

Signature Rewriting in Gröbner Basis Computation

Christian Eder and Bjarke Hammersholt Rouné

ABSTRACT

We introduce the **RB** algorithm for Gröbner basis computation, a simpler yet equivalent algorithm to **F5GEN**. **RB** contains the original unmodified **F5** algorithm as a special case, so it is possible to study and understand **F5** by considering the simpler **RB**. We present simple yet complete proofs of this fact and of **F5**'s termination and correctness.

RB is parametrized by a rewrite order and it contains many published algorithms as special cases, including **SB**. We prove that **SB** is the best possible instantiation of **RB** in the following sense. Let **X** be any instantiation of **RB** (such as **F5**). Then the S-pairs reduced by **SB** are always a subset of the S-pairs reduced by **X** and the basis computed by **SB** is always a subset of the basis computed by **X**.

1. INTRODUCTION

The **F5** algorithm [5] is a significant development in Gröbner basis computation and the difficulty of understanding it is a widely cited concern. We address this problem in this paper by introducing **RB**, an algorithm simpler than yet equivalent to **F5GEN** [9]. **RB** contains the original unmodified **F5** as a special case [9], so we can understand and prove statements about **F5** by considering the simpler **RB**. In this way we present simplified proofs of the termination and correctness of **F5**. Much of the paper is concerned with showing exactly how to derive **F5** as a special case of **RB**.

There are many publications that present signature-based algorithms for Gröbner basis computation. Seeing as these algorithms are based on signatures, they are all somehow related to **F5**, yet it is often not clear what the exact relationship is — is **F5** a special case of the new algorithm? If not, what are the exact differences? Pan, Hu and Wang [9] (PHW) addressed this problem by introducing the **F5GEN** algorithm. **F5GEN** contains many published algorithms as special cases which makes it possible to compare algorithms theoretically. PHW pose the open problem of which instantiation of **F5GEN** is faster:

Moreover, with this proved F5GEN algorithm,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

researchers can shift their focus on the different variants of the F5GEN algorithm and find out the fastest one. — Pan, Hu, Wang [9]

We answer this open problem by proving that the **SB** algorithm [10] is the best possible instantiation of **RB** (and hence **F5GEN**) in the following sense. Let **X** be any instantiation of **RB** (such as **F5**). Then the S-pairs reduced by **SB** are always a subset of the S-pairs reduced by **X** up to signature and the basis computed by **SB** is always a subset of the basis computed by **X** up to sig-lead pairs.

Rouné and Stillman's algorithm **SB** [10] is an improvement of Gao, Volny and Wang's algorithm **GVW** [8]. **SB** is equivalent to Arri and Perry's algorithm **AP** [2] and also to Volny's modification **GVWHS** [12] of **GVW**, so all statements made about **SB** also apply to **AP** and **GVWHS**.

Section 2 introduces terminology and **SB**. Section 3 introduces **RB** and proves that **SB** is the best possible instantiation of **RB**. Section 4 introduces **RB5**, a special case of **RB** that is more similar to **F5**. Section 5 proves that **F5** is a special case of **RB5** (and hence **RB**). Section 6 proves termination of **RB** (and hence **F5**).

REMARK 1. *The **F5** paper [6] requires that the input polynomials form a homogeneous regular sequence. Faugère must have known of an algorithm without these restrictions since the **F5** paper contains benchmarks on ideals that do not meet the restrictions. In this paper, the term **F5** refers exclusively to **F5** exactly as described in the **F5** paper [6] including the restriction to homogeneous regular sequences.*

2. THE SIGNATURE BASIS ALGORITHM

We briefly introduce **SB** in this section. We refer to [10] for a full treatment of **SB** with proofs. If we consider a polynomial ring R then the main differences between **SB** and the classic Buchberger algorithm is that everything is lifted from R to R^m , that polynomial reduction is constrained in that some reduction steps are not allowed and that the S-pair elimination criteria are more powerful. It is not actually necessary to represent polynomials in R^m when implementing **SB** on a computer (see Remark 2).

2.1 Notation and Terminology

Let R be a polynomial ring over a field κ . All polynomials $f \in R$ can be uniquely written as a finite sum $f = \sum_{cx^v \in M} cx^v$ where $c \in \kappa$, $x^v := \prod_i x_i^{v_i}$ and M is minimal. The elements of M are the *terms* of f . A *monomial* is a polynomial with exactly one term. A monomial with a coefficient of 1 is *monic*. Neither monomials nor terms

of polynomials are necessarily monic. We write $f \simeq g$ for $f, g \in R$ if there exists an $s \in \kappa$ such that $f = sg$.

Let R^m be a free R -module and let e_1, \dots, e_m be the standard basis of unit vectors in R^m . All module elements $\alpha \in R^m$ can be uniquely written as a finite sum $\alpha = \sum_{ae_i \in M} ae_i$ where the a are monomials and M is minimal. The elements of M are the *terms* of α . A *module monomial* is an element of R^m with exactly one term. A module monomial with a coefficient of 1 is *monic*. Neither module monomials nor terms of module elements are necessarily monic. Let $\alpha \simeq \beta$ for $\alpha, \beta \in R^m$ if there exists an $s \in \kappa$ such that $\alpha = s\beta$.

Let \leq denote two different orders – one for R^m and one for R . The order for R is a monomial order, which means that it is a well-order on the set of monomials in R such that $a \leq b$ implies $ca \leq cb$ for all monomials $a, b, c \in R$. The order for R^m is a module monomial order which means that it is a well-order on the set of module monomials in R^m such that $L \leq T$ implies $cL \leq cT$ for all module monomials $L, T \in R^m$ and monomials $c \in R$. We require the two orders to be compatible in the sense that $a \leq b$ if and only if $ae_i \leq be_i$ for all monomials $a, b \in R$ and $i = 1, \dots, m$.

Consider a finite sequence of polynomials $g_1, \dots, g_m \in R$ that we call the *input polynomials*. Define the homomorphism $\alpha \mapsto \bar{\alpha}$ from R^m to R by $\bar{\alpha} := \sum_{i=1}^m \alpha_i g_i$. $\alpha \in R^m$ is a *syzygy* if $\bar{\alpha} = 0$. Let $\mathcal{G} \subseteq R^m$ be finite and assume that $\mathfrak{s}(\alpha) \simeq \mathfrak{s}(\beta) \Rightarrow \alpha = \beta$ for $\alpha, \beta \in \mathcal{G}$. $\alpha, \beta \in R^m$ are *equal up to sig-poly pairs* if $(\mathfrak{s}(\alpha), \bar{\beta}) = (\mathfrak{s}(\beta), \bar{\alpha})$.

Let the *lead term* $\text{lt}(f)$ be the \leq -maximal term of $f \in R \setminus \{0\}$. Let the *signature* $\mathfrak{s}(\alpha)$ be the \leq -maximal term of $\alpha \in R^m \setminus \{0\}$. In this way every module element $\alpha \in R^m \setminus \{0\}$ has two main associated characteristics – the signature $\mathfrak{s}(\alpha) \in R^m$ and the lead term $\text{lt}(\bar{\alpha}) \in R$ of its image $\bar{\alpha}$ in R . Lead terms and signatures include a coefficient for convenience, though an implementation of **SB** need not store the signature coefficients. If $ae_i = \mathfrak{s}(\alpha)$ then $\text{ind}(\alpha) := i$ is the *index of signature* of α .

Consider the unique extension of the module monomial order $<$ to module monomials that may have negative exponents such that $ae_i < be_j \Leftrightarrow cae_i < cbe_j$ for all monomials $a, b, c \in R$. The *sig-lead ratio* of $\alpha \in R^m$ is $r_\alpha := \frac{\mathfrak{s}(\alpha)}{\text{lt}(\bar{\alpha})}$.

This notation is the **SB** notation [11], which is improved from the notation of the preceding paper [10]. In particular, there is no module R^n and no function $\phi: R^n \rightarrow R^m$.

REMARK 2. *The **SB** notation using module elements differs from the standard notation which considers sig-poly pairs $(\mathfrak{s}(\alpha), \bar{\alpha})$ in place of $\alpha \in R^m$. **SB** can also use sig-poly pairs (see Section 5.1). We believe that the **SB** notation makes signature algorithms easier to understand and reason about.*

2.2 Reduction With Signatures

Both **SB** and the classic Buchberger algorithm are based on reducing S-pairs. We first describe classic polynomial reduction and then describe **SB**'s reduction in similar terms.

Classic polynomial reduction

Let $f \in R$ and let t be a term of f . Then we can *reduce* t by $g \in R$ if $\text{lt}(g) | t$ or equivalently if

- there exists a monomial b such that $\text{lt}(bg) = t$.

The outcome of the reduction step is then $f - bg$ and g is the *reducer*. It holds that $\text{lt}(bg) \leq \text{lt}(f)$, but that is not listed as a requirement since it is implied. If $\text{lt}(bg) \simeq \text{lt}(f)$ then

the reduction step is a *top reduction step* and otherwise it is a *tail reduction step*.

The result of classic polynomial reduction of $f \in R$ is a polynomial $h \in R$ that has been calculated from f through a sequence of reduction steps such that h cannot be further reduced. The reduction is a *tail reduction* if only tail reduction steps are allowed and it is a *top reduction* if only top reduction steps are allowed.

The implied condition that $\text{lt}(bg) \leq \text{lt}(f)$ is equivalent to $\text{lt}(f - bg) \leq \text{lt}(f)$ so during classic polynomial reduction it is not allowed to increase the leading term. For tail reduction we perform only those reduction steps that do not change the leading term at all.

Signature reduction

Let $\alpha \in R^m$ and let t be a term of $\bar{\alpha}$. Then we can *s-reduce* t by $\beta \in R^m$ if

- there exists a monomial b such that $\text{lt}(\overline{b\beta}) = t$ and
- $\mathfrak{s}(b\beta) \leq \mathfrak{s}(\alpha)$.

The outcome of the \mathfrak{s} -reduction step is then $\alpha - b\beta$ and β is the *s-reducer*. The second condition is analogous to the implied condition $\text{lt}(bg) \leq \text{lt}(f)$ from classic polynomial reduction, the condition is just lifted to R^m so that it involves signatures. When β \mathfrak{s} -reduces t we also say for convenience that $b\beta$ \mathfrak{s} -reduces α . That way b is introduced implicitly instead of having to repeat the equation $\text{lt}(\overline{b\beta}) = t$.

Just as for classic polynomial reduction, if $\text{lt}(\overline{b\beta}) \simeq \text{lt}(\bar{\alpha})$ then the \mathfrak{s} -reduction step is a *top s-reduction step* and otherwise it is a *tail s-reduction step*. We need words for the analogous distinction for signatures, so if $\mathfrak{s}(b\beta) \simeq \mathfrak{s}(\alpha)$ then the reduction step is a *singular s-reduction step* and otherwise it is a *regular s-reduction step*.

The result of \mathfrak{s} -reduction of $\alpha \in R^m$ is a $\gamma \in R^m$ that has been calculated from α through a sequence of \mathfrak{s} -reduction steps such that γ cannot be further \mathfrak{s} -reduced. The reduction is a *tail s-reduction* if only tail \mathfrak{s} -reduction steps are allowed and it is a *top s-reduction* if only top \mathfrak{s} -reduction steps are allowed. The reduction is a *regular s-reduction* if only regular \mathfrak{s} -reduction steps are allowed. A module element $\alpha \in R^m$ is *s-reducible* if it can be \mathfrak{s} -reduced. If α \mathfrak{s} -reduces to γ and γ is a syzygy then we say that α *s-reduces to zero* even if $\gamma \neq 0$.

As for classic polynomial reduction, the implied condition $\text{lt}(\overline{b\beta}) \leq \text{lt}(\bar{\alpha})$ is equivalent to $\text{lt}(\overline{\alpha - b\beta}) \leq \text{lt}(\bar{\alpha})$, so during \mathfrak{s} -reduction it is not allowed to increase the leading term. For tail \mathfrak{s} -reduction we perform only those \mathfrak{s} -reduction steps that do not change the leading term at all. Analogously, the condition $\mathfrak{s}(b\beta) \leq \mathfrak{s}(\alpha)$ is equivalent to $\mathfrak{s}(\alpha - b\beta) \leq \mathfrak{s}(\alpha)$, so during \mathfrak{s} -reduction it is not allowed to increase the signature. For regular \mathfrak{s} -reduction, we perform only those \mathfrak{s} -reduction steps that do not change the signature at all.

Classic reduction is always with respect to a finite *basis* $B \subseteq R$. The reducers in classic polynomial reduction are chosen from the basis B . Analogously, \mathfrak{s} -reduction is always with respect to a finite *basis* $\mathcal{G} \subseteq R^m$. The \mathfrak{s} -reducers in \mathfrak{s} -reduction are chosen from the basis \mathcal{G} .

\mathcal{G} is a *signature Gröbner basis in signature T* if all $\alpha \in R^m$ with $\mathfrak{s}(\alpha) = T$ \mathfrak{s} -reduce to zero. \mathcal{G} is a *signature Gröbner basis up to signature T* if \mathcal{G} is a signature Gröbner basis in all signatures L such that $L < T$. \mathcal{G} is a *signature Gröbner basis* if it is a signature Gröbner basis in all signatures.

SB computes a signature Gröbner basis. If \mathcal{G} is a signature Gröbner basis then $\{\bar{\alpha} \mid \alpha \in \mathcal{G}\}$ is a Gröbner basis of $\langle g_1, \dots, g_m \rangle$, so we can use **SB** to compute Gröbner bases.

2.3 S-pairs and Syzygies

Let $f, g \in R$ and let $c := \gcd(\text{lt}(f), \text{lt}(g))$ be the monic greatest common divisor of $\text{lt}(f)$ and $\text{lt}(g)$. In the classic Buchberger algorithm the *S-polynomial* between f and g is

$$\frac{\text{lt}(g)}{c}f - \frac{\text{lt}(f)}{c}g.$$

The classic Buchberger algorithm proceeds by reducing S-polynomials. If an S-polynomial reduces to $h \neq 0$ then h is added to the basis so that the S-polynomial now reduces to zero by this larger basis. The classic Buchberger algorithm terminates once all S-polynomials between elements of the basis reduce to zero. The **SB** algorithm works the same way except that these computations are lifted from R to R^m in the following way.

Let $\alpha, \beta \in R$ and let $c := \gcd(\text{lt}(\bar{\alpha}), \text{lt}(\bar{\beta}))$ be the monic greatest common divisor of $\text{lt}(\bar{\alpha})$ and $\text{lt}(\bar{\beta})$. The *S-pair* between α and β is

$$\mathcal{S}(\alpha, \beta) := \frac{\text{lt}(\bar{\beta})}{c}\alpha - \frac{\text{lt}(\bar{\alpha})}{c}\beta.$$

If $\mathfrak{s}\left(\frac{\text{lt}(\bar{\beta})}{c}\alpha\right) \simeq \mathfrak{s}\left(\frac{\text{lt}(\bar{\alpha})}{c}\beta\right)$ then the S-pair is *singular* and otherwise it is *regular*. By ‘‘S-pair’’ we always mean ‘‘regular S-pair’’. Observe that $\mathcal{S}(\alpha, \beta) \in R^m$ and that $\mathcal{S}(\alpha, \beta)$ is the S-polynomial between $\bar{\alpha}$ and $\bar{\beta}$.

SB proceeds by \mathfrak{s} -reducing S-pairs. If an S-pair \mathfrak{s} -reduces to γ and $\bar{\gamma}$ is not zero then γ is added to the basis. Theorem 3 implies that if all S-pairs and all e_i \mathfrak{s} -reduce to zero then the basis is a signature Gröbner basis.

THEOREM 3. *Let T be a module monomial of R^m and let $\mathcal{G} \subseteq R^m$ be a finite basis. Assume that all S-pairs $p := \mathcal{S}(\alpha, \beta)$ with $\alpha, \beta \in \mathcal{G}$ and $\mathfrak{s}(p) < T$ \mathfrak{s} -reduce to zero and all e_i with $e_i < T$ \mathfrak{s} -reduce to zero. Then \mathcal{G} is a signature Gröbner basis up to signature T .*

The outcome of classic polynomial reduction depends on the choice of reducer, so the choice of reducer can change what the intermediate bases are in the classic Buchberger algorithm. Lemma 4 implies that all S-pairs with the same signature yield the same regular \mathfrak{s} -reduced result as long as we process S-pairs in order of increasing signature.

LEMMA 4. *Let $\alpha, \beta \in R^m$ and let \mathcal{G} be a signature Gröbner basis up to signature $\mathfrak{s}(\alpha) = \mathfrak{s}(\beta)$. If α and β are both regular top \mathfrak{s} -reduced then $\text{lt}(\bar{\alpha}) = \text{lt}(\bar{\beta})$ or $\bar{\alpha} = \bar{\beta} = 0$. If α and β are both regular \mathfrak{s} -reduced then $\bar{\alpha} = \bar{\beta}$.*

A signature Gröbner basis is *minimal* if no basis element top \mathfrak{s} -reduces any other basis element. Theorem 5 implies that the minimal signature Gröbner basis is unique and is contained in all signature Gröbner bases up to sig-lead pairs. **SB** computes a minimal signature Gröbner basis.

THEOREM 5. *Let A be a minimal signature Gröbner basis and let B be a signature Gröbner basis of g_1, \dots, g_m . Then it holds for all $\alpha \in A$ that there exists a scalar $c \in \kappa$ and a $\beta \in B$ such that $\mathfrak{s}(\alpha) = c\mathfrak{s}(\beta)$ and $\text{lt}(\bar{\alpha}) = c\text{lt}(\bar{\beta})$.*

```

SimpleSignatureBasisAlgorithm( $\{g_1, \dots, g_m\} \subseteq R$ )
 $\mathcal{G} \leftarrow \emptyset$  ( $\mathcal{G}$  will be the signature Gröbner basis)
 $P \leftarrow \{e_1, \dots, e_m\}$  ( $P$  is the set of pending reductions)
 $H \leftarrow \langle 0 \rangle \subseteq R^m$  ( $H$  will be the initial module of syzygies)
while  $P \neq \emptyset$  do
   $p \leftarrow$  an element of  $P$  with  $\leq$ -minimal signature
   $P \leftarrow P \setminus \{p\}$ 
   $p' \leftarrow$  result of regular reducing  $p$ 
  if  $p' = 0$  then
     $H \leftarrow H + \langle \mathfrak{s}(p') \rangle$ 
  else if  $p'$  is not singular top reducible then
     $P \leftarrow P \cup \{\mathcal{S}(\alpha, p') \mid \alpha \in \mathcal{G} \text{ and } \mathcal{S}(\alpha, p') \text{ is regular}\}$ 
     $\mathcal{G} \leftarrow \mathcal{G} \cup \{p'\}$ 
  end if
end while
return  $(\mathcal{G}, H)$ 

```

Figure 1: Pseudo code for a simple version of **SB**.

We can extract a Gröbner basis from a signature Gröbner basis, but that is not the only reason to be interested in signature Gröbner bases. It is also possible to extract a Gröbner basis of the syzygy module of the input basis $\{g_1, \dots, g_m\}$ from a signature Gröbner basis. Note that this is the syzygy module of the *input* basis rather than the syzygy module of a Gröbner basis of the same ideal. The former is in general much harder to compute than the latter.

The key to computing the syzygy module is Theorem 6 which implies that we can determine the initial module of the module of syzygies from looking at those S-pairs and e_i that regular \mathfrak{s} -reduce to zero. If we are computing with full elements of R^m (as opposed to sig-poly pairs) and we have stored the syzygies that were computed (as opposed to storing just their signatures) then those syzygies will form the minimal Gröbner basis of the syzygy module.

THEOREM 6. *Let $\alpha \in R^m$ be a syzygy and let \mathcal{G} be a signature Gröbner basis up to signature $\mathfrak{s}(\alpha)$. Then there exists a $\beta \in R^m$ with $\mathfrak{s}(\beta) \mid \mathfrak{s}(\alpha)$ such that β is an S-pair or has the form e_i and such that β regular \mathfrak{s} -reduces to zero.*

Figure 1 contains pseudo code for a simple version of **SB** that returns a signature Gröbner basis \mathcal{G} and also the initial submodule H of the syzygy module of g_1, \dots, g_m . This pseudo code is intended to succinctly state the essence of **SB** without getting bogged down in any of the complexities of implementation on a computer. Among many other things, a reasonable algorithm would use criteria to eliminate S-pairs (see Section 2.4) and use the sig-poly pair optimization (see Remark 2).

We stated that **SB** proceeds by \mathfrak{s} -reducing S-pairs and then adding the \mathfrak{s} -reduced result to the basis if it is not a syzygy. In Figure 1 we *regular* \mathfrak{s} -reduce the S-pair p and then add the *regular* \mathfrak{s} -reduced result p' to the basis if it is not a syzygy and not singular \mathfrak{s} -reducible. The reason for this is that if p' is singular \mathfrak{s} -reducible then \mathfrak{s} -reduction of p' is going to \mathfrak{s} -reduce p' to zero anyway so we might as well not spend the time on doing that \mathfrak{s} -reduction.

2.4 S-Pair Elimination

Three things can happen when **SB** regular \mathfrak{s} -reduces an S-pair in signature T and gets a result $\gamma \in R^m$.

Syzygy If γ is a syzygy then T is recorded in H .

Singular If γ is singular top \mathfrak{s} -reducible then γ \mathfrak{s} -reduces to zero so γ is discarded.

Basis Otherwise γ is recorded in \mathcal{G} as a new basis element.

For these three cases T is respectively a *syzygy*, *singular* or *basis* signature. This is well defined due to Lemma 4.

If L is a syzygy signature and $L|T$ then T is also a syzygy signature. **SB** eliminates an S-pair p by the *signature criterion* if $\mathfrak{s}(p) \in H$ since then $\mathfrak{s}(p)$ is syzygy.

The *Koszul syzygy* between $\alpha, \beta \in \mathcal{G}$ is $\mathcal{K}(\alpha, \beta) := \bar{\beta}\alpha - \bar{\alpha}\beta$. If $\mathfrak{s}(\beta\alpha) \not\leq \mathfrak{s}(\bar{\alpha}\beta)$ then the Koszul syzygy is *regular*. By “Koszul syzygy” we always mean “regular Koszul syzygy”. **SB** eliminates an S-pair p by the *Koszul criterion* if there exists a Koszul syzygy σ such that $\mathfrak{s}(p) = \mathfrak{s}(\sigma)$. In this case $\mathfrak{s}(p)$ is recorded in H since $\mathfrak{s}(p)$ is syzygy.

A signature T is *predictably syzygy* if $\mathfrak{s}(p) = \mathfrak{s}(\sigma)$ for a Koszul syzygy σ or if there exists a syzygy $\sigma \in R^m$ such that $\mathfrak{s}(\sigma) < T$ and $\mathfrak{s}(\sigma)|T$. The combined effect of the signature criterion and the Koszul criterion is to eliminate all S-pairs in predictably syzygy signatures.

If there are two or more S-pairs in the same signature T , then we only have to regular \mathfrak{s} -reduce one of them as they all regular \mathfrak{s} -reduce to the same thing by Lemma 4. Since \mathfrak{s} -reduction proceeds by decreasing the lead term, we can speed up the process by choosing an S-pair p in signature T whose lead term $\text{lt}(\bar{p})$ is minimal. If $\mathfrak{s}(\mathcal{S}(\alpha, \beta)) = \mathfrak{s}(a\alpha)$, then we get the same result from regular \mathfrak{s} -reducing $\mathcal{S}(\alpha, \beta)$ as for regular \mathfrak{s} -reducing $a\alpha$. So we should choose the $a\alpha \in M$ with minimal lead term $\text{lt}(\bar{a\alpha})$, where M is the set

$$\{a\alpha \mid \alpha \in \mathcal{G}, a \text{ is a monomial and } \mathfrak{s}(a\alpha) = T\}.$$

Note that α might not be involved in any S-pair in signature T . If $a\alpha$ is not regular top \mathfrak{s} -reducible, then T is a singular signature. In this case **RB** eliminates all the S-pairs in signature T by the *singular criterion*. The effect of the singular criterion is to eliminate all S-pairs in singular signatures.

3. THE REWRITE BASIS ALGORITHM

In this section we introduce the **RB** algorithm, which is simpler than yet equivalent to **F5GEN** [9]. Our motivation for studying **RB** is that it contains many published algorithms as special cases including **SB** and **F5**.

The difference between **RB** and **SB** is that **RB** uses the concept of *rewriting* instead of the concept of singular \mathfrak{s} -reduction. This has two main consequences. First, **RB** uses a different criterion for eliminating S-pairs. Second, **RB** adds a regular \mathfrak{s} -reduced $\alpha \in R^m$ to the basis even if α is singular top \mathfrak{s} -reducible. **SB** does not add such α to the basis because such α \mathfrak{s} -reduce to zero anyway. **RB** must add such α to the basis because rewriting requires it.

RB is parametrized by a *rewrite order*. Despite the differences between **RB** and **SB**, PHW [9] show that **GVWHS** (and hence **SB**) is an instantiation of **F5GEN** (and hence **RB**). We simplify their proof in the setting of **SB** and **RB** and extend the result to say that **SB** is the best possible instantiation of **RB** in the sense of Theorem 13.

Figure 2 contains pseudo code for a simple version of **RB**.

3.1 Rewriting and S-pair Elimination

Let \preceq be a *rewrite order*, which means that \preceq is a total order on \mathcal{G} such that $\mathfrak{s}(\alpha)|\mathfrak{s}(\beta) \Rightarrow \alpha \preceq \beta$. Such an order always exists due to our assumption that $\mathfrak{s}(\alpha) \simeq \mathfrak{s}(\beta) \Rightarrow$

```

SimpleRewriteBasisAlgorithm( $\{g_1, \dots, g_m\} \subseteq R$ )
 $\mathcal{G} \leftarrow \emptyset$  ( $\mathcal{G}$  will be the rewrite basis)
 $P \leftarrow \{\mathbf{e}_1, \dots, \mathbf{e}_m\}$  ( $P$  is the set of pending reductions)
 $H \leftarrow \langle 0 \rangle \subseteq R^m$  ( $H$  will be the initial module of syzygies)
while  $P \neq \emptyset$  do
   $p \leftarrow$  an element of  $P$  with  $\leq$ -minimal signature
   $P \leftarrow P \setminus \{p\}$ 
  if not RBEliminated( $p$ ) then
     $p' \leftarrow$  result of regular reducing  $p$ 
    if  $\bar{p}' = 0$  then
       $H \leftarrow H + \langle \mathfrak{s}(p') \rangle$ 
    else
       $P \leftarrow P \cup \{\mathcal{S}(\alpha, p') \mid \alpha \in \mathcal{G} \text{ and } \mathcal{S}(\alpha, p') \text{ is regular}\}$ 
       $\mathcal{G} \leftarrow \mathcal{G} \cup \{p'\}$  (insert  $p'$  into the basis even when  $p'$ 
        is singular top  $\mathfrak{s}$ -reducible)
    end if
  end if
end while
return  $(\mathcal{G}, H)$ 

```

```

RBEliminated(S-pair  $a\alpha - b\beta \in R^m$ )
if  $a\alpha$  is rewritable or  $a\alpha$  is predictably syzygy then
  return true (eliminate S-pair)
end if
if  $b\beta$  is rewritable or  $b\beta$  is predictably syzygy then
  return true (eliminate S-pair)
end if
return false (do not eliminate S-pair)

```

Figure 2: Pseudo code for a simple version of **RB**.

$\alpha = \beta$. A basis element $\alpha \in \mathcal{G}$ is a *rewriter in signature* T if $\mathfrak{s}(\alpha)|T$. If $\mathfrak{s}(a\alpha) = T$ for a monomial a we also say for convenience that $a\alpha$ is a rewriter of T . The \preceq -maximal rewriter in signature T is the *canonical rewriter*. A basis element multiple $a\alpha$ is *rewritable* if α is not the canonical rewriter of $\mathfrak{s}(a\alpha)$. An S-pair $a\alpha - b\beta$ is eliminated by **RB** if $a\alpha$ is predictably syzygy or rewritable, or if $b\beta$ is predictably syzygy or rewritable.

Note that **RB**'s S-pair elimination criterion applies equally to both components $a\alpha$ and $b\beta$ of an S-pair $a\alpha - b\beta$. No criterion in **SB** can eliminate an S-pair based on $\mathfrak{s}(b\beta)$ where $\mathfrak{s}(a\alpha) > \mathfrak{s}(b\beta)$. Can **RB** eliminate some S-pairs that **RB** cannot due to this difference? Since both **SB** and **RB** regular \mathfrak{s} -reduce at most one S-pair in any given signature, what matters is the ability to eliminate *all* S-pairs in a signature. We will prove that **SB** has a stronger S-pair elimination criterion than **RB** does in the sense that if **SB** must regular \mathfrak{s} -reduce an S-pair in a signature then so must **RB**.

3.2 Rewrite bases

\mathcal{G} is a *rewrite basis in signature* T if the canonical rewriter in signature T is not regular top \mathfrak{s} -reducible or if T is a syzygy signature. \mathcal{G} is a *rewrite basis up to signature* T if \mathcal{G} is a rewrite basis for all signatures L such that $L < T$. \mathcal{G} is a *rewrite basis* if it is a rewrite basis in all signatures. We prove Theorem 7 later in this section.

THEOREM 7. **RB** computes a rewrite basis.

LEMMA 8. If \mathcal{G} is a rewrite basis up to signature T then \mathcal{G} is also a signature Gröbner basis up to signature T .

PROOF. Suppose to get a contradiction that \mathcal{G} is a rewrite

basis up to signature T but that it is not a signature Gröbner basis up to signature T . Since the monomial order is a well-order there exists an $\alpha \in R^m$ with minimal signature $\mathfrak{s}(\alpha) < T$ such that α does not \mathfrak{s} -reduce to zero. Then \mathcal{G} is a signature basis up to signature $\mathfrak{s}(\alpha)$ and it is a rewrite basis in signature $\mathfrak{s}(\alpha)$.

Let β be the result of regular \mathfrak{s} -reducing α and let $c\gamma$ be the canonical rewriter in signature $\mathfrak{s}(\alpha)$. Then $\mathfrak{s}(c\gamma) = \mathfrak{s}(\beta)$ and both $c\gamma$ and β are not regular top \mathfrak{s} -reducible, so $\text{lt}(\overline{c\gamma}) = \text{lt}(\overline{\beta})$ by Lemma 4. Perform the singular \mathfrak{s} -reduction step $\beta - c\gamma$. Since $\mathfrak{s}(\beta - c\gamma) < \mathfrak{s}(\alpha)$ and \mathcal{G} is a signature Gröbner basis up to signature $\mathfrak{s}(\alpha)$ we then get that $\beta - c\gamma$ \mathfrak{s} -reduces to zero. Hence α \mathfrak{s} -reduces to zero which is a contradiction. \square

SB computes a minimal signature Gröbner basis so the **SB** basis is a subset of the **RB** basis up to sig-poly pairs. Both **SB** and **RB** have to perform a regular \mathfrak{s} -reduction in all syzygy signatures that are not predictably syzygy, so there are no differences between the two algorithms in terms of how many S-pairs are reduced to zero. For non-syzygy signatures T , both algorithms will add an element to the basis with signature T if and only if they reduce an S-pair in signature T .¹ This implies that if **SB** regular \mathfrak{s} -reduces an S-pair in signature T then so does **RB**. In other words, the S-pairs reduced by **SB** form a subset of the S-pairs reduced by **RB** up to signature. If A is the basis computed by **SB and B is the basis computed by **RB**, then **RB** regular \mathfrak{s} -reduced $|B| - |A| \geq 0$ more S-pairs than **SB** did. If $|A| = |B|$ then **SB** and **RB** reduced the same S-pairs up to signature.**

Our proof of Theorem 7 is based on the following series of lemmas. Lemma 10 connects S-pairs to rewrite bases. Lemma 11 is an important technical lemma that we use here to construct S-pairs and which we use again in Section 4.1. Lemma 12 gives a precise criterion for when **RB** regular \mathfrak{s} -reduces an S-pair in a signature.

LEMMA 9. *Let \mathcal{G} be a rewrite basis up to signature T . Let $a\alpha$ be the canonical rewriter in signature T and let $b\beta$ be a regular top \mathfrak{s} -reducer of $a\alpha$. Then $\mathcal{S}(\alpha, \beta) = a\alpha - b\beta$ and $\mathfrak{s}(\mathcal{S}(\alpha, \beta)) = T$.*

PROOF. If $g := \gcd(a, b)$ then $a\alpha - b\beta = g\mathcal{S}(\alpha, \beta)$. Suppose to get a contradiction that $g \neq 1$. Let $c\gamma$ be the canonical rewriter in signature $\frac{T}{g}$. Then $\gamma \succeq \alpha$ since $\frac{a}{g}\alpha$ is a rewriter in signature $\frac{T}{g}$. Also $\gamma \preceq \alpha$ since $gc\gamma$ is a rewriter in signature T . Hence $\alpha = \gamma$. This is a contradiction since then $\frac{b}{g}\beta$ is a regular top \mathfrak{s} -reducer of $\frac{a}{g}\alpha = c\gamma$ but $c\gamma$ is not regular top \mathfrak{s} -reducible. \square

LEMMA 10. *\mathcal{G} is a rewrite basis up to signature T if \mathcal{G} is a rewrite basis in all signatures $\mathfrak{s}(p) < T$ where p is an S-pair or $p = e_i$.*

PROOF. Suppose to get a contradiction that \mathcal{G} is not a rewrite basis up to signature T . Since the module monomial order $<$ is a well-order there exists a minimal non-syzygy signature $L < T$ such that \mathcal{G} is not a rewrite basis in signature L . Then \mathcal{G} is a rewrite basis up to signature L . Then there exists an S-pair with signature L by Lemma 9 so \mathcal{G} is a rewrite basis in signature L which is a contradiction. \square

¹If an S-pair p regular \mathfrak{s} -reduces to p' then it is true that **SB** will not add p' to the basis if p' is singular top \mathfrak{s} -reducible. However, the singular criterion always eliminates such S-pairs p before the regular \mathfrak{s} -reduction happens.

LEMMA 11. *Let $\alpha \in R^m$, let \mathcal{G} be a rewrite basis up to signature $\mathfrak{s}(\alpha)$ and let t be a regular \mathfrak{s} -reducible term of $\overline{\alpha}$. Let M be the set of $c\gamma$ that regular \mathfrak{s} -reduce t . Let $b\beta$ be the canonical rewriter in signature $L := \min_{c\gamma \in M} \mathfrak{s}(c\gamma)$. Then*

- $b\beta$ is a regular \mathfrak{s} -reducer of t ,
- $b\beta$ is not regular top \mathfrak{s} -reducible,
- $b\beta$ is not rewritable and
- $\mathfrak{s}(b\beta)$ is not syzygy.

PROOF. $b\beta$ is the canonical rewriter in signature $L < \mathfrak{s}(\alpha)$ and \mathcal{G} is a rewrite basis up to signature $\mathfrak{s}(\alpha)$, which establishes the two middle statements.

$b\beta \in M$: Let $c\gamma \in M$ such that $\mathfrak{s}(c\gamma) = L$. Suppose to get a contradiction that $d\delta$ regular top \mathfrak{s} -reduces $c\gamma$. Then $\text{lt}(\overline{d\delta}) = \text{lt}(\overline{c\gamma}) = t$ and $\mathfrak{s}(d\delta) < \mathfrak{s}(c\gamma) < \mathfrak{s}(\alpha)$, so $d\delta \in M$ and $\mathfrak{s}(d\delta) < L$ which is a contradiction. Then $\text{lt}(\overline{b\beta}) = \text{lt}(\overline{c\gamma}) = t$ by Lemma 4 so $b\beta \in M$.

$\mathfrak{s}(b\beta)$ is not syzygy: Suppose to get a contradiction that there exists a syzygy σ such that $\mathfrak{s}(\sigma) = \mathfrak{s}(b\beta)$. Since $\mathfrak{s}(b\beta - \sigma) < \mathfrak{s}(\alpha)$ there exists a top \mathfrak{s} -reducer $c\gamma$ of $b\beta - \sigma$. Then $\mathfrak{s}(c\gamma) \leq \mathfrak{s}(b\beta - \sigma) < \mathfrak{s}(b\beta)$ so $c\gamma$ is a regular top \mathfrak{s} -reducer of $b\beta$ but $b\beta$ is not regular top \mathfrak{s} -reducible. \square

LEMMA 12. *Let \mathcal{G} be a rewrite basis up to signature T . Let $a\alpha$ be the canonical rewriter in signature T . Then **RB** \mathfrak{s} -reduces an S-pair in signature T if and only if $a\alpha$ is regular top \mathfrak{s} -reducible and T is not predictably syzygy.*

PROOF. **if:** Let $b\beta$ be the regular top \mathfrak{s} -reducer of $a\alpha$ from Lemma 11 so that $b\beta$ is not rewritable and $\mathfrak{s}(b\beta)$ is not predictably syzygy. Then $\mathcal{S}(\alpha, \beta) = a\alpha - b\beta$ by Lemma 9 and this S-pair is not eliminated by **RB**.

only if: Let $\mathcal{S}(\alpha, \beta) = a\alpha - b\beta$ such that $T = \mathfrak{s}(a\alpha) > \mathfrak{s}(b\beta)$ and such that **RB** does not eliminate $\mathcal{S}(\alpha, \beta)$. The latter implies that $a\alpha$ is the canonical rewriter in signature T and that T is not predictably syzygy. Observe that $a\alpha$ is regular top \mathfrak{s} -reducible since $b\beta$ regular top \mathfrak{s} -reduces it. \square

PROOF OF THEOREM 7. Suppose that **RB** has terminated with the basis \mathcal{G} . Suppose to get a contradiction that p is an S-pair such that \mathcal{G} is not a rewrite basis in signature $\mathfrak{s}(p)$. Let $a\alpha$ be the canonical rewriter in signature $\mathfrak{s}(p)$. Then $a\alpha$ is regular top \mathfrak{s} -reducible and $\mathfrak{s}(a\alpha) = \mathfrak{s}(p)$ is not syzygy, so **RB** has reduced an S-pair in signature $\mathfrak{s}(p)$ by Lemma 12. Then **RB** has also added a basis element β with signature $\mathfrak{s}(p)$ whereby $\beta = a\alpha$ is the canonical reducer in signature $\mathfrak{s}(p)$. This is a contradiction since β is not regular top \mathfrak{s} -reducible. Then \mathcal{G} is a rewrite basis by Lemma 10. \square

3.3 SB is the Best Possible Instantiation of RB

Let the *sig-lead ratio order* \preceq_r be the order on \mathcal{G} such that $\alpha \preceq_r \beta$ if $r_\alpha \leq r_\beta$ or if $r_\alpha = r_\beta$ and $\mathfrak{s}(\alpha) \leq \mathfrak{s}(\beta)$.

THEOREM 13. ***RB** is equivalent to **SB** when using the sig-lead ratio rewrite order. If **RB** uses any other rewrite order, then **SB** computes a basis that is a subset of the basis computed by **RB** up to sig-poly pairs and **SB** regular \mathfrak{s} -reduces a subset of the S-pairs regular \mathfrak{s} -reduced by **RB** up to signature.*

PROOF. Let **RB** use \preceq_r for rewriting and let \mathcal{G} be the minimal signature Gröbner basis. We will prove that \preceq_r is

a rewrite order and that \mathcal{G} is a rewrite basis. Then **RB** computes \mathcal{G} as it only regular \mathfrak{s} -reduces an S-pair in a non-syzygy signature T if the basis is not already a rewrite basis in signature T . Thus **SB** and **RB** compute the same basis and then the result follows from the statements in Section 3.2.

\preceq_r is a rewrite order: If $\mathfrak{s}(\alpha) \mid \mathfrak{s}(\beta)$ then $r_\alpha \leq r_\beta$ by Lemma 14 and $\mathfrak{s}(\alpha) \leq \mathfrak{s}(\beta)$ so $\alpha \preceq_r \beta$.

\mathcal{G} is a rewrite basis: Let $a\alpha$ be a canonical rewriter. We need to prove that $a\alpha$ is not regular top \mathfrak{s} -reducible. Let γ be the result of regular \mathfrak{s} -reducing $a\alpha$. \mathcal{G} is a signature Gröbner basis, so there exists a $b\beta$ that singular top \mathfrak{s} -reduces γ . Then $r_\alpha \leq r_\beta$ by Lemma 14 so $\alpha = \beta$. \square

LEMMA 14. Let \mathcal{G} be a signature Gröbner basis up to signature $\mathfrak{s}(\alpha) = \mathfrak{s}(b\beta)$ where $\alpha, \beta \in \mathcal{G}$ and a, b are monomials. If $b\beta$ is not regular top \mathfrak{s} -reducible then $r_\alpha \leq r_\beta$.

PROOF. Let γ be the result of regular \mathfrak{s} -reducing $a\alpha$. $b\beta$ and γ are not regular top \mathfrak{s} -reducible so $\text{lt}(\overline{a\alpha}) \geq \text{lt}(\overline{\gamma}) = \text{lt}(\overline{b\beta})$ by Lemma 4. Also $\mathfrak{s}(\gamma) = \mathfrak{s}(a\alpha) = \mathfrak{s}(b\beta)$, so

$$r_\alpha = \frac{\mathfrak{s}(a\alpha)}{\text{lt}(\overline{a\alpha})} \leq \frac{\mathfrak{s}(\gamma)}{\text{lt}(\overline{\gamma})} = \frac{\mathfrak{s}(b\beta)}{\text{lt}(\overline{b\beta})} = r_\beta. \quad \square$$

4. \mathbf{RB}_5 — BRINGING RB CLOSER TO \mathbf{F}_5

In this section we define \mathbf{RB}_5 , which is a special case of **RB**. \mathbf{RB}_5 is more similar to \mathbf{F}_5 than **RB** is. We will show why the differences between \mathbf{RB}_5 and **RB** are not significant, except for the fact that **RB** is more general. In Section 5 we resolve the differences between \mathbf{RB}_5 and \mathbf{F}_5 , which completes our proof that \mathbf{F}_5 is a special case of **RB**.

\mathbf{RB}_5 is specialized from **RB** in three main ways. First, the module monomial order must be position-first, which means that if $i < j$ then $ae_i < be_j$ for all monomials $a, b \in R$. Second, the input basis elements g_1, \dots, g_m must form a regular sequence. Third, the rewrite order must be the \mathbf{F}_5 rewrite order \preceq_r . Let $ae_i := \mathfrak{s}(\alpha)$ and $\mathfrak{s}(be_j) := \beta$. Then $\alpha \preceq_r \beta$ if $i < j$ or if $i = j$ and the total degree of α is strictly smaller than the total degree of b . Break ties arbitrarily.

4.1 Reduction in \mathbf{RB}_5

Let $\alpha \in R^m$ and let t be a term of $\overline{\alpha}$. We can \mathfrak{t} -reduce t by a basis element $\beta \in \mathcal{G}$ if

- there exists a monomial b such that $\text{lt}(\overline{b\beta}) = t$,
- $\mathfrak{s}(b\beta) < \mathfrak{s}(\alpha)$,
- $\mathfrak{s}(b\beta)$ is not syzygy and
- $b\beta$ is not rewritable.

The outcome of the \mathfrak{t} -reduction step is then $\alpha - b\beta$ and β is the \mathfrak{t} -reducer. When β \mathfrak{t} -reduces a term of $\overline{\alpha}$ we also say for convenience that $b\beta$ \mathfrak{t} -reduces α . If $t = \text{lt}(\overline{\alpha})$ then the \mathfrak{t} -reduction step is a top \mathfrak{t} -reduction step.

The first condition on \mathfrak{t} -reducers is the same as for \mathfrak{s} -reduction. The second condition requires $<$ where \mathfrak{s} -reduction requires only \leq , so \mathfrak{t} -reduction is implicitly regular. The final two conditions are not used by **SB**. In Section 5.3 we motivate these extra conditions in terms of S-pair elimination. All \mathfrak{t} -reducers are regular \mathfrak{s} -reducers but not vice versa.

LEMMA 15. Let $\alpha \in R^m$, let t be a term of $\overline{\alpha}$ and let \mathcal{G} be a rewrite basis up to signature $\mathfrak{s}(\alpha)$. Then t is regular \mathfrak{s} -reducible if and only if t is \mathfrak{t} -reducible

PROOF. If t is regular \mathfrak{s} -reducible then by Lemma 11 there exists a regular \mathfrak{s} -reducer $b\beta$ of t such that $\mathfrak{s}(b\beta)$ is not syzygy and $b\beta$ is not rewritable. So $b\beta$ is an \mathfrak{t} -reducer. \square

During regular \mathfrak{s} -reduction there can be a choice of which regular \mathfrak{s} -reducer to use. Lemma 15 shows that the effect of the extra conditions imposed on \mathfrak{t} -reducers is to exclude some of the \mathfrak{s} -reducers from consideration, but it is never the case that all of the \mathfrak{s} -reducers are excluded. The outcome of regular \mathfrak{s} -reduction does not depend on the choice of \mathfrak{s} -reducer up to sig-poly pairs, which proves Corollary 16.

COROLLARY 16. Let $\alpha \in R^m$ and let \mathcal{G} be a rewrite basis up to signature $\mathfrak{s}(\alpha)$. Let α_s be the result of regular \mathfrak{s} -reducing α and let α_t be the result of \mathfrak{t} -reducing α . Then $\mathfrak{s}(\alpha_s) = \mathfrak{s}(\alpha_t)$ and $\overline{\alpha_s} = \overline{\alpha_t}$.

4.2 S-pair elimination in \mathbf{RB}_5

Due to the extra assumptions on the input, \mathbf{RB}_5 can eliminate all S-pairs that reduce to zero by considering only the Koszul syzygies. This is stated already in the \mathbf{F}_5 paper [6]. **RB** uses condition (2) of Theorem 17 while \mathbf{RB}_5 uses condition (4). Theorem 17 implies that these two are equivalent.

THEOREM 17. Let \mathcal{G} be a signature Gröbner basis up to signature e_i . Assume that the input basis element g_i is not a zero divisor of $R/\langle g_1, \dots, g_{i-1} \rangle$. Assume that the module monomial order $<$ on R^m has the property that $be_j < ae_i < ce_k$ for all j, k such that $j < i < k$ and for all monomials $a, b, c \in R$. Then the following statements are equivalent for all monomials $a \in R$.

1. The signature ae_i is syzygy.
2. The signature ae_i is predictably syzygy.
3. $\exists \alpha \in \mathcal{G}$: $\text{ind}(\alpha) < i$ and $\mathfrak{s}(\mathcal{K}(e_i, \alpha)) \mid ae_i$.
4. $\exists \alpha \in \mathcal{G}$: $\text{ind}(\alpha) < i$ and $\text{lt}(\overline{\alpha}) \mid a$.

PROOF. Let $G := \{\alpha \in \mathcal{G} \mid \text{ind}(\alpha) < i\}$, $\overline{G} := \{\overline{\alpha} \mid \alpha \in G\}$ and $F := \langle g_1, \dots, g_{i-1} \rangle$. That (3) \Rightarrow (2) \Rightarrow (1) is immediate.

\overline{G} is a Gröbner basis of F : If $\alpha \in G$ then $\alpha_j = 0$ for $j \geq i$ so it is immediate that $\langle \overline{G} \rangle \subseteq \langle g_1, \dots, g_{i-1} \rangle$. To prove the other inclusion, let $f \in F$. We need to prove that f reduces to zero on classic polynomial reduction by \overline{G} .

There exists an $\alpha \in R^m$ such that $f = \overline{\alpha}$ and $\alpha_j = 0$ for $j \geq i$. Then $\text{ind}(\alpha) < i$ so $\mathfrak{s}(\alpha) < e_i$ whereby α \mathfrak{s} -reduces to zero when \mathfrak{s} -reducing by \mathcal{G} . All the \mathfrak{s} -reducers $\gamma \in \mathcal{G}$ in that \mathfrak{s} -reduction must have $\mathfrak{s}(\gamma) \leq \mathfrak{s}(\alpha)$ whereby $\text{ind}(\gamma) < i$. Thus α also \mathfrak{s} -reduces to zero when \mathfrak{s} -reducing by G . Then $\overline{\alpha} = f$ reduces to zero on classic polynomial reduction by \overline{G} .

(1) \Rightarrow (4): Let $\beta \in R^m$ be a syzygy such that $\mathfrak{s}(\beta) = ae_i$. Then $\beta_j = 0$ for all $j > i$ whereby $\beta_i g_i \in F$ since

$$0 = \overline{\beta} = \sum_{j=1}^m \beta_j g_j = \beta_i g_i + \sum_{j=1}^{i-1} \beta_j g_j.$$

Hence $\beta_i \in F$ as g_i is not a zero divisor. As \overline{G} is a Gröbner basis of F there exists an $\alpha \in G$ such that $\text{lt}(\overline{\alpha}) \mid \text{lt}(\beta_i) = a$.

(4) \Rightarrow (3): Let $\alpha \in \mathcal{G}$ such that $\text{ind}(\alpha) < i$ and $\text{lt}(\overline{\alpha}) \mid a$. By definition $\mathcal{K}(e_i, \alpha) = \overline{\alpha} e_i - g_i \alpha$. As $\text{lt}(\overline{\alpha}) e_i > \text{lt}(g_i) \mathfrak{s}(\alpha)$ we see that $\mathfrak{s}(\mathcal{K}(e_i, \alpha)) = \text{lt}(\overline{\alpha}) e_i \mid ae_i$. \square

$\phi_i \in \mathcal{G}$	reduced from	$\text{lt}(\overline{\phi_i})$	$\mathfrak{s}(\phi_i)$
ϕ_1	e_1	y^3	e_1
ϕ_2	e_2	xyz	e_2
ϕ_3	$y^2\phi_2 - xz\phi_1 = \mathcal{S}(\phi_2, \phi_1)$	x^3z^2	y^2e_2
ϕ_4	e_3	yz^2	e_3
ϕ_5	$x\phi_3 - z\phi_2 = \mathcal{S}(\phi_3, \phi_2)$	xz^3	xe_3
ϕ_6	$y^2\phi_3 - z^2\phi_1 = \mathcal{S}(\phi_3, \phi_1)$	x^2z^3	y^2e_3
ϕ_7	$y\phi_5 - z^2\phi_2 = \mathcal{S}(\phi_5, \phi_2)$	x^2y^2t	xye_3
ϕ_8	$x\phi_5 - \phi_6 = \mathcal{S}(\phi_5, \phi_6)$	z^5	x^2e_3
ϕ_9	$x\phi_6 - z\phi_3 = \mathcal{S}(\phi_6, \phi_3)$	x^4zt	xy^2e_3
ϕ_{10}	$y\phi_8 - z^3\phi_4 = \mathcal{S}(\phi_8, \phi_4)$	x^3y^2t	x^2ye_3
ϕ_{11}	$x^3\phi_4 - y\phi_3 = \mathcal{S}(\phi_4, \phi_3)$	x^4yt	x^3e_3
ϕ_{12}	$z\phi_{11} - x^3\phi_2 = \mathcal{S}(\phi_{11}, \phi_2)$	x^3zt^3	x^3ze_3
ϕ_{13}	$y\phi_{10} - x^3\phi_1 = \mathcal{S}(\phi_{10}, \phi_1)$	x^5zt	$x^2y^2e_3$
ϕ_{14}	$x\phi_{12} - \phi_9 = \mathcal{S}(\phi_{12}, \phi_9)$	x^4t^4	x^4ze_3

Figure 3: Computations for \mathbf{RB}_5 in Example 19.

4.3 Examples

Example 18 shows a homogeneous ideal where \mathbf{RB}_5 (and hence \mathbf{F}_5) computes a larger basis than \mathbf{SB} does. Example 19 shows details of how \mathbf{RB}_5 computes a basis.

EXAMPLE 18. Consider the homogenized *Eco-6* ideal using the graded reverse lexicographic monomial order. Here \mathbf{RB}_5 computes a basis with 100 elements while \mathbf{SB} computes a basis with 87 elements. Both \mathbf{RB}_5 and \mathbf{SB} compute a basis element α with $\mathfrak{s}(\alpha) = x_0x_2e_5$. \mathbf{RB}_5 additionally computes basis elements β and γ with $\mathfrak{s}(\beta) = x_0x_2x_3^2e_5$ and $\mathfrak{s}(\gamma) = x_0x_2x_3x_4e_5$. As β and γ are not necessary to have a signature Gröbner basis, we can replacing them with a single element with signature $x_0x_2x_3e_5$ and thus get a rewrite basis that contains fewer elements than the basis computed by \mathbf{RB}_5 . Furthermore, the basis computed by \mathbf{SB} is not a rewrite basis with respect to the \mathbf{F}_5 rewrite order.

EXAMPLE 19. Let κ be the finite field with 13 elements and let $R := \kappa[x, y, z, t]$. Let $<$ be the graded reverse lexicographic monomial order. Consider the three input elements

$$\begin{aligned} g_1 &:= -2y^3 - x^2z - 2x^2t - 3y^2t, & g_2 &:= 3xyz + 2xyt, \\ g_3 &:= 2xyz - 2yz^2 + 2z^3 + 4yzt. \end{aligned}$$

Figure 3 shows the τ -reductions performed by \mathbf{RB}_5 to compute the rewrite basis $\{\phi_1, \dots, \phi_{14}\}$. \mathbf{SB} regular \mathfrak{s} -reduce the same S -pairs except for the ones that lead to ϕ_{10} and ϕ_{13} , so the basis computed by \mathbf{SB} contains 12 elements. The S -pair $\mathcal{S}(\phi_8, \phi_4)$ leads to ϕ_{10} and it is eliminated by the singular criterion since $\mathfrak{s}(x\phi_7) = \mathfrak{s}(\mathcal{S}(\phi_8, \phi_4))$ and $x\phi_7$ is not regular top \mathfrak{s} -reducible. The S -pair $\mathcal{S}(\phi_{10}, \phi_1)$ leads to ϕ_{13} and it is not considered by \mathbf{SB} since ϕ_{10} is not part of \mathbf{SB} 's basis.

5. \mathbf{F}_5 IS EQUIVALENT TO \mathbf{RB}_5

The \mathbf{F}_5 paper [6] assumes that the input basis elements g_1, \dots, g_m are homogeneous. In this section we show that \mathbf{F}_5 and \mathbf{RB}_5 become the same algorithm with this assumption.

5.1 The Sig-Poly Pair Optimization

The notation for \mathbf{RB}_5 concerns elements $\alpha \in R^m$ while the \mathbf{F}_5 paper uses notation based on sig-poly pairs $(\mathfrak{s}(\alpha), \overline{\alpha})$. We believe that the \mathbf{SB} notation makes signature computations easier to understand and reason about. However, the

overhead in time and space for computations on elements of R^m are higher than for sig-poly pairs.

Consider that \mathbf{RB}_5 never performs singular \mathfrak{s} -reduction steps, so we never need to know any of the terms in R^m other than the leading one — the signature. The singular \mathfrak{s} -reduction steps appear only in theorems and proofs. So it is possible to apply the *sig-poly pair optimization* when implementing \mathbf{RB}_5 (and \mathbf{RB} and \mathbf{SB}), which involves replacing $\alpha \in R^m$ with the sig-poly pair $(\mathfrak{s}(\alpha), \overline{\alpha})$. In \mathbf{F}_5 this idea is used in the notation while in \mathbf{RB}_5 we present it as an optimization. The outcome is the same.

5.2 \mathbf{F}_5 maintains a list of rewriting rules

The \mathbf{F}_5 paper [6] does not contain the phrases “canonical rewriter” and “rewrite order”. Instead, there is a pseudo code function named `Rewritten` that takes a signature as a parameter and returns the canonical rewriter in that signature. The concept of the canonical rewriter is thus represented implicitly in the \mathbf{F}_5 paper as that thing which `Rewritten` returns. We do not specify how the \mathbf{F}_5 rewrite order breaks ties since \mathbf{F}_5 's implicit rule for this seems more an arbitrary outcome of how the pseudo code was written than an intentional choice (it is not a function of signature).

In \mathbf{F}_5 there is a list of *rewriting rules*. A rewriting rule is added to the front of the list when an S -polynomial is calculated. When `Rewritten` is called, it goes through the rewriting rules and returns the first rewriter it finds. The first rewriter found is also the rewriter that was most recently added to the basis, so the list of rewriting rules represents the rewrite order \preceq . S -polynomials are only calculated for S -pairs whose signature has the current index and the current total degree, so the rewriting rules are added to the front of the list first in order of index of the signature and then in order of total degree of the signature. So \mathbf{F}_5 is indeed using what we have called the \mathbf{F}_5 rewrite order.

The only exception to this description is that \mathbf{F}_5 does not record $\alpha \in \mathcal{G}$ into the list of rewriters if $\mathfrak{s}(\alpha) \simeq e_i$. However, this does not change whether a basis element multiple is rewritable, so this exception does not matter.

5.3 \mathbf{F}_5 is written to resemble \mathbf{F}_4 reduction

\mathbf{F}_4 is an algorithm that speeds up classic polynomial reduction by using a symbolic preprocessing step to turn the reduction of many polynomials into row reduction of a single matrix [4]. The more reductions in the same homogeneous degree that can be done at the same time, the larger the matrix becomes and the more of a speed up there is from using \mathbf{F}_4 over using classic polynomial reduction. Albrecht and Perry describe a variation of \mathbf{F}_5 that uses \mathbf{F}_4 reduction [1].

The pseudo code in the \mathbf{F}_5 paper is written without \mathbf{F}_4 reduction. However, it is still written in a way that heavily suggests how to adapt the symbolic preprocessing step from \mathbf{F}_4 to \mathbf{F}_5 . For example, the \mathbf{F}_5 pseudo code is written to process all S -pairs in a homogeneous degree at the same time, which is important when using \mathbf{F}_4 . This is done by having a set P of all S -pairs in the current homogeneous degree. In \mathbf{F}_4 τ -reduction of all the S -pairs in P would then be carried out simultaneously within one big matrix.

The above scheme for τ -reducing all S -pairs in a homogeneous degree does not succeed in getting *all* of the S -pairs in that homogeneous degree. When reducing an S -pair to a new basis element α , it is possible that α can form an S -pair in the same homogeneous degree with another basis element.

Since \mathbf{F}_5 requires the input ideal to be homogeneous, this happens if and only if there is another basis element β such that $b\beta$ would top τ -reduce α except that $\mathfrak{s}(b\beta) > \mathfrak{s}(\alpha)$. Call such an S-pair $\mathcal{S}(\beta, \alpha)$ a *late S-pair*. The late S-pairs will not initially be included in P since they only appear after some τ -reductions have already been carried out. This is unfortunate when using \mathbf{F}_4 τ -reduction because for \mathbf{F}_4 τ -reduction we want to identify *all* of the S-pairs in each homogeneous degree up front — including the late ones.

In \mathbf{F}_5 there is a mechanism for discovering late S-pairs while τ -reducing another S-pair. The mechanism is that \mathbf{F}_5 allows reducers to increase the signature of the reduction. In that case \mathbf{F}_5 will split the τ -reduction into two τ -reductions — one τ -reduction in the old signature that continues to be carried out and another τ -reduction in the new signature that is recorded in P . These higher signature τ -reductions are precisely the late S-pairs and this mechanism would allow an \mathbf{F}_4 version of \mathbf{F}_5 to get all the S-pairs in a degree before doing any τ -reduction. Note that it is possible for such a variant of \mathbf{F}_5 to detect more late S-pairs than actually exist because the symbolic preprocessing step does not take account of cancellations.

We have seen that the only effect of \mathbf{F}_5 's higher signature reducers is to indicate how to translate \mathbf{F}_5 into a variant that uses \mathbf{F}_4 reduction. The pseudo code for reduction in \mathbf{F}_5 looks different from τ -reduction because that pseudo code is not just doing τ -reduction — it is also constructing late S-pairs. This also explains the purpose of the extra conditions on τ -reducers — they are S-pair elimination criteria. So what \mathbf{F}_5 is doing remains equivalent to τ -reduction, which proves that \mathbf{F}_5 is a special case of \mathbf{RB}_5 despite \mathbf{F}_5 's higher-signature reducers and its multiple results of reduction.

6. TERMINATION OF RB AND HENCE \mathbf{F}_5

The proof of termination in the \mathbf{F}_5 paper [6] is incorrect. Proving termination of \mathbf{F}_5 had been an open problem for a decade before it was settled by Galkin [7]. PHW later proved that $\mathbf{F5GEN}$ (and hence \mathbf{RB}) terminates [9], which implies that \mathbf{F}_5 terminates. Our proof of termination is similar to but significantly simpler than Galkin's and PHW's.

THEOREM 20. *The \mathbf{RB} algorithm terminates.*

PROOF. Let ϕ_1, ϕ_2, \dots be the sequence of basis elements computed by \mathbf{RB} and let $\mathcal{G} := \{\phi_1, \phi_2, \dots\}$. \mathbf{RB} processes S-pairs in increasing order of signature so $\mathfrak{s}(\phi_1) < \mathfrak{s}(\phi_2) < \dots$. Partition \mathcal{G} into sets $R_r := \{\phi_i \mid r_{\phi_i} = r\}$. We prove that there are only finitely many non-empty sets R_r and that each R_r is finite. Hence \mathcal{G} is finite whereby \mathbf{RB} terminates.

Only finitely many R_r are non-empty: If $\alpha \in \mathcal{G}$ then $\mathcal{G}_\alpha := \{\beta \in \mathcal{G} \mid \mathfrak{s}(\beta) < \mathfrak{s}(\alpha)\}$ is a signature Gröbner basis up to signature $\mathfrak{s}(\alpha)$. Call $\alpha \in \mathcal{G}$ *minimal* if there is no other $\beta \in \mathcal{G}$ such that $\mathfrak{s}(\beta) \mid \mathfrak{s}(\alpha)$ and $\text{lt}(\bar{\beta}) \mid \text{lt}(\bar{\alpha})$. It follows from Lemma 21 that a non-minimal $\alpha \in \mathcal{G}$ is top \mathfrak{s} -reducible by \mathcal{G}_α . No basis element in \mathbf{RB} is regular top \mathfrak{s} -reducible, so if $\alpha \in \mathcal{G}$ is non-minimal, then α must be singular top \mathfrak{s} -reducible by \mathcal{G}_α . Thus there exists a $\beta \in \mathcal{G}_\alpha$ and a monomial $m \neq 1$ such that $\mathfrak{s}(m\beta) = \mathfrak{s}(\alpha)$ and $\text{lt}(\overline{m\beta}) = \text{lt}(\bar{\alpha})$ whereby α and β lie in the same set R_r . This shows that there are exactly as many non-empty sets R_r as there are minimal basis elements in \mathcal{G} . Both R and R^m are Noetherian, so there are only finitely many minimal basis elements. Hence there are also only finitely many non-empty sets R_r .

Each R_r is finite: We prove by induction on the finitely many non-empty sets R_r that each R_r is finite. Assume by induction that all sets $R_{r'}$ with $r' < r$ are finite. We need to prove that R_r is finite. The base case is immediate.

Let $\gamma \in R_r$ and let $\mathcal{S}(\alpha, \beta) = a\alpha - b\beta$ be the S-pair that \mathbf{RB} regular \mathfrak{s} -reduced to get γ where $\mathfrak{s}(a\alpha) > \mathfrak{s}(b\beta)$. Let c be the monic greatest common divisor of $\text{lt}(\bar{\alpha})$ and $\text{lt}(\bar{\beta})$. Then $\frac{1}{c}\text{lt}(\bar{\beta}) \mathfrak{s}(\alpha) = \mathfrak{s}(a\alpha) > \mathfrak{s}(b\beta) = \frac{1}{c}\text{lt}(\bar{\alpha}) \mathfrak{s}(\beta)$ so $r_\beta > r_\alpha$. Also $\text{lt}(\bar{\gamma}) < \text{lt}(\overline{a\alpha})$ and $\mathfrak{s}(\gamma) = \mathfrak{s}(a\alpha)$ so

$$r = \frac{\mathfrak{s}(\gamma)}{\text{lt}(\bar{\gamma})} > \frac{\mathfrak{s}(a\alpha)}{\text{lt}(\overline{a\alpha})} = \frac{\mathfrak{s}(\alpha)}{\text{lt}(\bar{\alpha})} > \frac{\mathfrak{s}(\beta)}{\text{lt}(\bar{\beta})}. \quad (1)$$

Hence $\alpha \in R_{r'}$ and $\beta \in R_{r''}$ where $r', r'' < r$, so R_r contains at most as many elements as there are pairs of elements in the set $\cup_{r' < r} R_{r'}$ which is finite by induction. There may also be one more element with signature e_i since those elements do not come from S-pairs. So by induction R_r is finite. \square

LEMMA 21 (EDER AND PERRY [3]). *Let $\alpha \in R^m$ and let \mathcal{G} be a signature Gröbner basis up to signature $\mathfrak{s}(\alpha)$. If there exists a $\beta \in \mathcal{G}$ such that $\text{lt}(\bar{\beta}) \mid \text{lt}(\bar{\alpha})$ and $\mathfrak{s}(\beta) \mid \mathfrak{s}(\alpha)$ then α is top \mathfrak{s} -reducible by \mathcal{G} .*

7. REFERENCES

- [1] Albrecht, M. and Perry, J. F4/5. 2010. <http://arxiv.org/abs/1006.4933>.
- [2] Arri, A. and Perry, J. The F5 Criterion revised. *Journal of Symbolic Computation*, 46(2):1017–1029, June 2011.
- [3] Eder, C. and Perry, J. Signature-based Algorithms to Compute Gröbner Bases. In *Proceedings of the 2011 international symposium on Symbolic and algebraic computation*, pages 99–106, 2011.
- [4] Faugère, J.-C. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra*, 139(1–3):61–88, June 1999.
- [5] Faugère, J.-C. A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). pages 75–83, 2002.
- [6] Faugère, J.-C. A new efficient algorithm for computing Gröbner bases without reduction to zero F5. In *Proceedings of the 2002 international symposium on Symbolic and algebraic computation*, pages 75–82, July 2002. Revised version from <http://fgbrs.lip6.fr/jcf/Publications/index.html>.
- [7] Galkin, V. Termination of original F5. <http://arxiv.org/abs/1203.2402>, 2012.
- [8] Gao, S., Volny IV, F., and Wang, D. A new algorithm for computing Groebner bases. <http://eprint.iacr.org/2010/641>, 2010.
- [9] Pan, S., Hu, Y., and Wang, B. The Termination of Algorithms for Computing Gröbner Bases. 2012. <http://arxiv.org/abs/1202.3524>.
- [10] Roune, B. H. and Stillman, M. Practical Gröbner basis computation. In *Proceedings of the 2011 international symposium on Symbolic and algebraic computation*.
- [11] Roune, B. H. and Stillman, M. Practical signature Gröbner basis computation. In preparation.
- [12] Frank Volny IV. *New Algorithms for Computing Gröbner Bases*. PhD thesis, Clemson University, 2011. http://etd.lib.clemson.edu/documents/1306872881/Volny_clemson_0050D_11180.pdf.