On Affine Tropical F5 Algorithms

Tristan Vaccon[a], Thibaut Verron[b], Kazuhiro Yokoyama[c]

[a] *XLIM, Université de Limoges, France*
[b] *Institute for Algebra, JKU, Linz, Austria*
[c] *Departement of Mathematics, Rikkyo University, Japan*

Abstract

Let K be a field equipped with a valuation. Tropical varieties over K can be defined with a theory of Gröbner bases taking into account the valuation of K. Because of the use of the valuation, the theory of tropical Gröbner bases has proved to provide settings for computations over polynomial rings over a p-adic field that are more stable than that of classical Gröbner bases.

Beforehand, these strategies were only available for homogeneous polynomials. In this article, we extend the F5 strategy to a new definition of tropical Gröbner bases in an affine setting. We also provide a competitor with an adaptation of the F4 strategy to tropical Gröbner bases computations.

We provide numerical examples to illustrate time-complexity and p-adic stability of this tropical F5 algorithm. We also illustrate its merits as a first step before an FGLM algorithm to compute (classical) lex bases over p-adics.

Keywords: Algorithms, Tropical geometry, Gröbner bases, F5 algorithm, p-adic precision

1. Introduction

Tropical geometry as we understand it has not yet reached half a century of age. It has nevertheless spawned significant applications to very various domains, from algebraic geometry to combinatorics, computer science, economics, non-archimedean geometry (see [MS15], [EKL06]) and even attempts at proving the Riemann hypothesis (see [C15]).

Effective computation over tropical varieties makes decisive use of Gröbner bases. Since Chan and Maclagan's definition of tropical Gröbner bases taking into account the valuation in [C13, CM13], computations of tropical Gröbner bases are available over fields with trivial or non-trivial valuation, but only in the context of homogeneous ideals.

Email addresses: `tristan.vaccon@unilim.fr` (Tristan Vaccon), `thibaut.verron@jku.at` (Thibaut Verron), `kazuhiro@rikkyo.ac.jp` (Kazuhiro Yokoyama)

URL: `www.unilim.fr/pages_perso/tristan.vaccon/` (Tristan Vaccon), `https://thibautverron.github.io/index.html` (Thibaut Verron)

On the other hand, for classical Gröbner bases, numerous algorithms have been developed allowing for more and more efficient computations. The latest generation of algorithms for computing Gröbner bases is the family of signature-based algorithms, which keep track of where the polynomials come from in order to anticipate useless reductions. This idea was initiated in Algorithm F5 [F02], and has since then been widely studied and generalized ([BFS14, EF17]).

Most of those algorithms, including the original F5 algorithm, are specifically designed for homogeneous systems, and adapting them to affine (or inhomogeneous) systems requires special care (see [E13]).

An F5 algorithm computing tropical Gröbner bases without any trivial reduction to 0, inspired by the classical F5 algorithm, has been described in [VY17]. The goal of this paper is to extend the definition of tropical Gröbner bases to inhomogeneous ideals, and describe ways to adapt the F5 algorithm in this new setting.

The core motivation is the following. It has been proved in [V15] that computing tropical Gröbner bases, taking into account the valuation, is more stable for polynomial ideals over a p-adic field than classical Gröbner bases.

Thus, an affine variant of tropical Gröbner bases is highly desirable to handle non-homogeneous ideals over p-adics. For classical Gröbner bases, it is always possible to homogenize the input ideal, compute a homogeneous Gröbner basis, and dehomogenize the result. This technique is not always optimal, because the algorithm may end up reaching a higher degree than needed, computing points at infinity of the system, but it always gives a correct result and, in the case of signature Gröbner basis algorithms, is able to eliminate useless reductions. However, in a tropical setting, terms are ordered with a tropical term order, taking into account the valuation of the coefficients. As far as we know, there is no way to dehomogenize a system in a way that would preserve the tropical term order. Indeed, no tropical term order can be an elimination order.

Moreover, the FGLM algorithm can be adapted to the tropical case (see Chap. 9 of [V*]), making it possible to compute a lexicographical (classical) Gröbner basis from a tropical one. We provide numerical data to estimate the loss in precision for the computation of a lex Gröbner basis using a tropical F5 algorithm followed by an FGLM algorithm, in an affine setting.

Another classical strategy to compute Gröbner bases is Faugère's F4 algorithm [F99]. We provide a competitor to the F5 strategy with a tropical adaptation of F4, along with numerical data to compare their respective merits.

Up to our knowledge, this is the first study of an F4 algorithm in a tropical context. Regarding to computation at finite precision, there is a strong motivation for developping F4 algorithms along with F5 algorithms. In F4 algorithms, there is considerably more freedom in the choice of pivots during reduction. With more choices of reductors, a trade-off between speed and loss of precision is possible.

1.1. Related works

A canonical reference for an introduction to computational tropical algebraic geometry is the book of Maclagan and Sturmfels [MS15].

The computation of tropical varieties over $\mathbb{Q}$ with trivial valuation is available in the Gfan package by Anders Jensen (see [Gfan]), by using standard Gröbner bases computations. Chan and Maclagan have developed in [CM13] a Buchberger algorithm to compute tropical Gröbner bases for homogeneous input polynomials (using a special division algorithm). Following their work, still for homogeneous polynomials, a Matrix-F5 algorithm has been proposed in [V15] and a Tropical F5 algorithm in [VY17]. Markwig and Ren have provided a completely different technique of computation using projection of standard bases in [MR16], again only for homogeneous input polynomials.

In the classical Gröbner basis setting, many techniques have been studied to make the computation of Gröbner bases more efficient. In particular, Buchberber's algorithm is frequently made more efficient by using the *sugar-degree* (see [GMNRT91, BCM11]) instead of the actual degree for selecting the next pair to reduce. This technique was a precursor of modern signature techniques, in the sense that the sugar-degree of a polynomial is exactly the degree of its signature. General signature-based algorithms for computing classical Gröbner bases of inhomogeneous ideals have been extensively studied in [E13].

A shorter version of this article has been published in the Proceedings of the 43th International Symposium on Symbolic and Algebraic Computation (ISSAC 2018) with [VVY18]§. The main additions are the exposition of the results and proofs in Section 3 and 4, which have been improved and extended, and a new section with Section 7, introducing an F4 algorithm to compute tropical GB.

1.2. Specificities of computating tropical GB

The computation of tropical GB, even by a Buchberger-style algorithm, is not as straightforward as for classical Gröbner bases. One way to understand it is the following: even for homogeneous ideals, there is no equivalence between tropical Gröbner bases and row-echelon linear bases at every degree. Indeed,we can remark that $(f_1, f_2) = (x + y, 2x + y)$ is a tropical GB over $\mathbb{Q}[x, y]$ with 2-adic valuation, $w = [0, 0]$ and grevlex ordering. Nevertheless, the corresponding 2×2 matrix in the vector space of homogeneous polynomials of degree 2 is not under tropical row-echelon form.

As a consequence, reduction of a polynomial by a tropical GB is not easy. In [C13, CM13], Chan and Maclagan relied on a variant of Mora's tangent cone algorithm to obtain a division algorithm. In [V15, VY17], the authors relied on linear algebra and the computation of (tropical) row-echelon form. In this article, we extend their method to the computation of tropical Gröbner bases in an affine setting, through an F5 algorithm.

1.3. Main idea and results

Extending the tropical F5 algorithm to inhomogeneous inputs poses two difficulties. First, as mentioned, tropical Gröbner bases used to be only defined and computed for homogeneous systems. Even barebones algorithms such as Buchberger's algorithm are not available for inhomogeneous systems. The second problem is a general problem of signature Gröbner bases with inhomogeneous input. The idea of signature algorithms is to compute polynomials with increasing signatures, and the F5 criterion detects trivial reductions to 0 by matching candidate signatures with existing leading terms. For homogeneous ideals, the degree of the signature of a polynomial and the degree of the polynomial itself are correlated. This is what makes the F5 criterion applicable.

The survey paper [E13] has shown that Algorithm F5, using the *position over term* ordering on the signatures, has to reach a tradeoff between eliminating all reductions to 0 and performing other useless reductions.

More precisely, let $f_1, \dots, f_m$ be homogeneous polynomials with coefficients in a field with valuation K, and define $I_{k,d}$ the vector space of polynomials in $\langle f_1, \dots, f_k \rangle$ with degree at most d. With the usual computational strategy, the algorithm computes a basis of $I_{1,1}$, then $I_{2,1}$, and so on until $I_{m,1}$, and then $I_{1,2}$, and so on. In a lot of situations ideals with more generators have a Gröbner basis with lower degree [BFS04], and this strategy ensures that the algorithm does not reach a degree higher than needed.

However, the same algorithm for affine system will, at each step, merely compute a *set of polynomials* in each $I_{k,d}$. This set needs not be a generating set because of degree falls. To obtain a basis instead, one has to proceed up to some $I_{k,\delta}$ with $\delta \geq d$. When $\delta > d$, some polynomials will be missing for the F5 criterion in degree less than δ, and the corresponding trivial reductions to 0 will not be eliminated.

In this paper, we show that the tropical F5 algorithm [VY17] works in an affine setting, and we characterize those trivial reductions to 0 which are eliminated by the F5 criterion. In particular, we show that the Macaulay matrices built at each step of the computations are Macaulay matrices of all polynomials with a given *sugar-degree*.

Compared to [VY17], the overall presentation of the F5 algorithms is clarified. It can now be summarized as the following strategy: filtration, signature, F5 elimination criterion, Buchberger-F5 criterion and finally the F5 algorithm.

Theorem 1.1. *Given a set of (non-necessarily homogeneous) polynomials $f_1, \dots, f_m \in K[X_1, \dots, X_n]$, the Tropical F5 algorithm (Algorithm 3) computes a tropical Gröbner basis of I, without reducing to 0 any trivial tame syzygy (Def. 3.1).*

We also examine an incremental affine version of the homogeneous tropical F5-algorithm and an affine tropical F4, and we compare their performances on several examples. Even in a non-homogeneous setting, the loss in precision of the tropical F5 algorithm remains satisfyingly low.

1.4. Organization of the paper

Section 2 introduces notations and nonhomogeneous tropical Gröbner bases. Section 3 then introduces the filtration on ideals necessary for F5 algorithms in this context. Section 4 provides a Buchberger-F5 criterion on which Section 5 elaborates a first tropical F5 algorithm. Section 6 briefly presents another variant of the F5 algorithm to compute nonhomogeneous tropical Gröbner bases, and Section 7 introduces a tropical adaptation of the F4 algorithm. Finally, Section 8 displays numerical results related to the precision behaviour and time-complexity of all the algorithms we have described.

Acknowledgements

We thank Jean-Charles Faugère, Marc Mezzarobba, Pierre-Jean Spaenlehauer, Masayuki Noro, and Naoyuki Shinohara for fruitful discussions.

2. Affine Tropical GB

2.1. Notations

Let k be a field with valuation val. The polynomial ring $k[X_1, \dots, X_n]$ will be denoted by A. Let T be the set of monomials of A. For $u = (u_1, \dots, u_n) \in \mathbb{Z}^n_{\geq 0}$, we write x^u for $X_1^{u_1} \dots X_n^{u_n}$ and $|f|$ for the degree of $f \in A$. In A^s, $(e_i)_{i=1}^s$ is the canonical basis.

The matrix of a list of polynomials written in a basis of monomials is called a *Macaulay matrix.*

Given a mapping $\phi : U \to V$, $Im(\phi)$ denotes the image of ϕ. For a matrix M, $Rows(M)$ is the list of its rows, and $Im(M)$ denotes the left-image of M (*i.e.* $Im(M) = span(Rows(M))$). For $w \in Im(val)^n \subset \mathbb{R}^n$ and $\leq_1$ a monomial order on A, we define $\leq$ a tropical term order as in the following definition:

Definition 2.1. Given $a, b \in k^*$ and x^α and x^β two monomials in A, we write $ax^\alpha < bx^\beta$ if:

- $|x^\alpha| < |x^\beta|$, or
- $|x^\alpha| = |x^\beta|$, and $val(a)+w\cdot\alpha > val(b)+w\cdot\beta$, or $val(a)+w\cdot\alpha = val(b)+w\cdot\beta$ and $x^\alpha <_1 x^\beta$.

For u of valuation 0, we write $ax^\alpha =_\leq uax^\alpha$. Accordingly, $ax^\alpha \leq bx^\beta$ if $ax^\alpha < bx^\beta$ or $ax^\alpha =_\leq bx^\beta$.

Throughout this article, we are interested in computing a tropical Gröbner basis of $I = \langle f_1, \dots, f_s \rangle$ for some given $f_1, \dots, f_s \in A$ (ordered increasingly by degree).

2.2. Tropical GB

A tropical term order provides an order on the terms of the polynomials $f \in A$.

Definition 2.2. For $f \in A$, we define $LT(f)$ to be the biggest term of f. We define $LM(f)$ to be the monomial corresponding to $LT(f)$ and $LC(f)$ the corresponding coefficient.

We define $LM(I)$ to be the monomial ideal generated by the monomials $LM(f)$ for $f \in I$.

We can then naturally define what is a tropical Gröbner basis (*tropical GB* for short):

Definition 2.3. $G \subset I$ is a tropical GB of I if span($LM(g)$ for $g \in G$) $= LM(I)$.

We can remark that for homogeneous polynomials this definition coincides with the one given in [VY17].

3. Filtration and $\mathfrak{S}$-GB

3.1. Definition and elimination criterion

One of the main ingredients for F5 algorithms is the definition of a vector space filtration of the ideal I. It is defined from the initial polynomials $F = (f_1, \dots, f_s)$ generating I. For simplicity, we assume that they are ordered by increasing degree.

First, we extend $\leq$ to the monomials of the vector space A^s. To that intent, we highlight some monomials which appear as leading monomial of a syzygy.

Definition 3.1. Let $(a_1, \dots, a_s) \in A^s$ and $i \in \{1, \dots, s\}$ be such that:

1. $\sum_j a_j f_j = 0$.
2. $a_i \neq 0$ and $a_j = 0$ for $j > i$.
3. for all $j < i$, $|a_j f_j| \leq |a_i f_i|$.

We call such a syzygy a *tame syzygy* and we define $LM(a_i)e_i$ to be its leading monomial. We define $LM(TSyz(F))$ as the module in A^s generated by the leading monomials of the tame syzygies. Trivial tame syzygies are the tame syzygies that are also trivial (*i.e.* in the module generated by the $f_i e_j - f_j e_i$).

The F5 criterion that we use in this article is designed to recognize some of the tame syzygies and use this knowledge to avoid useless reduction to zero of some polynomials. It is the main motivation for defining a filtration on the vector space A^s. We use a degree-refining monomial ordering $\leq_m$ on A.[1] We define a total order on the monomials of A^s.

Definition 3.2. We write that $x^\alpha e_i \leq_{sign} x^\beta e_j$ if:

1. if $i < j$, or
2. if $i = j$ and $|x^\alpha f_i| < |x^\beta f_j|$, or
3. if $i = j$ and $|x^\alpha f_i| = |x^\beta f_j|$, and
    - $x^\alpha \notin LM(TSyz(F))$ and $x^\beta \in LM(TSyz(F))$, or
    - both $x^\alpha, x^\beta \in LM(TSyz(F))$ and $x^\alpha \leq_m x^\beta$, or
    - both $x^\alpha, x^\beta \notin LM(TSyz(F))$ and $x^\alpha \leq_m x^\beta$.

[1] $\leq_m$ is not necessarily related to $\leq_1$ and $\leq$.

Definition 3.3. We consider the vector space

$$I_{\leq_{sign} x^\alpha e_i} := Span(\{x^\beta f_j,\ \text{s.t. } x^\beta e_j \leq_{sign} x^\alpha e_i\})$$

and the vector space $I_{<_{sign} x^\alpha e_i}$ defined accordingly. We define $I = \bigcup_{\uparrow x^\alpha e_i} I_{\leq_{sign} x^\alpha e_i}$ as an increasing vector space filtration of I.

We then have a very natural definition of signature. In litterature, this notion of signature is sometimes called *minimal signature*.

Definition 3.4. For $f \in I$, the smallest $x^\alpha e_i$ such that $f \in I_{\leq_{sign} x^\alpha e_i}$ is called the **signature** of f and noted $S(f)$.

The degree $|x^\alpha f_i|$ is called the **sugar-degree** of $x^\alpha e_i$. [2] For the purpose of Algorithm 3, we design a filtration compatible with the sugar-degree.

Example 3.5. Let $F = (x^2y+1, x^3+1, y^3-x, y^3-y) \in \mathbb{Q}_2[x,y]$ with $w = (0,0)$. We use $\leq_m$ and $\leq_1$ to be the graded lexicographical ordering. Then $S(x-y) = ye_2$.

Definition 3.6. We consider the vector space

$$I^{\leq d} = Span(\{x^\beta f_j,\ \text{s.t. } |x^\beta e_j| \leq d\})$$

We then define, for $x^\alpha e_i$ with sugar-degree d, the vector space
$I^{\leq d}_{\leq_{sign} x^\alpha e_i} = Span(\{x^\beta f_j,\ \text{s.t. } x^\beta e_j \leq_{sign} x^\alpha e_i \text{ and } |x^\beta f_j| \leq d\})$.

$I = \bigcup_{\uparrow d} I^{\leq d}$ is also a vector space filtration. $I^{\leq d}$ can itself be filtrated by the $I^{\leq d}_{\leq_{sign} x^\alpha e_i}$. We have a compatible notion of signature:

Definition 3.7. For $d \in \mathbb{Z}_{>0}$ and $f \in I^{\leq d}$, the smallest $x^\alpha e_i$ such that $f \in I^{\leq d}_{\leq_{sign} x^\alpha e_i}$ is called the d-**signature** of f and noted $S_d(f)$.

Remark 3.8. $S_d(f)$ can be different from $S(f)$ for small d, but given an f, all $S_d(f)$ are equal when d is large. Moreover, for any $f \in I^{\leq d}$, $S(f) \leq_{sign} S_d(f)$.

Example 3.9. We set $F = (x^2y+1, x^3+1, y^3-x, y^3-y) \in \mathbb{Q}_2[x,y]$ as in Example 3.5. Then $S_3(x-y) = e_4$ and $S_4(x-y) = ye_2 = S(x-y)$.

The main motivation for defining the vector spaces $I^{\leq d}_{\leq_{sign} x^\alpha e_i}$ is their finite dimension. Their compatibility with the sugar-degree allows the F5 algorithm to compute only one Macaulay matrix per sugar-degree d.

The goal of the F5 criterion is to recognize some $x^\alpha e_i$ such that the filtration is constant at $I^{\leq d}_{\leq_{sign} x^\alpha e_i}$. As a consequence, this knowledge allows to skip some calculation as, because of this constancy, they will not provide any new leading monomial. We can then state a first version of the F5 elimination criterion:

[2] Sugar-degree has been introduced and explored in [GMNRT91, BCM11].

Proposition 3.10 ([F02])**.** *If x^α is such that $x^\alpha e_i \in LM(Tsyz(F))$, $d \geq |x^\alpha f_i|$, then the filtration is constant at $I^{\leq d}_{\leq_{sign} x^\alpha e_i}$. If $x^\alpha \in LM(I^{\leq d}_{\leq_{sign} x^\beta e_j})$ for some x^β and j such that $|x^\beta f_j| \leq |x^\alpha|$, then $x^\alpha e_i \in LM(Tsyz(F))$ for any $i > j$.*

Proof. For the first part, we can write $(x^\alpha + g) f_i = \sum_{j<i} a_j f_j$, with $LT(g) < x^\alpha$ and for all $j < i$, $|a_j f_j| \leq |x^\alpha f_i|$. Then:

$$x^\alpha f_i = (-g) f_i + \sum_{j=1}^{i-1} a_j f_j.$$

By linear algebra and a complete reduction using as pivot the $x^\beta e_j \in LM(Tsyz(F))$, we can assume that g has no monomial in $LM(TSyz(F))$ and obtain: $x^\alpha f_i \in I^{\leq d}_{<x^\alpha e_i}$, and therefore, the filtration is constant at $I^{\leq d}_{\leq x^\alpha e_i}$.

For the second part, we can write $x^\alpha + g = \sum_{k \leq j} a_k f_k$, with $LT(g) < x^\alpha$ and for all $k \leq j$, $|a_k f_k| \leq |x^\beta f_j| \leq |x^\alpha|$. Then $(x^\alpha + g) f_i - \sum_{k \leq j} (a_k f_i) f_k = 0$ and we do have $|x^\alpha f_i| \geq |(a_k f_i) f_k|$ for all $k \leq j$. □

If all the f_i's are homogeneous, this coincides with the usual F5 elimination criterion, as for example stated in [VY17], which eliminates all trivial reductions to zero in the course of the algorithm. For affine polynomials, the F5 criterion only eliminates those trivial reductions which are tame.

3.2. Tropical $\mathfrak{S}$-GB

In order to take advantage of the F5 elimination criterion to compute tropical Gröbner bases, we focus on the computation of tropical Gröbner bases which are compatible with the filtration: tropical $\mathfrak{S}$-GB. We first need the definition of reductions compatible with the filtration and the corresponding irreducible polynomials.

Definition 3.11 ($\mathfrak{S}$-reduction)**.** Let $e, g \in I$, $h \in I$. We say that e **$\mathfrak{S}$-reduces** to g with h, $e \rightarrow^h_{\mathfrak{S}} g$, if there are $t \in T$ and $\alpha \in k^*$ such that:

- $LT(g) < LT(e)$, $LM(g) \neq LM(e)$ and $e - \alpha t h = g$ and
- $S(th) <_{sign} S(e)$.

We should remark that in the previous definition, we ask for the condition $LM(g) \neq LM(e)$ so as to ensure that the leading monomial of e is eliminated. For example, reducing x^2 by $x + y$ into $2x^2 + xy$ would never be acceptable.

We can naturally define then what is an $\mathfrak{S}$-irreducible polynomial.

Definition 3.12 ($\mathfrak{S}$-irreducible polynomial)**.** We say that $g \in I$ is $\mathfrak{S}$-irreducible if there is no $h \in I$ which $\mathfrak{S}$-reduces g. If there is no ambiguity, we might omit the $\mathfrak{S}$-.

Example 3.13. We use $F = (x^2 y + 1, x^3 + 1, y^3 - x, y^3 - y) \in \mathbb{Q}_2[x, y]$ as in Example 3.5. Then $x^3 y - y$ $\mathfrak{S}$-reduces to $x - y$, which is irreducible (its signature is $y e_2$).

Definition 3.14 (Tropical $\mathfrak{S}$-Gröbner basis)**.** We say that $G \subset I$, a set of $\mathfrak{S}$-irreducible polynomials, is a **tropical $\mathfrak{S}$-Gröbner basis** (or tropical $\mathfrak{S}$-GB, or just $\mathfrak{S}$-GB for short when there is no ambiguity) of I with respect to a given tropical term order, if for each $\mathfrak{S}$-irreducible polynomial $h \in I$, there exist $g \in G$ and $t \in T$ such that $LM(tg) = LM(h)$ and $tS(g) = S(h)$.

Definition 3.15. Definitions 3.11, 3.12 and 3.14 have natural analogues when one restricts to the vector space $I^{\leq d}$ and S_d with $\mathfrak{S}_d$-reduction, $\mathfrak{S}_d$-irreducible polynomial and tropical $\mathfrak{S}_d$-GB.

Example 3.16. We use $F = (x^2y+1, x^3+1, y^3-x, y^3-y) \in \mathbb{Q}_2[x,y]$ as in Example 3.5. Then $G = (x^2y+1, x^3+1, x-y, y^3-1, y-1)$, with respective signatures $(e_1, e_2, ye_2, xy^2e_2, e_3, e_4)$, is a tropical $\mathfrak{S}$-GB of I.

Proposition 3.17. *If G is a tropical $\mathfrak{S}$-Gröbner basis, then for any nonzero $h \in I$, there exist $g \in G$ and $t \in T$ such that:*

- $LM(tg) = LM(h)$
- $S(tg) = tS(g) = S(h)$ *if h is irreducible, and* $S(tg) = tS(g) <_{sign} S(h)$ *otherwise.*

Hence, there is an $\mathfrak{S}$-reductor for h in G when h is not irreducible.

Proof. For h irreducible, the result is simply a consequence of the definition of a tropical $\mathfrak{S}$-GB. When h is reducible, it means that it can be reduced by h_2 a polynomial with same leading term but such that $S(h_2) <_{sign} S(h)$. If h_2 is irreducible, we are done. Otherwise, there exists h_3 with same leading term and such that $S(h_3) <_{sign} S(h_2)$, and so on. A consequence of Definition 3.7 is that there is no strictly decreasing sequence of monomials for $\leq_{sign}$. Consequently, after a finite number of iterations, we find h' with same leading term as h, irreducible and such that $S(h') <_{sign} S(h)$. We can then apply the definition of a tropical $\mathfrak{S}$-GB to conclude. □

Corollary 3.18. *If G is a tropical $\mathfrak{S}$-Gröbner basis, then G is a tropical Gröbner basis of I, for $<$.*

As a consequence computing a tropical $\mathfrak{S}$-GB provides a tropical GB, and we can use the F5 elimination criterion 3.10 to our advantage when computing these tropical $\mathfrak{S}$-GB. Moreover, we also have the following finiteness result:

Proposition 3.19. *Every tropical $\mathfrak{S}$-Gröbner basis contains a finite tropical $\mathfrak{S}$-Gröbner basis.*

Proof. It is a property of the structure of monomial ideals, and the proof in the classical case from [AP, Prop. 14], using Dickson's lemma can directly be transposed to the tropical setting. □

3.3. Linear algebra and existence

For $x^\alpha \in T$ and $1 \leq i \leq n$, let us denote by $Mac_{\leq_{sign} x^\alpha e_i}(F)$ the Macaulay matrix of the polynomials $x^\beta f_j$ such that $x^\beta f_j \leq x^\alpha f_i$, ordered increasingly for the order on the $x^\beta e_j$'s. One can perform a tropical LUP algorithm on $Mac_{\leq d}(F)$ (see Algo. 2) and obtain all the leading monomials in $I_{\leq_{sign} x^\alpha e_i}$. This can be (theoretically) performed for all $x^\alpha e_i$ to obtain the existence of an $\mathfrak{S}$-GB of I.

3.4. More on signatures

We define Σ to be the set of signatures.

Thanks to Proposition 3.10, not all $x^\alpha e_i$ can be a signature:

Remark 3.20. If $x^\alpha e_i \in LM(TSyz(F))$ then $x^\alpha e_i \notin \Sigma$.

We provide three lemmata to understand the compatibility of Σ with basic operations on polynomials.

Lemma 3.21. *If $f, g \in I$ are such that $S(f) = S(g)$ and $LM(f) \neq LM(g)$, then there exist $a, b \in k^*$ such that $S(af + bg) < S(f)$ and $af + bg \neq 0$.*

If one takes the point of view of linear algebra, the proof is direct. This Lemma has a direct consequence on $\mathfrak{S}$-irreducible polynomials.

Lemma 3.22. *All $\mathfrak{S}$-irreducible polynomials of a given signature share the same leading monomial.*

We finally enounce one last lemma about the compatibility of signatures with multiplication.

Lemma 3.23. *If $g \in I$ and $\tau \in T$ then $S(\tau g) \leq \tau S(g)$. If moreover $\tau S(g) \in \Sigma$, then $S(\tau g) = \tau S(g)$ and for all $\mu \in T$ such that μ divides τ, $S(\mu g) = \mu S(g)$.*

Proof. The first part is direct. For the second part, one can show that it is possible to write that $\tau g = h + r$ for some $h \in I$ of signature $\tau S(g)$, irreducible, and $r \in I_{<_{sign} \tau S(g)}$ and conclude that $S(\tau g) = \tau S(g)$.

For the last statement, assume that there exists a $\mu \in T$ dividing τ such that $S(\mu g) < \mu S(g)$. Then $S(\tau g) = S(\frac{\tau}{\mu} \mu g) \leq \frac{\tau}{\mu} S(\mu g) < \frac{\tau}{\mu} \mu S(g) = \tau S(g)$, which is a contradiction. $\square$

4. Buchberger-F5 criterion

In this section, we explain a criterion, the Buchberger-F5 criterion, on which we build our F5 algorithm to compute tropical $\mathfrak{S}$-Gröbner bases. It is an analogue of the Buchberger criterion which includes the F5 elimination criterion.

We need a slightly different notion of S-pairs, called here normal pairs and can then state the Buchberger-F5 criterion.

Definition 4.1 (Normal pair). Given $g_1, g_2 \in I$, let $Spol(g_1, g_2) = u_1 g_1 - u_2 g_2$ be the S-polynomial of g_1 and g_2, with for $i \in \{1, 2\}$, $u_i = \frac{lcm(LM(g_1), LM(g_2))}{LT(g_i)}$. We say that (g_1, g_2) is a **normal pair** if:

1. the g_i's are $\mathfrak{S}$-irreducible polynomials.
2. $S(u_i g_i) = LM(u_i) S(g_i)$ for $i = 1, 2$.
3. $S(u_1 g_1) \neq S(u_2 g_2)$.

We define accordingly d-normal pairs in $I^{\leq d}$.

Theorem 4.2 (Buchberger-F5 criterion)**.** *Suppose that G is a finite set of $\mathfrak{S}$-irreducible polynomials of $I = \langle f_1, \ldots, f_s \rangle$ such that:*

1. for all $\forall i \in [\![1, s]\!]$, there exists $g \in G$ such that $S(g) = e_i$.

2. for any $g_1, g_2 \in G$ such that (g_1, g_2) is a normal pair, there exists $g \in G$ and $t \in T$ such that tg is $\mathfrak{S}$-irreducible and $tS(g) = S(tg) = S(Spol(g_1, g_2))$.

Then G is a $\mathfrak{S}$-Gröbner basis of I. The analogue result using d-normal pairs to recognize an $\mathfrak{S}_d$-GB in $I^{\leq d}$ is also true.

Remark 4.3. The converse of this result is clear.

Theorem 4.2 is an analogue of the Buchberger criterion for tropical $\mathfrak{S}$-Gröbner bases. To prove it, we adapt the classical proof of the Buchberger criterion and the proof of the tropical Buchberger algorithm of Chan and Maclagan (Algorithm 2.9 of [C13]). First we need two lemmata.

Lemma 4.4. *Let $x^\alpha, x^\beta, x^\gamma, x^\delta \in T$ and $P, Q \in A$ be such that $LM(x^\alpha P) = LM(x^\beta Q) = x^\gamma$ and $x^\delta = lcm(LM(P), LM(Q))$. Then*

$$Spol(x^\alpha P, x^\beta Q) = x^{\gamma-\delta} Spol(P, Q).$$

Proof. Exercice 8 in Section 2.6 of [CLO]. □

Lemma 4.5. *Let G be an $\mathfrak{S}$-Gröbner basis of I up to some signature σ. Let $h \in I$, be such that $S(h) \leq \sigma$. Then there exist $r \in \mathbb{N}$, $g_1, \ldots, g_r \in G$, $Q_1, \ldots, Q_r \in A$ such that for all i and x^α a monomial of Q_i, $S(x^\alpha g_i) = x^\alpha S(g_i) \leq S(h)$ and $LT(Q_i g_i) \leq LT(h)$, the $x^\alpha S(g_i)$'s are all distinct and non-zero, and, finally, we have*

$$h = \sum_{i=1}^{r} Q_i g_i.$$

Proof. It is clear by linear algebra. One can form a Macaulay matrix whose rows correspond to polynomials τg with $\tau \in T, g \in G$ such that $S(\tau g) = \tau S(g) \leq S(h)$. Only one row is possible per non-zero signature, and each monomial in $LM(I_{\leq \sigma})$ is reached as leading term by only one row. It is then enough to add a row corresponding to h at the bottom of this matrix and perform a tropical LUP form computation (see Algorithm 2) to read the Q_i's on the reduction of h. □

Proof of Theorem 4.2. We prove the main result by induction on the signature. We follow the order $\leq_{sign}$ for the induction. It is clear for $\sigma = e_1$ and also for the fact we can pass from an $\mathfrak{S}$-GB up to $<_{sign} e_i$ to $\leq_{sign} e_i$. We write the elements of G as $g_1, \dots, g_r$ for some $r \in \mathbb{Z}_{>0}$.

Let us assume that G is an $\mathfrak{S}$-GB up to signature $<_{sign} \sigma$ for some signature σ and let us prove it is a $\mathfrak{S}$-GB up to $\leq_{sign} \sigma$. We can assume that all $g \in G$ satisfy $LC(g) = 1$. Let $P \in I$ be irreducible, with $LC(P) = 1$ and such that $S(P) = \sigma$. We prove that there is $\tau \in T, g \in G$ such that $LM(P) = LM(\tau g)$ and $S(\tau g) = \tau S(g) = \sigma$.

Our first assumption for G implies that there exist at least one $g \in G$ and some $\tau \in T$ such that $\tau S(g) = S(P) = \sigma$.

If $LM(\tau g) =_{\leq} LM(P)$ we are done. Otherwise, by Lemma 3.21, there exist some $a, b \in k^*$ such that $S(aP + b\tau g) = \sigma'$ for some $\sigma' <_{sign} \sigma$.

We can apply Lemma 4.5 to $aP + b\tau g$ and obtain that there exist $h_1, \dots, h_r \in A$, such that $P = \sum_{i=1}^{r} h_i g_i$, and for all i, and x^γ monomial of h_i, the $x^\gamma S(g_i) = S(x^\gamma g_i) \leq_{sign} \sigma$ are all distincts. We remark that $LT(P) \leq \max_i(LT(g_i h_i))$. We denote by $m_i := LT(g_i h_i)$.

Among all such possible ways of writing P as $\sum_{i=1}^{r} h_i g_i$, we define β as the **minimum** of the $\max_i(LT(g_i h_i))$'s. Such a β exists thanks to Lemma 2.10 in [CM13] (adaptation to the non-homogeneous setting is for free). We write $x^u = LM(\beta)$.

If $LT(P) =_{\leq} \beta$, then we are done. Indeed, there is then some i and τ in the terms of h_i such that $LM(\tau g_i) = LM(P)$ and $S(\tau g_i) = \tau S(g_i) \leq_{sign} \sigma$.

We now show that $LT(P) < \beta$ leads to a contradiction.

We can renumber the g_i's so that:

- $\beta =_{\leq} m_1 =_{\leq} \cdots =_{\leq} m_d$.
- $\beta > m_i$ for $i > d$.

We can assume that among the set of possible $(h_1, \dots, h_r)$ that reaches β, we take one such that this d is minimal. Since $LT(P) < \beta$, we have $d \geq 2$.

We can write

$$Spol(g_1, g_2) = LC(g_2)\frac{lcm(LM(g_1), LM(g_2))}{LM(g_1)} g_1 - LC(g_1)\frac{lcm(LM(g_1), LM(g_2)}{LM(g_2)} g_2.$$

By construction, $LM(h_1)S(g_1) \neq LM(h_2)S(g_2)$, so $(LM(h_1)g_1, LM(h_2 g_2)$ is a normal pair. By Lemma 4.4, there exists a term μ such that $\mu \frac{lcm(LM(g_1), LM(g_2))}{LM(g_i)} = LM(h_i)$ for $i \in \{1, 2\}$. So by Lemma 3.23, (g_1, g_2) is a normal pair as well.

If $S(Spol(g_1, g_2)) = \sigma$, by the second property of the F5 criterion, we are done. Indeed, it provides us with some t_0 and g_0 such that $t_0 g_0$ is $\mathfrak{S}$-irreducible of signature σ, and as all $\mathfrak{S}$-irreducible polynomials of a given signature share the same leading monomial, thanks to Lemma 3.22, we have with t_0 and g_0 a satisfying τ and g.

Otherwise, $S(Spol(g_1, g_2)) <_{sign} \sigma$. Moreover, let

$$L = \frac{LC(h_1 g_1)}{LC(g_1)LC(g_2)} \frac{x^u}{lcm(LM(g_1), LM(g_2))}.$$

Then we have $S(L \cdot Spol(g_1, g_2)) \leq_{sign} \sigma$ thanks to Lemma 4.4. Using the same construction as before with the second assumption of the F5 criterion and Lemmata 3.21 and 4.5, we obtain some h_i''s such that $L \cdot Spol(g_1, g_2) = \sum_{i=1}^r h_i' g_i$, $LT(h_i' g_i) \leq LT(L \cdot Spol(g_1, g_2)) < \beta$ for all i. Furthermore, the signatures $S(x^\alpha g_i) = x^\alpha S(g_i)$ for $i \in \{1, \ldots, r\}$ and x^α in the support of h_i' are all distincts.

We then get:

$$\begin{aligned} P &= \sum_{i=1}^r h_i g_i, \\ &= \sum_{i=1}^r h_i g_i - L \cdot Spol(g_1, g_2) + \sum_{i=1}^r h_i' g_i, \\ &= \left(h_1 - \frac{LC(h_1 g_1)}{LC(g_1)} \frac{x^u}{LM(g_1)} + h_1' \right) g_1 \\ &\quad + \left(h_2 + \frac{LC(h_1 g_1)}{LC(g_2)} \frac{x^u}{LM(g_2)} + h_2' \right) g_2 + \sum_{i=3}^r (h_i + h_i') g_i, \\ &=: \sum_{i=1}^r \widetilde{h}_i g_i, \end{aligned}$$

where the $\widetilde{h}_i$'s are defined naturally.

By construction, $LT(\widetilde{h}_1 g_1) < LT(h_1 g_1) = \beta$ and $LT(\widetilde{h}_i g_i) \leq \beta$ for $i \leq d$ and $LT(\widetilde{h}_i g_i) < \beta$ for $i > d$.

As a consequence, we have obtained a new expression for f with either $max_i(LT(\widetilde{h}_i g_i)) < \beta$ or this term attained stricly less than d times, which is in either case a contradiction with their definitions as minima. So $LT(P) =_{\leq} \beta$, which concludes the proof of the main result. It is then direct to adapt the previous proof to the case of an $\mathfrak{S}_d$-GB. $\square$

This theorem holds also for $\mathfrak{S}$-GB (or $\mathfrak{S}_d$-GB) up to a given signature. We have the following variant for compatibility with sugar-degree:

Proposition 4.6. *Suppose that $d \in \mathbb{Z}_{>0}$, and G is a finite set of polynomials of I such that:*

1. *Any $g \in G$ is $\mathfrak{S}_d$-irreducible in $I^{\leq d}$.*
2. *For all $g_1, g_2 \in G$ we have $g_1, g_2 \in I^{\leq d}$ along with $Spol(g_1, g_2)$ and both sides of the S-pair for g_1 and g_2.*

3. For all $i \in [\![1, s]\!]$, there exists $g \in G$ such that $S_d(g) = e_i$.

4. for any $g_1, g_2 \in G$ such that (g_1, g_2) is a d-normal pair, there exists $g \in G$ and $t \in T$ such that tg is $\mathfrak{S}_d$-irreducible and $tS_d(g) = S_d(tg) = S_d(Spol(g_1, g_2))$.

Then G is an $\mathfrak{S}$-Gröbner basis of I.

Proof. Using Theorem 4.2, it is clear that such a G is an $\mathfrak{S}_d$-GB.

It remains to prove that G is then an $\mathfrak{S}$-GB.

A first remark if σ is a signature then for any $f \in I^{\leq d}$, if $S_d(f) = \sigma$ then $S(f) = \sigma$. It is a consequence of Remark 3.8:

A second remark is that if $f \in I^{\leq d}$ is such that $S_d(f) \leq_{sign_d} \sigma$, then $S(f) \leq_{sign} \sigma$.

One can then use the same proof on induction on σ as for Theorem 4.2 except for the decomposition of $LSpol(g_1, g_2)$ in signature $\leq_{sign} \sigma$ in terms of the g_i's. If $S(LSpol(g_1, g_2)) <_{sign} \sigma$, then it is true by induction. Otherwise, $S(LSpol(g_1, g_2)) = \sigma$, and as (g_1, g_2) is a d-normal pair, then $S_d(LSpol(g_1, g_2)) = \sigma$ by the first remark. One can then use the $\mathfrak{S}_d$-GB property, and the second remark. Given the decomposition, all arguments are the same. □

Remark 4.7. When going from $\mathfrak{S}_d$ to $\mathfrak{S}$, one should be aware that signatures can drop, even for $\mathfrak{S}_d$-irreducible polynomials. For instance, looking back to Example 3.9, $x - y$ is $\mathfrak{S}_d$-irreducible, with $S_d(x - y) = e_4$ when $d = 3$. Then, $x - y$ is also $\mathfrak{S}$-irreducible, but of signature ye_2. If instead, we had taken a small variation with $F = [x^2y + 1, x^3 + 1, y^3 + x, y^3 + y]$, then e_4 is not even a signature. It is a an $\mathfrak{S}_3$-signature but disappears in sugar-degree 4. Fortunately, Proposition 4.6 is enough to tell when d is big enough for an $\mathfrak{S}_d$-GB to be an $\mathfrak{S}$-GB (even though signatures might still drop going from $\mathfrak{S}_d$ to $\mathfrak{S}$).

5. F5 algorithm

In this section, we present our F5 algorithm. To this intent, we need to discuss some crucial algorithmic points: how to recognize with which pairs to proceed and how to build the Macaulay matrices and reduce them. Some algorithms are on the following page.

5.1. Admissible pairs and guessed signatures

The second condition in the Definition 4.1 of normal pairs is not possible to check in advance in an F5 algorithm. One needs an $\mathfrak{S}$-Gröbner basis up to the corresponding signature to be able to certify it. To circumvent this issue, we use the weaker notion of admissible pair.

Definition 5.1 (d-Admissible pair)**.** Given $g_1, g_2 \in I^{\leq d}$, let $Spol(g_1, g_2) = u_1 g_1 - u_2 g_2$ be the S-polynomial of g_1 and g_2. We have

$$u_i = \frac{lcm(LM(g_1), LM(g_2))}{LT(g_i)}.$$

We say that (g_1, g_2) is a d-**admissible pair** if:

1. $LM(u_i)S_d(g_i) = x_i^{\alpha} e_{j_i} \notin LM(TSyz)$.
2. $LM(u_1)S_d(g_1) \neq LM(u_2)S_d(g_2)$.

To certify that a set is an $\mathfrak{S}_d$-GB, handling d-admissible pairs instead of d-normal pairs is harmless. Indeed, d-normal pairs in $I^{\leq d}$ are contained inside the d-admissible pairs. Whether a pair is d-admissible can be checked easily before proceeding to reduction.

During the execution of the algorithm, when a polynomial $x^{\alpha} g$ is processed, it is at first not possible to know what is its signature. Algorithm 3 (on page 18) has computed $S_d(g)$ beforehand. Thanks to the F5 elimination criterion (Prop 3.10), we can detect some of the $x^{\alpha} g$ such that $S(x^{\alpha} g) \neq x^{\alpha} S(g)$ and eliminate them. For the processed polynomials, we use $x^{\alpha} S_d(g)$ as a **guessed signature** in the algorithm. Once an $\mathfrak{S}$-GB up to signature $< x^{\alpha} S_d(g)$ is computed, we have the following alternative. First case: $S_d(x^{\alpha} g) < x^{\alpha} S_d(g)$ and $x^{\alpha} g$ reduces to zero (by the computed $\mathfrak{S}_d$-GB up to d-signature $< x^{\alpha} S_d(g)$). The guessed signature was wrong but it is harmless as the polynomial is useless anyway. Second case: $S_d(x^{\alpha} g) = x^{\alpha} S_d(g)$, and then the guessed signature is certified. Once the criterion of Proposition 4.6 is satisfied, all signatures are certified.

What happens when we can obtain f with signature $S_d(f) = x^{\alpha} e_i$ in degree d, and $S_{d+1}(f) = x^{\beta} e_j <_{sign} x^{\alpha} e_i$ in degree $d+1$? Thanks to the way Algorithm 1 (on page 16) handles polynomials, always looking for smallest signature available, f and its multiples will then be built using only the second way. The first way of writing will at most appear so as to be reduced by the second one.

For instance, if we look back at Example 3.9, $(e_4, x-y)$ will only be used in sugar-degree 3. In sugar-degree 4 and more, it is $(ye_2, x-y)$ that will appear. $(e_4, x-y)$ may only appear in sugar-degree 4 to be reduced to 0 by $(ye_2, x-y)$.

5.2. Symbolic Preprocessing and Rewritten criterion

One of the main parts of the F5 algorithm 3 (on page 18) is the Symbolic Preprocessing: Algorithm 1 on page 16. From the current set of S-pairs, sugar-degree d, and the current $\mathfrak{S}_{d-1}$-GB, it produces a Macaulay matrix. One can read on the tropical reduction of this matrix new polynomials to append to the current basis to obtain an $\mathfrak{S}_d$-GB. It mostly consists of detecting which pairs are admissible and selecting a (complete) set of reductors.

A special part of the algorithm is the use of Rewritten techniques (due to Faugère (see [F02])).

The idea is the following. Once a polynomial has passed the F5 elimination criterion and is set to appear in a Macaulay matrix, it can be replaced by any other multiple of an element of G of the same d-signature. Indeed, assuming correctness of the algorithm without any rewriting technique, if one of them, h, is of d-signature $x^{\alpha} e_i$, the algorithm computes a tropical $\mathfrak{S}$-Gröbner basis up to d-signature $<_{sign} x^{\alpha} e_i$. Hence, h can be replaced by any other polynomial of same signature: it would be reduced to the same polynomial. By induction, one can prove that all of them can be replaced at the same time. We also remark that this is still valid for replacing a row of a given guessed d-signature by another of the same guessed d-signature.

One efficient way is to replace a polynomial $t \times g$ by the polynomial $x^{\beta}h$ ($h \in G$) of same (guessed) d-signature $tS_d(g)$ such that x^{β} has smallest degree.[3] Taking the sparsest available is another possibility. It actually leads to a substantial reduction of the running time of the F5 algorithm.

Algorithm 1: Symbolic-Preprocessing-Rewritten

---

input : P, a set of $d-1$-admissible pairs of sugar-degree d and G such that $G \cap I^{\leq d-1}$ is an $\mathfrak{S}_{d-1}$-GB
output: A Macaulay matrix
for *Q polynomial in P* **do**
  Replace Q in P by the polynomial $(uS(g), u \times g)$ with g latest added to G reaching the same guessed signature ;
$C \leftarrow$ the set of the **monomials** of the polynomials in P ;
$U \leftarrow$ the polynomials of P with their signature, except only one polynomial is taken by guessed signature ;
$D \leftarrow \emptyset$;
while $C \neq D$ **do**
  $m \leftarrow \max(C \setminus D)$;
  $D \leftarrow D \cup \{m\}$;
  $V \leftarrow \emptyset$;
  for $g \in G$ **do**
    if $LM(g) \mid m$ **then**
      $V \leftarrow V \cup \{(g, \frac{m}{LM(g)})\}$;
  $(g, \delta) \leftarrow$ the element of V with $\delta \times g$ of smallest guessed signature not already in the signatures of U, with tie-breaking by taking minimal δ (for degree then for $\leq_{sign}$) ;
  $U \leftarrow U \cup \{\delta \times g\}$;
  $C \leftarrow C \cup \{\text{monomials of } \delta \times g\}$;
$M \leftarrow$ the polynomials of U, written in a Macaulay matrix and ordered by increasing guessed signature ;
Return M ;

5.3. Linear algebra

To reduce the Macaulay matrices while respecting the signatures, we use the following tropical LUP algorithm from [V15]: Algorithm 2 on page 17. If the rows correspond to polynomials ordered by increasing signature, it computes a row-reduction, respecting the signatures with each non-zero row with a different leading monomial.

---

[3]Indeed, such an h can be considered as one of the most reduced possible.

Algorithm 2: The tropical LUP algorithm

input : M, a Macaulay matrix of degree d in A, with n_{row} rows and n_{col} columns, and mon a list of monomials indexing the columns of M.

output: $\widetilde{M}$, the U of the tropical LUP-form of M

$\widetilde{M} \leftarrow M$;

if $n_{col} = 1$ *or* $n_{row} = 0$ *or* M *has no non-zero entry* **then**

  Return $\widetilde{M}$;

else

  for $i = 1$ *to* n_{row} **do**

    Find j such that $\widetilde{M}_{i,j}$ has the greatest term $\widetilde{M}_{i,j}x^{mon_j}$ for $\leq$ of the row i;

    Swap the columns 1 and j of $\widetilde{M}$, and the 1 and j entries of mon;

    By **pivoting** with the first row, eliminates the coefficients of the other rows on the first column;

    Proceed recursively on the submatrix $\widetilde{M}_{i\geq 2, j\geq 2}$;

  Return $\widetilde{M}$;

5.4. A Complete Algorithm

We now provide with Algorithm 3 on page 18 a complete version of an F5 algorithm wich uses Buchberger-F5 criterion and all the techniques introduced in this section.

Theorem 5.2. *Algorithm 3 (on page 18) computes an $\mathfrak{S}$-GB of I. It avoids trivial tame syzygies.*

Proof. It relies on Theorem 4.2 and then Proposition 4.6. The proof is by induction on the sugar-degree, then i, then the $x^\alpha e_i$. One first proves that at the end of the main *while* loop any guessed signature is correct, or its row has reduced to zero, and then that $\mathfrak{S}_d$-GB are computed, signature by signature. One can then apply 4.6 on the output to conclude. Termination is a consequence of correctness and Prop. 3.19. For the syzygies, it is a consequence of Prop. 3.10 and the fact that trivial syzygies of leading monomial $x^\alpha e_i$ are such that $x^\alpha \in LM(\langle f_1, \ldots, f_{i-1} \rangle)$. □

Remark 5.3. Condition 1 of 4.2 and 3 of 4.6 is not satisfied when for some i, $f_i \in \langle f_1, \ldots, f_{i-1} \rangle$. This is harmless as: 1. As soon as it is found by computation, no signature in e_i will appear anymore. 2. The Buchberger-F5 criterion can be applied omitting f_i.

Algorithm 3: A complete F5 algorithm

input : $f_1, \dots, f_s$ polynomials, ordered by degree
output: A tropical $\mathfrak{S}$-GB G of $\langle f_1, \dots, f_s \rangle$
$G \leftarrow \{(e_i, f_i) \text{ for i in } [\![1, s]\!]\}$;
$B \leftarrow \{\text{S-pairs of } G\}$; $d \leftarrow 1$;
while $B \neq \emptyset$ **do**
  if *there is* i *s.t.* $|f_i| = d$ **then**
    Replace the occurence of f_i in G by its reduction modulo $G \cap \langle f_1, \dots, f_{i-1} \rangle$;
  P **receives** the pop of the $d-1$-admissible pairs in B of sugar-degree d. Suppress from B the others of sugar-degree d;
  Write them in a Macaulay matrix M_d, along with their $\mathfrak{S}_d$-reductors obtained from G (one per signature) by **Symbolic-Preprocessing-Rewritten**(P, G) (Algorithm 1);
  Apply Algorithm 2 to compute the U in the tropical LUP form of M (no choice of pivot) ;
  Add to G all the polynomials obtained from $\widetilde{M}$ that provide new leading monomial up to their d-signature ;
  Add to B the corresponding new d-admissible pairs ;
  $d \leftarrow d + 1$;
Return G ;

6. Iterative F5

In this section, we present briefly another way of extending the F5 algorithm to the affine setting: a completely iterative way in the initial polynomials. The idea is to compute tropical Gröbner bases for $\langle f_1 \rangle, \langle f_1, f_2 \rangle, \dots, \langle f_1, \dots, f_s \rangle$.

This corresponds to using the position over term ordering on the signatures, or in terms of filtration, to the following filtration on A^s:

Definition 6.1. We write that $x^\alpha e_i \leq_{incr} x^\beta e_j$ if:

1. if $i < j$.
2. if $i = j$ and $|x^\alpha f_i| < |x^\beta f_j|$.
3. if $i = j$ and $|x^\alpha f_i| = |x^\beta f_j|$, and
   - $x^\alpha \notin LM(I_{i-1})$ and $x^\beta \in LM(I_{i-1})$, or
   - both $x^\alpha, x^\beta \in LM(I_{i-1})$ and $x^\alpha \leq x^\beta$, or
   - both $x^\alpha, x^\beta \notin LM(I_{i-1})$ and $x^\alpha \leq x^\beta$.

Proposition 6.2 ([F02])**.** *If $x^\alpha \in LM(I_{i-1})$, then the filtration is constant at*

$$I_{\leq x^\alpha e_i}.$$

Proof. We can write $x^\alpha + g = \sum_{j<i} a_j f_j$, with for all j $a_j \in I$, and $g \in I$ with no monomial in $LM(I_{i-1})$. Then: $x^\alpha f_i = (-g) f_i + \sum_{j=1}^{i-1} (a_j f_i) f_j$, and the filtration is constant at $I_{\leq x^\alpha e_i}$. $\square$

It is then possible to state a Buchberger-F5 criterion and provide an adapted F5 algorithm. The two algorithms will then differ in the following way.

1. For a given x^α and e_i, the vector space $I_{<x^\alpha e_i}$ is much bigger in the iterative setting, often of infinite dimension. Thus, polynomials of signature $x^\alpha e_i$ can be more deeply reduced.
2. More syzygies can be avoided in the iterative setting.
3. However, many more matrices are to be produced: one for each i and each necessary degree. Construction of the matrices is not mutualised by degree anymore.

7. The F4 approach

Another way to compute tropical Gröbner bases for affine polynomials is to adapt Faugère's F4 algorithm [F99].

Roughly, the F4 algorithm is an adaptation of Buchberger's algorithm such that all S-polynomials of a given degree are processed and reduced together in a big Macaulay matrix, along with their reducers. The algorithm carries on the computation until there is no S-polynomials to reduce.

In a tropical setting, we have adapted the so-called "normal strategy" of F4 using the tropical LUP algorithm to reduce the Macaulay matrices.

For simplicity of exposition, we have only used Algorithm 2 to reduce the Macaulay matrices.

Nevertheless, so-called tropical row-echelon forms (Algorithm 3.2.2 and 3.7.3 of [V15]) are also possible. This enables a trade-off between speed, thoroughness of the reduction and loss in precision. Indeed, contrary to the F5 strategy where one has to reduce with respect to signatures, any reasonable strategy of reduction is possible.

In this section, we present the algorithm, along with proof of correctness and termination.

7.1. Buchberger criterion and reduction

Defining algorithm for reduction in a tropical setting is not straightforward. A classical example comes from trying to use a naive adaptation of the classical polynomial reduction algorithm to reduce x by the family $G = (x + y, y + 2x)$, in $\mathbb{Q}_2[x, y]$ with $w = (0, 0)$. Even though G is a minimal Gröbner basis, the corresponding algorithm would not terminate, going through all the $2^n x$'s and $-2^n y$'s as intermediate polynomials to reduce.

Yet, we need an algorithm for reduction so as to state a Buchberger criterion, on which the correctness of F4 is based. As in [V15, VY17], this issue can be circumvented using matrix reduction.

We begin by adapting the Symbolic-Preprocessing algorithm, removing any call for signatures. The corresponding version is presented in Algorithm 4 on page 20.

Algorithm 4: F4-Symbolic-Preprocessing

---

input : P, a list of pairs of elements of A, G a list of elements in A.
output: A Macaulay matrix
$C \leftarrow$ the set of the **monomials** of the polynomials in P ;
$U \leftarrow$ the polynomials of P ;
$D \leftarrow \emptyset$;
while $C \neq D$ **do**
  $m \leftarrow \max(C \setminus D)$;
  $D \leftarrow D \cup \{m\}$;
  $V \leftarrow \emptyset$;
  for $g \in G$ **do**
    if $LM(g) \mid m$ **then**
      $V \leftarrow V \cup \{(g, \frac{m}{LM(g)})\}$;
  $(g, \delta) \leftarrow$ the element (g, δ) of V with $\delta \times g$ of biggest leading term, with tie-breaking by taking minimal δ (for $\leq$) ;
  $U \leftarrow U \cup \{\delta \times g\}$;
  $C \leftarrow C \cup \{\text{monomials of } \delta \times g\}$;
$M \leftarrow$ the polynomials of U, written in a Macaulay matrix ;
Return M ;

Algorithm 4 allows us to construct matrices on which we can perform reduction. Algorithm 5 on page 20 is then build upon it. It is used to define the reduction of a polynomial by a family of polynomials. It will then be used only for theoretical purpose as a way to provide a Buchberger criterion for termination of F4.

Algorithm 5: Tropical polynomial reduction

---

input : $f \in A$, G a list of elements in A.
output: $\widetilde{f}$ a reduction of f by G.
$M \leftarrow$ **F4-Symbolic-Preprocessing**$(\{f\}, G)$;
Swap rows of M so that f corresponds to the last row of M ;
$\widetilde{M} \leftarrow$ **Tropical-LUP**(M);
$\widetilde{f} \leftarrow$ last row of $\widetilde{M}$;
Return $\widetilde{f}$;

Proposition 7.1 (Buchberger's criterion)**.** *G is a Tropical Gröbner basis of $\langle G \rangle$ if and only if every S-polynomial of G reduces to zero using Algorithm 5.*

Proof. We already have the Buchberger-F5 criterion, Theorem 4.2. It is clear that Buchberger's criterion can be obtained as a weak variant of the Buchberger-F5 criterion, getting rid of the signatures. □

7.2. An algorithm and its proof

We can now present a Tropical F4 Algorithm with Algorithm 6 on page 21.

Algorithm 6: A tropical F4 algorithm

---

input : $f_1, \ldots, f_s$ in A
output: A tropical GB G of $\langle f_1, \ldots, f_s \rangle$
$G \leftarrow (f_1, \ldots, f_s)$;
$B \leftarrow \{\text{pairs of } G\}$;
while $B \neq \emptyset$ **do**
  $d \leftarrow$ the smallest degree of an lcm of leading terms in a pair of B ;
  P **receives** the pop of the pairs of degree d in B;
  Write them in a Macaulay matrix M, along with their reductors obtained by **F4-Symbolic-Preprocessing**(P, G) (Algorithm 4);
  Apply Algorithm 2 to compute $\widetilde{M}$ the row-echelon form of M (with choice of pivot) ;
  Add to G all the polynomials obtained from $\widetilde{M}$ that provide leading monomials not in $\langle \{LM(g) \text{ for } g \in G\} \rangle$;
  Add to B the corresponding new pairs ;
Return G ;

Proposition 7.2. *The Tropical F4 algorithm, Algorithm 6 terminates and computes a tropical GB of I.*

Proof. If Algorithm 6 does not terminate for some initial $(f_1, \ldots, f_s)$, it means that there is an infinite sequence of polynomials $(h_i)_{i \in \mathbb{N}}$ that are added to G. Indeed, otherwise, the algorithm would terminate as it would run out of S-pairs.

For those h_i to be added in the algorithm, they have to provide an $LM(h_i)$ not already in the monomial ideal $\langle LM(g),\, g \in G \rangle$ for the current G. This would mean an infinite number of generator in a monomial ideal, which is a contradiction. Hence the algorithm terminates.

As for correctness, it is an application of the Buchberger criterion. After every turn in the **while** loop, all already processed S-polynomial of an S-pair reduce to 0 by G. Hence, after termination, it is the case for all the S-pairs of G and we can apply Buchberger's criterion. □

8. Numerical experiments

A toy implementation of our algorithms in Sagemath [Sage] is available on `https://gist.github.com/TristanVaccon`. We have gathered some numerical results in the following arrays. Timings are in seconds of CPU time.[4]

8.1. Benchmarks

Here, the base field is $\mathbb{Q}$ with 2-adic valuation. We have applied the tropical F5 algorithm, Algorithm 3, an iterative tropical F5, and a tropical F4 algorithm on the Katsura n and Cyclic n systems for varying n. Dots mean no conclusion in decent time.

w=[0,...,0]	Katsura 4	5	6	7	Cyclic 4	5	6
Trop F5	.16	1.2	1371	•	0.4	21	•
Iterative trop F5	0.3	1.9	1172	•	0.4	21	•
Trop F4	.5	5	30	•	1.7	112	•

$w=[(-2)^{i-1}]$	Katsura 4	5	6	7	Cyclic 4	5	6
Trop F5	0.15	0.8	17	•	0.18	11	•
Iterative trop F5	0.18	1.1	20	•	0.18	11	•
Trop F4	0.2	1.7	15	•	1	65	•

8.2. Tropical F5 + FGLM

For a given p, we take three polynomials with random coefficients in $\mathbb{Z}_p$ (using the Haar measure) in $\mathbb{Q}_p[x,y,z]$ of degree $2 \leq d_1 \leq d_2 \leq d_3 \leq 4$. We first compute a tropical Gröbner basis for the weight $w = [0,0,0]$[5] and the grevlex monomial ordering, and then apply an FGLM algorithm (tropical to classical as in Chapter 9 of [V*]) to obtain a lex GB. For any given choice of d_i's, we repeat the experiment 50 times. Coefficients of the initial polynomials are all given at some high-enough precision $O(p^N)$ for no precision issue to appear. We can not provide a certificate on the monomials of the output basis though. Results are compiled in the following arrays.

Firstly, an array for timings given as couples: average of the timings for the tropical F5 part and for the FGLM part, with $D = d_1+d_2+d_3-2$, the Macaulay bound. We add that for $p = 2, 3$, there is often a huge standard deviation on the timings of the F5 part.

	$D=4$		5		6		7		8		9	
$p=2$	.7	0.2	2.5	0.5	18	2.3	300	11	50	37	145	138
3	.8	.2	.9	.5	4	2	9	11	16	37	80	144
101	0.3	.2	.5	.5	1	2	3	10	4.6	37	11	150
65519	.4	.2	.6	.6	1.3	2.6	3.5	11	5	39	10	132

[4]Everything was performed on a Ubuntu 16.04 with 2 processors of 2.6GHz and 16 GB of RAM.

[5]Efficiency of this choice regarding to the loss in precision was studied in the extended version of [V15]

Coefficients of the output tropical GB or classical GB are known at individual precision $O(p^{N-m})$. We compute the total mean and max on those m's on the obtained GB. Results are compiled in the following array as couples of mean and max. The first array is for the F5 part and the second for the precision on the final result.

	$D = 4$		5		6		7		8		9	
$p = 2$	1.3	13	1.3	13	1.3	14	1.5	13	1.4	17	1.3	15
3	.6	6	.7	8	.7	7	.6	7	.6	7	.6	10
101	0	1	0	1	0	1	0	2	0	2	0	1
65519	0	0	0	0	0	1	0	0	0	0	0	0

	$D = 4$		5		6		7		8		9	
$p = 2$	8	71	17	170	58	393	167	913	290	1600	570	3900
3	5	38	13	114	27	230	81	640	167	1600	430	3100
101	.2	11	0	2	1.3	80	4	210	8	407	0	2
65519	0	0	0	0	0	0	0	0	0	0	0	0

Most of the loss in precision appears in the FGLM part. In comparison, the F5 part is quite stable, and hence, our goal is achieved.

References

[AP] Alberto Arri and John Perry. The F5 criterion revised. Journal of Symbolic Computation, 2011, plus *corrigendum* in 2017.

[BFS04] Magali Bardet, Jean-Charles Faugère and Bruno Salvy. On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations Proceeding of the 2004 International Conference on Polynomial System Solving (ICPSS'04), Paris, France, pages 71–75, 2004.

[BFS14] Magali Bardet, Jean-Charles Faugère, and Bruno Salvy. On the Complexity of the F5 Gröbner basis Algorithm. Journal of Symbolic Computation, pages 1–24, September 2014. 24 pages.

[BCM11] Bigatti A.M. and Caboara M. and Robbiano L. Computing inhomogeneous Gröbner bases, Journal of Symbolic Computation, 2011.

[C13] Chan, Andrew J., Gröbner bases over fields with valuations and tropical curves by coordinate projections, PhD Thesis, University of Warwick, August 2013.

[CM13] Chan, Andrew J. and Maclagan, Diane Gröbner bases over fields with valuations, to appear in Mathematics of Computation.

[C15] Connes, Alain, An essay on the Riemann Hypothesis, `http://arxiv.org/pdf/1509.05576`, 2015.

[CLO] Cox, D., Little, J., & O'shea, D. , Ideals, varieties, and algorithms (4th edition), Springer, 2015.

[E13] Eder Christian An analysis of inhomogeneous signature-based Gröbner basis computations, Journal of Symbolic Computation, 2013

[EF17] Eder Christian and Faugère Jean-Charles A survey on signature-based algorithms for computing Gröbner bases, Journal of Symbolic Computation, 2017

[EKL06] Einsiedler, Manfred and Kapranov, Mikhail and Lind, Douglas Non-archimedean amoebas and tropical varieties, Journal für die reine und angewandte Mathematik (Crelles Journal), 2006.

[F99] Faugère, Jean-Charles, A new efficient algorithm for computing Gröbner bases (F4), Journal of Pure and Applied Algebra, 1999

[F02] Faugère, Jean-Charles, A new efficient algorithm for computing Gröbner bases without reduction to zero (F5), Proceedings of the 2002 international symposium on Symbolic and algebraic computation, ISSAC '02, Lille, France.

[GMNRT91] Giovini Alessandro, Mora Teo, Niesi Gianfranco, Robbiano Lorenzo, and Traverso Carlo. "One sugar cube, please" or selection strategies in the Buchberger algorithm. In Proceedings of the 1991 international symposium on Symbolic and algebraic computation (ISSAC '91).

[Gfan] Jensen, Anders N. Gfan, a software system for Gröbner fans and tropical varieties, Available at `http://home.imf.au.dk/jensen/software/gfan/gfan.html`.

[MS15] Maclagan, Diane and Sturmfels, Bernd, Introduction to tropical geometry, Graduate Studies in Mathematics, volume 161, American Mathematical Society, Providence, RI, 2015, ISBN 978-0-8218-5198-2.

[MR16] Markwig, Thomas and Ren, Yue Computing tropical varieties over fields with valuation, `http://arxiv.org/pdf/1612.01762`, 2016.

[Sage] SageMath, the Sage Mathematics Software System (Version 8.3.beta1), The Sage Development Team, 2018, http://www.sagemath.org.

[V14] Vaccon Tristan, Matrix-F5 algorithms over finite-precision complete discrete valuation fields, Proceedings of 39th International Symposium on Symbolic and Algebraic Computation, ISSAC'14, Kobe, Japan.

[V*] Vaccon Tristan, Précision p-adique, thèse de l'Université de Rennes 1, https://tel.archives-ouvertes.fr/tel-01205269.

[V15] Vaccon Tristan, Matrix-F5 Algorithms and Tropical Gröbner Bases Computation, Proceedings of the 40th International Symposium on Symbolic and Algebraic Computation, ISSAC 2015, Bath, United Kingdom. Extended version in the Journal of Symbolic Computation, Dec. 2017.

[VY17] Vaccon Tristan and Yokoyama Kazuhiro, A Tropical F5 algorithm, Proceedings of the 42th International Symposium on Symbolic and Algebraic Computation, ISSAC 2017, Kaiserslautern, Germany.

[VVY18] Vaccon Tristan, Verron Thibaut and Yokoyama Kazuhiro, On Affine Tropical F5 Algorithms, Proceedings of the 42th International Symposium on Symbolic and Algebraic Computation, ISSAC 2018, New York City, USA.