-polynomial $S(g_k, g_\ell) \in \tilde{I}^+$ is contained in $\mathcal{O}^+$. Hence there are coefficients $c_j \in K$ such that $h := S(g_k, g_\ell) - \sum_{j=1}^{\nu} c_j g_j$ has its support in $\mathcal{O}$. Then $h \in \tilde{I}^+ \cap \langle \mathcal{O} \rangle_K = \tilde{I}^+ \cap L \cap \langle \mathcal{O} \rangle_K = \tilde{I} \cap \langle \mathcal{O} \rangle_K = \{0\}$. By the Buchberger Criterion for Border Bases, G is a border basis of $\langle \tilde{I} \rangle_P = I$. □

We have to deal with one more technicality. The following reduction process transforms a suitable set of polynomials into the wanted border basis.

Proposition 17 (Final Reduction Algorithm)
Let $\mathcal{F} = \{f_1, \ldots, f_s\} \subset P$ be a system of generators of a zero-dimensional ideal I. Let σ be a degree-compatible term ordering. Let L be an order ideal (e.g. $L = \mathbb{T}^n_{\leq d}$), $\mathcal{V}$ a vector basis of the span $\mathcal{F}_L$ with pairwise different leading terms, and let $\mathcal{O} := L \setminus \mathrm{LT}_\sigma(\mathcal{V})$ such that

$$L = \mathcal{F}_L \oplus \langle \mathcal{O} \rangle_K \qquad \text{and} \qquad \partial\mathcal{O} \subseteq L.$$

Then the following algorithm computes the $\mathcal{O}_\sigma\{I\}$-border basis $\{g_1, \ldots, g_\nu\}$.

(F1) Let $\mathcal{V}_R := \emptyset$.
(F2) If $\mathcal{V} = \emptyset$ then go to step (F8).
(F3) Determine in $\mathcal{V}$ the element v with minimal leading term. Remove it from $\mathcal{V}$.
(F4) Let $\mathcal{H} := \mathrm{Supp}(v) \setminus (\{\mathrm{LT}_\sigma(v)\} \cup \mathcal{O})$.
(F5) If $\mathcal{H} = \emptyset$ then append $v/\mathrm{LC}_\sigma(v)$ to $\mathcal{V}_R$ and go to step (F2).
(F6) For each $h \in \mathcal{H}$ determine $w_h \in \mathcal{V}_R$ and $c_h \in K$ such that $\mathrm{LT}_\sigma(w) = h$ and $h \notin \mathrm{Supp}(v - c_h \cdot w_h)$.
(F7) Replace v with $v - \sum_h c_h \cdot w_h$, append $v/\mathrm{LC}_\sigma(v)$ to $\mathcal{V}_R$ and go to step (F2).
(F8) Let $\partial\mathcal{O} = \{b_1, \ldots, b_\nu\}$. Determine for each $b_j \in \partial\mathcal{O}$ the polynomial $g_j \in \mathcal{V}_R$ with $b_j = \mathrm{LT}_\sigma(g_j)$. Return $g_1, \ldots, g_\nu$.

Proof. The hypotheses give a vector basis $\mathcal{V}$ of $\mathcal{F}_L$ so that for each $b_j \in \partial\mathcal{O}$ there is an $h_j \in \mathcal{V}$ with $b_j = \mathrm{LT}_\sigma(h_j)$. We have almost reached our goal, but it is still possible that $\mathrm{Supp}(h_j)$ contains terms outside $\{b_j\} \cup \mathcal{O}$. The algorithm reduces these unwanted terms.

The loop of steps (F2)–(F7) maintains the invariant $\langle \mathcal{V} \cup \{v\} \cup \mathcal{V}_R \rangle_K = \mathcal{F}_L$ where the set $\mathcal{V} \cup \{v\} \cup \mathcal{V}_R$ has pairwise different leading terms. (In the beginning, when v is undefined, interpret $\{v\}$ as the empty set.) Moreover, the elements of $\mathcal{V}_R$ are polynomials g with $\mathrm{Supp}(g) \subseteq \{\mathrm{LT}_\sigma(g)\} \cup \mathcal{O}$ and $\mathrm{LC}_\sigma(g) = 1$. This invariant property holds prior to the first iteration since the first part of the algorithm computed $\mathcal{V}$ as a vector basis of $\mathcal{F}_L$ with pairwise different leading terms and step (F1) defines $\mathcal{V}_R$ as the empty set. Each iteration removes from $\mathcal{V}$ the element v with minimal leading term. Thus, if $\mathrm{Supp}(v)$ contains a term outside $\{\mathrm{LT}_\sigma(v)\} \cup \mathcal{O}$ then it is necessarily the leading term of some element in $\mathcal{V}_R$: it must be in $\mathbb{T}^n_{\leq d} \setminus \mathcal{O} = \mathrm{LT}_\sigma\{\mathcal{F}_L\}$, which equals $\mathrm{LT}_\sigma\{\mathcal{V} \cup \{v\} \cup \mathcal{V}_R\}$, while it cannot be in $\mathrm{LT}_\sigma\{\{v\} \cup \mathcal{V}\}$ by the minimal property of $\mathrm{LT}_\sigma(v)$. Hence w and c in step (F6) do exist as stated. The loop of steps (F2)–(F7) is finite since each iteration removes one element from the finite set $\mathcal{V}$. At termination the invariant proves that we have obtained a vector basis $\mathcal{V}_R$ of $\mathcal{F}_L$ with pairwise different leading terms and that $\mathrm{Supp}(g) \subseteq \{\mathrm{LT}_\sigma(g)\} \cup \mathcal{O}$ for all $g \in \mathcal{V}_R$.

We have found polynomials $g_j \in \mathcal{V}_R \subseteq \mathcal{F}_L \subseteq I$ with $\mathrm{Supp}(g_j) \subseteq \mathrm{LT}_\sigma(g_j) \cup \mathcal{O}$, with $\mathrm{LC}_\sigma(g_j) = 1$, and $\{\mathrm{LT}_\sigma(g_1), \ldots \mathrm{LT}_\sigma(g_\nu)\} = \partial\mathcal{O}$. By our hypotheses and Proposition 16, the order ideal $\mathcal{O}$ supports a border basis. Due to the border basis uniqueness, the computed polynomials $g_1, \ldots, g_\nu$ constitute this $\mathcal{O}$-border basis. □

This Final Reduction Algorithm is more subtle than it may appear at first sight. The information contained in $\mathcal{V}$ is not limited to having for each border term one polynomial having this term in its support. In trying to get rid of the term ordering, we contemplated the following question.

Let I be a zero-dimensional ideal and $\mathcal{O}$ an order ideal that supports a border basis of I. Let $\partial\mathcal{O} = \{b_1 \ldots, b_\nu\}$ and $v_1, \ldots, v_\nu \in I$ be polynomials with $b_j \in \mathrm{Supp}(v_j)$ for all $1 \leq j \leq \nu$. Can there be an algorithm that computes the $\mathcal{O}$-border basis $\{g_1 \ldots, g_\nu\}$ of I?

The answer is negative. For example, consider the monomial ideal $I := \langle x^2, y \rangle_P$ in $P := K[x, y]$. The order ideal $\mathcal{O} = \{1, x\}$ supports the border basis $\{y, xy, x^2\}$ and I contains the polynomials $y, xy, x^2 + x^3$. Clearly, there is no way to obtain the basis polynomial x^2 via reductions using only the given polynomials.

Finally, we are ready to assemble the main algorithm.

Proposition 18 (Border Basis Algorithm)

Let $\mathcal{F} = \{f_1, \ldots, f_s\} \subset P$ be a set of polynomials that generates a zero-dimensional ideal $I = \langle \mathcal{F} \rangle_P$. Let σ be a degree-compatible term ordering. The following algorithm computes the $\mathcal{O}_\sigma\{I\}$-border basis $\{g_1, \ldots, g_\nu\}$.

(B1) Let $d := \max\{\deg(f_i) \mid 1 \leq i \leq s\}$ and $L := \langle \mathbb{T}^n_{\leq d} \rangle_K$.
(B2) Compute a vector space basis $\mathcal{V} = \{v_1, \ldots, v_r\}$ of $\langle \mathcal{F} \rangle_K$ with pairwise different leading terms.
(B3) Compute a basis extension $\mathcal{W}' := \{v'_{r+1}, \ldots, v'_{r+\varrho'}\}$ for $\langle \mathcal{V} \rangle_K \subseteq \langle \mathcal{V}^+ \rangle_K$ so that the elements of $\mathcal{V} \cup \mathcal{W}'$ have pairwise different leading terms.
(B4) Let $\mathcal{W} = \{v_{r+1}, \ldots, v_{r+\varrho}\} = \{v \in \mathcal{W}' \mid \deg(v) \leq d\}$.
(B5) If $\varrho > 0$ then replace $\mathcal{V}$ with $\mathcal{V} \cup \mathcal{W}$, increase r by ϱ, and go to step (B3).
(B6) Let $\mathcal{O} := \mathbb{T}^n_{\leq d} \setminus \{\mathrm{LT}_\sigma(v_1) \ldots \mathrm{LT}_\sigma(v_r)\}$.
(B7) If $\partial\mathcal{O} \not\subseteq L$ then increase d by one, update $L := \langle \mathbb{T}^n_{\leq d} \rangle_K$, and continue with step (B3).
(B8) Apply the Final Reduction Algorithm and return the polynomials $g_1, \ldots, g_\nu$ it computes.

Proof. Step (B1) initializes L so that $\mathcal{F} \subseteq L$. By Proposition 13, steps (B2)–(B5) compute a vector basis $\mathcal{V}$ of the stable span $\mathcal{F}_L$ with pairwise different leading terms. By Proposition 15, step (B6) defines an order ideal.

Now consider the loop of steps (B3)–(B7). Each new iteration starts with the updated universe $L := \langle \mathbb{T}^n_{\leq d} \rangle_K$ and a vector basis $\tilde{\mathcal{V}} := \mathcal{V}$ with pairwise different leading terms of the stable span $\mathcal{F}_U$ with respect to the preceding universe $U := \langle \mathbb{T}^n_{\leq d-1} \rangle_K$. Applying Proposition 13 to the set of polynomials $\tilde{\mathcal{V}}$, we see that steps (B3)–(B5) compute a vector basis $\mathcal{V}$ of the stable span $\tilde{\mathcal{V}}_L$, and Lemma 11 gives $\tilde{\mathcal{V}}_L = (\mathcal{F}_U)_L = \mathcal{F}_L$. Therefore each iteration ends with an updated vector basis $\mathcal{V}$ of $\mathcal{F}_L$ and an updated order ideal $\mathcal{O}$ such that $L = \mathcal{F}_L \oplus \langle \mathcal{O} \rangle_K$. Next we check that only finitely many iterations occur. Though the order ideal $\mathcal{O}_\sigma\{I\}$ and the border basis polynomials g_j have not been computed yet, they do exist. In particular, $\partial(\mathcal{O}_\sigma\{I\}) = \{\mathrm{LT}_\sigma(g_1), \ldots \mathrm{LT}_\sigma(g_\nu)\}$, and there are polynomials $h_{j1}, \ldots h_{js}$ such that

$$g_j = h_{j1} f_1 + \ldots + h_{js} f_s \qquad \text{for } 1 \leq j \leq \nu. \tag{2}$$

Let $\tilde{d} := \max(\{d\} \cup \{\deg(h_{ji} f_i) \mid 1 \leq i \leq s, 1 \leq j \leq \nu\})$. It suffices to consider the case that the loop has not terminated prior to reaching the iteration with parameter value $d = \tilde{d}$. This iteration uses the universe $L = \langle \mathbb{T}^n_{\leq \tilde{d}} \rangle_K$ and computes a vector basis $\mathcal{V}$ of $\mathcal{F}_L$ with pairwise different leading terms. By the choice of $\tilde{d}$, all summands $h_{ji} f_i$ in the expansions (2) are in L and hence $g_1, \ldots, g_\nu \in \mathcal{F}_L$. We have $\partial(\mathcal{O}_\sigma\{I\}) = \{\mathrm{LT}_\sigma(g_1), \ldots \mathrm{LT}_\sigma(g_\nu)\} \subseteq \mathrm{LT}_\sigma\{\mathcal{F}_L\}$. Since $\mathrm{LT}_\sigma\{\mathcal{F}_L\}$ is closed under multiples in L (Proposition 15), we deduce

$\mathbb{T}^n_{\leq \tilde{d}} \setminus \mathcal{O}_\sigma\{I\} \subseteq \mathrm{LT}_\sigma\{\mathcal{F}_L\}$. Therefore, $\mathcal{O}_\sigma\{I\}$ contains the order ideal $\mathcal{O} := \mathbb{T}^n_{\leq \tilde{d}} \setminus \mathrm{LT}_\sigma\{\mathcal{F}_L\}$ which is determined by the current iteration. This leads to $\partial\mathcal{O} \subseteq \mathcal{O}_\sigma\{I\} \cup \partial(\mathcal{O}_\sigma\{I\}) \subseteq L$ and the loop terminates.

Having reached step (B8), we have a vector basis $\mathcal{V}$ of $\mathcal{F}_L$ that satisfies all hypotheses of Proposition 17. Hence the Final Reduction Algorithm computes the $\mathcal{O}$-border basis. Above we showed $\mathcal{O} \subseteq \mathcal{O}_\sigma\{I\}$. Both order ideals support border bases of I; in particular, they must consist of the same finite number of terms and therefore coincide. □

The $\mathcal{O}_\sigma\{I\}$-border basis computed by the algorithm is an extension of the reduced σ-Gröbner basis. Hence we better show an example in which this algorithm performs better than Buchberger's algorithm.

Example 19 We consider the zero-dimensional ideal I generated by $\mathcal{F} := \{f_1, \ldots, f_5\}$ where $f_1 = x^3 - x$, $f_2 = y^3 - y$, $f_3 = x^2y - \frac{1}{2}y - \frac{1}{2}y^2$, $f_4 = xy - x - \frac{1}{2}y + x^2 - \frac{1}{2}y^2$, and $f_5 = xy^2 - x - \frac{1}{2}y + x^2 - \frac{1}{2}y^2$; this is the $\mathcal{O}_7$-border basis in Example 6. Let σ be the degree-lexicographic term ordering `DegLex` on $\mathbb{T}^2$. First we compute the $\mathcal{O}_\sigma\{I\}$-border basis according to the steps of the above Border Basis Algorithm.

(B1) The generators induce the universe $L := \langle \mathbb{T}^2_{\leq 3} \rangle_{\mathbb{Q}}$.
(B2) The set of generators is a vector basis of $\langle \mathcal{F} \rangle_{\mathbb{Q}}$ with pairwise different leading terms, hence (rewritten with respect to `DegLex`) $\mathcal{V} = \{x^3 - x, y^3 - y, x^2y - \frac{1}{2}y^2 - \frac{1}{2}y, x^2 + xy - \frac{1}{2}y^2 - x - \frac{1}{2}y, xy^2 + x^2 - \frac{1}{2}y^2 - x - \frac{1}{2}y\}$.
(B3) We obtain $\mathcal{V}^+ = \{x^4 - x^2, x^3y - xy, xy^3 - xy, y^4 - y^2, x^3y - \frac{1}{2}xy^2 - \frac{1}{2}xy, x^2y^2 - \frac{1}{2}y^3 - \frac{1}{2}y^2, x^3 + x^2y - \frac{1}{2}xy^2 - x^2 - \frac{1}{2}xy, x^2y + xy^2 - \frac{1}{2}y^3 - x - \frac{1}{2}y^2, x^2y^2 + x^2 - \frac{1}{2}xy^2 - x^2 - \frac{1}{2}xy, xy^3 + x^2y - \frac{1}{2}y^3 - xy - \frac{1}{2}y^2\}$. A basis extension with pairwise different leading terms is $\mathcal{W}' = \{x^4 - x^2, x^3y - xy, xy^3 - xy, y^4 - y^2, x^2y^2 - \frac{1}{2}y^3 - \frac{1}{2}y^2\}$.
(B4) $\mathcal{W} = \mathcal{W}' \cap L = \emptyset$.
(B6) The algorithm computes the order ideal $\mathcal{O} = \{1, x, y, y^2, xy\}$ with border $\partial\mathcal{O} = \{x^2, xy^2, x^2y, y^3\}$ which is contained in the universe.
(B8) Let $\mathcal{V}_R := \emptyset$. The Final Reduction Algorithm processes the elements of $\mathcal{V}$ in the order $v = x^2 + xy - \frac{1}{2}y^2 - x - \frac{1}{2}y$ ($\mathcal{H} = \emptyset$), $v = y^3 - y$ ($\mathcal{H} = \emptyset$), $v = xy^2 + x^2 - \frac{1}{2}y^2 - x - \frac{1}{2}y$ ($\mathcal{H} = \{x^2\}$; replace v with $v - 1 \cdot (x^2 + xy - \frac{1}{2}y^2 - x - \frac{1}{2}y)$, thus $v = xy^2 - xy$.), $v = x^2y - \frac{1}{2}y^2 - \frac{1}{2}y$ ($\mathcal{H} = \emptyset$), and $v = x^3 - x$ ($\mathcal{H} = \emptyset$). Eventually, the border basis $\{x^2 + xy - \frac{1}{2}y^2 - x - \frac{1}{2}y, y^3 - y, xy^2 - xy, x^2y - \frac{1}{2}y^2 - \frac{1}{2}y\}$ is returned.

On the other hand, the Buchberger algorithm computes the S-polynomials $S_{12} = -x^6 + x^3y^3 + x^4 - xy^3$, $S_{13} = -x^4 + x^3y + x^2 - xy$, $S_{14} = -x^4 + x^3 + x^2 - x$, $S_{15} = -x^5 + x^3y^2 + x^3 - xy^2$, $S_{23} = x^2y^3 - y^5 - x^2y + y^3$, $S_{24} = -y^6 + x^2y^3 + y^4 - x^2y$, $S_{25} = xy^3 - y^4 - xy + y^2$, $S_{34} = -x^2y^2 + x^2y + \frac{1}{2}y^3 - \frac{1}{2}y$, $S_{35} =$

$-x^3y+x^2y^2+\frac{1}{2}xy^2-\frac{1}{2}y^3+\frac{1}{2}xy-\frac{1}{2}y^2$, $S_{45} = x^2y^2+xy^3-\frac{1}{2}y^4-x^3-x^2y-\frac{1}{2}xy^2-\frac{1}{2}y^3+x^2+\frac{1}{2}xy$. The pairs of generators (f_1, f_2) and (f_2, f_4) have relatively prime leading terms; we let the Buchberger algorithm justifiably disregard them. All S-polynomials reduce to zero, hence the system of generators is already a Gröbner basis.

How do these calculations differ? The border basis computation requires only the terms up to degree 3. In the Buchberger computation the nine additional terms $x^5, x^3y^2, x^2y^3, y^5, x^4, x^3y, x^2y^2, xy^3, y^4$ appear (This list excludes the terms x^6, x^3y^3, and y^6 that appear in the S-polynomials whose generators have relatively prime leading terms). So, this calculation produces terms up to degree 5 which subsequently need to be reduced. Thus, even in this small example we observe a redundancy in the Buchberger algorithm that is avoided by the border basis algorithm.

Let us compute another, more complicated example.

Example 20 This time we consider the vanishing ideal of the points $(-1, 0, 0)$, $(0, 0, 0)$, $(1, 0, 0)$, $(3, 0, 0)$, $(5, 0, 0)$, $(4, 4, 4)$, and $(0, 0, 7)$ in $\mathbb{A}^3(\mathbb{Q})$. It is generated by the set of polynomials $\{z^2 + 3y - 7z,\ yz - 4y,\ xz - 4y,\ y^2 - 4y,\ xy - 4y,\ x^5 - 8x^4 + 14x^3 + 8x^2 - 15x + 15y\}$. Let $\sigma :=$ `DegRevLex`. The Border Basis Algorithm starts with the universe $\langle \mathbb{T}^3_{\leq 5} \rangle_K$ which consists of 56 terms. To compute the stable span, the algorithm performs four linear basis extensions. Then it obtains the order ideal $\mathcal{O} = \{1,\ x,\ x^2,\ x^3,\ x^4,\ y,\ z\}$ whose border is already contained in the universe. The universe need not be enlarged. The border basis is the set of 12 polynomials $\{z^2+3y-7z,\ yz-4y,\ xz - 4y,\ y^2 - 4y,\ xy - 4y,\ x^2z - 16y,\ x^2y - 16y,\ x^3z - 64y,\ x^3y - 64y,\ x^4z - 256y,\ x^4y - 256y,\ x^5 - 8x^4 + 14x^3 + 8x^2 - 15x + 15y\}$.

The Buchberger Algorithm applied to this example works with S-polynomials up to degree 6 (There are S-polynomials of degree 7, but they belong to pairs of polynomials with relatively prime leading terms). The difference is not particularly striking here, but it is still there.

5 Some Optimizations of the Border Basis Algorithm

The following improved version of the Border Basis Algorithm replaces the use of $\mathbb{T}^n_{\leq d}$ as a computational universe with order ideals which are kept as small as possible.

Proposition 21 (Improved Border Basis Algorithm)
Let $\mathcal{F} = \{f_1, \ldots, f_s\} \subset P$ be polynomials that generate a zero-dimensional

ideal $I := \langle \mathcal{F} \rangle_P$. *Let* σ *be a degree-compatible term ordering. The following algorithm computes the* $\mathcal{O}_\sigma\{I\}$*-border basis* $\{g_1 \ldots, g_\nu\}$.

(I1) Let $\mathcal{L}$ *be the order ideal spanned by* $\bigcup_{i=1}^r \operatorname{Supp}(f_i)$.
(I2) Compute a vector basis $\mathcal{V}$ *of* $\langle \mathcal{F} \rangle_K$ *with pairwise different leading terms.*
(I3) Compute a basis extension $\mathcal{W}'$ *for* $\langle \mathcal{V} \rangle_K \subseteq \langle \mathcal{V}^+ \rangle_K$ *so that the elements of* $\mathcal{V} \cup \mathcal{W}'$ *have pairwise different leading terms.*
(I4) Let $\mathcal{W} := \{w \in \mathcal{W}' \mid \mathrm{LT}_\sigma(w) \in \mathcal{L}\}$.
(I5) If $\bigcup_{w \in \mathcal{W}} \operatorname{Supp}(w) \not\subseteq \mathcal{L}$ *then replace* $\mathcal{L}$ *with the order ideal spanned by* $\mathcal{L}$ *and* $\bigcup_{w \in \mathcal{W}} \operatorname{Supp}(w)$ *and continue with step (I4).*
(I6) If $\mathcal{W} \neq \emptyset$ *then replace* $\mathcal{V}$ *with* $\mathcal{V} \cup \mathcal{W}$ *and continue with step (I3).*
(I7) Let $\mathcal{O} := \mathcal{L} \setminus \{\mathrm{LT}_\sigma(v) \mid v \in \mathcal{V}\}$.
(I8) If $\partial\mathcal{O} \not\subseteq \mathcal{L}$ *then replace* $\mathcal{L}$ *with the order ideal* $\mathcal{L}^+$ *and continue with step (I3).*
(I9) Apply the Final Reduction Algorithm and return the polynomials $g_1, \ldots, g_\nu$ *computed by it.*

Proof. To show that the procedure terminates and that the algorithm is correct, we consider its loops in order of their appearance. The subloop of steps (I4)–(I5) is finite because $\mathcal{W} \subseteq \mathcal{W}'$ implies that each instance of $\mathcal{L}$ is contained in the invariant order ideal spanned by $\cup_{v \in \mathcal{V} \cup \mathcal{W}'} \operatorname{Supp}(v)$. Since each new iteration corresponds to an enlargement of $\mathcal{L}$ inside this invariant finite set, there can be only finitely many iterations.

When this subloop terminates, we have $\langle \mathcal{V} \cup \mathcal{W} \rangle_K = \langle \mathcal{V} \cup \mathcal{W}' \rangle_K \cap \langle \mathcal{L} \rangle_K$. The left-hand side is contained in the right-hand side because the universe enlargements in (I5) insure that the premise $w \in \mathcal{W}$, i.e. $\mathrm{LT}_\sigma(w) \in \mathcal{L}$, implies $\operatorname{Supp}(w) \subseteq \mathcal{L}$. For the reverse inclusion, let $v = \alpha_1 v_1 + \ldots + \alpha_r v_r + \beta_1 w_1 + \ldots + \beta_s w_s$ be in $\langle \mathcal{L} \rangle_k$, where $\alpha_1, \ldots, \alpha_r, \beta_1, \ldots, \beta_s \in K \setminus \{0\}$, $v_1, \ldots, v_r \in \mathcal{V}$, and $w_1, \ldots, w_s \in \mathcal{W}'$. The vectors are assumed to be pairwise different. We need to show that $w_1, \ldots, w_s \in \mathcal{W}$. Since the leading terms of $\mathcal{V} \cup \mathcal{W}'$ are pairwise different, $\mathrm{LT}_\sigma(v)$ equals some $\mathrm{LT}_\sigma(v_i)$ or some $\mathrm{LT}_\sigma(w_j)$. We know $\operatorname{Supp}(v_i) \subseteq \mathcal{L}$; hence, in the former case, we obtain $v - \alpha_i v_i \in \langle \mathcal{V} \cup \mathcal{W}' \rangle_K \cap \langle \mathcal{L} \rangle_K$. In the latter case, we deduce from $\mathrm{LT}_\sigma(w_j) = \mathrm{LT}_\sigma(v) \in \mathcal{L}$ that $w_j \in \mathcal{W}$. We get $v - \beta_j w_j \in \langle \mathcal{V} \cup \mathcal{W}' \rangle_K \cap \langle \mathcal{L} \rangle_K$. The desired inclusion follows by induction.

Next we show that the loop of steps (I3)-(I6) is finite. At the beginning of an arbitrary iteration let $\mathcal{L}$ be contained in some $\mathbb{T}^n_{\leq d}$. Then the subset selection criterion $\mathrm{LT}_\sigma(w) \in \mathcal{L}$ and σ being degree-compatible yield $\operatorname{Supp}(w) \subseteq \mathbb{T}^n_{\leq d}$. Thus, for $\mathcal{L} \subseteq \mathbb{T}^n_{\leq d}$ at the beginning of the first iteration, all linear basis extensions take place in the finite-dimensional space $\langle \mathbb{T}^n_{\leq d} \rangle_K$.

At termination of this loop, we have $\langle \mathcal{V} \rangle_K = \langle \mathcal{V} \rangle_K^+ \cap \langle \mathcal{L} \rangle_K$ due to the following identities. The basis extension in step (I3) gives $\langle \mathcal{V} \cup \mathcal{W}' \rangle_K = \langle \mathcal{V} \rangle_K^+$. Since we passed the subloop (I4)–(I5), we have $\langle \mathcal{V} \cup \mathcal{W} \rangle_K = \langle \mathcal{V} \cup \mathcal{W}' \rangle_K \cap \langle \mathcal{L} \rangle_K$.

Finally, when exiting the subloop in step (I6), we have $\mathcal{W} = \emptyset$ and hence $\langle \mathcal{V} \rangle_K = \langle \mathcal{V} \cup \mathcal{W} \rangle_K$.

Now, we show that the definition of $\mathcal{O}$ in step (I9) produces an order ideal. Analogously to the proof of Proposition 15 let $t \in \mathcal{L} \setminus \mathcal{O}$ and consider the case $x_i\, t \in \mathcal{L}$. Since $t \notin \mathcal{O}$ there is a $v \in \mathcal{V}$ with $t = \mathrm{LT}_\sigma(v)$. We have $x_i\, v \in \mathcal{V}^+$ and by case consideration $\mathrm{LT}_\sigma(x_i v) = x_i t \in \mathcal{L}$, i.e. $x_i v \in \mathcal{V} \cup \mathcal{W}$. Having passed the subloop (I4)–(I5), we infer $\mathrm{Supp}(x_i v) \subseteq \mathcal{L}$. Thus $x_i v \in \langle \mathcal{V} \rangle_K^+ \cap \langle \mathcal{L} \rangle_K$. By the argument in the preceding paragraph this intersection equals $\langle \mathcal{V} \rangle_K$. As the leading terms of $\mathcal{V}$ are pairwise different, $\mathrm{LT}_\sigma(x_i v) \in \mathrm{LT}_\sigma\{\mathcal{V}\}$. This shows $x_i\, t \in \mathcal{L} \setminus \mathcal{O}$.

The loop of steps (I3)–(I8) terminates because, with each call of a new iteration in step (I8), the universe $\mathcal{L}$ becomes strictly larger. Unless the loop has terminated before, eventually the universe becomes sufficiently large to contain the polynomials (2) as in the proof of the original Border Basis Algorithm. By the same argument as there, the loop terminates.

This covers all changes in the algorithm. □

The following example shows what kind of improvement can be expected.

Example 22 For comparison we apply the Improved Border Basis Algorithm to the set of generators stated in Example 20. The algorithm starts with the universe $\{1, z, z^2, y, yz, x, xz, y^2, xy, x^2, x^3, x^4, x^5\}$ consisting of 13 terms. The first basis extension produces a nonempty $\mathcal{W}'$, but the restriction to elements with leading term in the universe leads to $\mathcal{W} = \emptyset$ and to the order ideal $\mathcal{O} = \{1, z, y, x, x^2, x^3, x^4\}$. The border $\partial\mathcal{O}$ is not contained in the universe and hence we enlarge the universe. From now on we are working in a universe with 29 terms. Next, four linear basis extensions are computed and we obtain again the order ideal $\mathcal{O}$ as above. Of course, this time its border is contained in the universe and the border basis is computed.

So, instead of computing four linear basis extensions and the final reduction in a 56-dimensional space (cf. Example 20), the Improved Border Basis Algorithm computes one extension in a 13-dimensional space as well as four extensions and the final reduction in a 29-dimensional space.

We can do even better. In step (I8) of the Improved Border Basis Algorithm we enlarge the universe $\mathcal{L}$ more than is required to fit in $\partial\mathcal{O}$.

Corollary 23 Let N be a positive integer. Replace step (I8) in the Improved Border Basis Algorithm with

(I8') If $\partial\mathcal{O} \not\subseteq \mathcal{L}$ then replace $\mathcal{L}$ with the order ideal spanned by $\mathcal{L}$ and $\partial\mathcal{O}$; every Nth time this is done, replace $\mathcal{L}$ with $\mathcal{L}^+$ instead. Continue with step (I3).

The instructions (I1)–(I7), (I8'), (I9) form an algorithm that computes a border basis $g_1, \ldots, g_\nu$.

We included the $\mathcal{L}^+$ replacement every Nth time as a safeguard. In that way the above termination argument also applies to the loop (I3)–(I8'). Without this there is a theoretical chance that the modified procedure may run in an infinite loop: potentially, universe enlargements to accomodate $\partial\mathcal{O}$ may always act in the x-direction while the wanted reduction information is along the y-direction. This problematic behaviour is avoided in the Improved Border Basis Algorithm, since the enlargement $\mathcal{L}^+$ lets the universe grow in all directions. However, we have not met this problematic behaviour in any of the examples computed and we strongly believe that the use of (I8') without the safeguard does not produce termination problems. In other words, our N is huge.

Using this primed version of the Improved Border Basis Algorithm we compute the preceding example once more. Of course, we start with the same universe of 13 terms as before and compute one linear basis extension that leads to $\mathcal{W} = \emptyset$. Now, the enlargement leads to a universe with only 19 terms, in which four linear basis extensions are computed until the border basis is found.

References

[1] W. Auzinger, H. J. Stetter. An elimination algorithm for the computation of all zeros of a system of multivariate polynomial equations, in: R.G. Agarwal, Y.M. Chow, S.J. Wilson (eds.), Int. Conf. on Numerical Mathematics, Singapore 1988, Birkhäuser ISNM **86**, Basel 1988, pp. 11–30.

[2] W. Bruns, J. Herzog. *Cohen-Macaulay Rings*, Cambridge University Press, Cambridge 1993.

[3] J.C. Faugère, P. Gianni, D. Lazard, T. Mora. Efficient computation of zero-dimensional Gröbner bases by change of ordering, *Journal of Symbolic Computation* **16** (1993), pp. 329-344.

[4] J.C. Faugère. A new efficient algorithm for computing Gröbner bases (F_4), *Journal of Pure and Applied Algebra* **139** (1999), pp. 61–88.

[5] J.C. Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F_5), in: T. Mora (ed.), Proc. Conf. ISSAC 2002, ACM Press, New York 2002, pp. 75–83.

[6] A. Kehrein, M. Kreuzer. Characterizations of border bases, *Journal of Pure and Applied Algebra* **196** (2005), pp. 251–270.

[7] A. Kehrein, M. Kreuzer, L. Robbiano. An algebraist's view on border bases, in: I. Emiris and A. Dickenstein (eds.). *Solving Polynomial Equations*, Alg. and Comp. in Math. **14**, Springer, Heidelberg 2005, pp. 169-202.

[8] M. Kreuzer, L. Robbiano. *Computational Commutative Algebra 1*, Springer, Heidelberg 2000.

[9] H. M. Möller. Systems of algebraic equations solved by means of endomorphisms, in: G. Cohen *et al.* (eds.), Applied algebra, algebraic algorithms and error-correcting codes, Proc. Conf. AAECC-10, LNCS **673**, Springer, Heidelberg 1993, pp. 43–46.

[10] B. Mourrain. A new criterion for normal form algorithms, in: M. Fossorier, H. Imai, S. Lin, A. Poli (eds.), Proc. Conf. AAECC-13, Honolulu 1999, LNCS **1719**, Springer, Heidelberg 1999, pp. 440–443.

[11] H. J. Stetter. *Numerical polynomial algebra*, SIAM, Philadelphia 2004.

[12] B. Sturmfels. *Gröbner Bases and Convex Polytopes*, AMS, University Lecture series **8**, Providence R.I. 1996.