A NEW ALGORITHM FOR COMPUTING GRÖBNER BASES

SHUHONG GAO, FRANK VOLNY IV, AND MINGSHENG WANG

Abstract. Buchberger's algorithm for computing Gröbner bases was introduced in 1965, and subsequently there have been extensive efforts in improving its efficiency. Major algorithms include F4 (Faugère 1999), XL (Courtois et al. 2000) and F5 (Faugère 2002). F5 is believed to be the fastest algorithm known in the literature. Most recently, Gao, Guan and Volny (2010) introduced an incremental algorithm (G2V) that is two to ten times faster than F5. In this paper, a new algorithm is presented that matches Buchberger's algorithm in simplicity and yet is more flexible than G2V. Given a list of polynomials, the new algorithm computes simultaneously a Gröbner basis for the ideal generated by the polynomials and a Gröbner basis for the leading terms of the syzygy module of the polynomials. For any term order for the ideal, one may vary the term order for the syzygy module. Under one term order for the syzygy module, the new algorithm specializes to the G2V algorithm, and under another term order for the syzygy module, the new algorithm may be several times faster than G2V, as indicated by computer experiments on benchmark examples.

1. Introduction

Polynomial systems are ubiquitous in mathematics, science and engineering, and Gröbner basis theory is one of the most powerful tools for solving polynomial systems. Buchberger introduced in 1965 the first algorithm for computing Gröbner bases, and it has been implemented in most computer algebra systems (e.g., Maple, Mathematica, Magma, Sage, Singular, Macaulay 2, CoCoA, etc). Computing Gröbner bases is a basic routine that is essential in many computational tasks in algebra and algebraic geometry. As time has witnessed the importance of Gröbner bases, Bruno Buchberger was awarded the 2007 ACM Paris Kanellakis Theory and Practice Award and the Golden Medal of Honor for his original development of Gröbner basis theory.

There has been extensive effort in finding more efficient algorithms for computing Gröbner bases. In Buchberger's original algorithm (1965, [1]), one has to reduce many "useless" S-polynomials (i.e., those that reduce to 0 via long division), and each reduction is time consuming. It is natural to avoid useless reductions as much as possible. Buchberger [2, 3] discovered two simple criteria for detecting useless S-polynomials. Note that a reduction of an S-polynomial to 0 corresponds to a syzygy (for the initial list of polynomials). Möller, Mora and Traverso (1992, [13] go a step further to present an algorithm using the full module of syzygies, however, their algorithm is not very efficient. Faugère (2002, [8]) introduced the

Date: October 9, 2010.

Gao and Volny were partially supported by National Science Foundation under Grants DMS-1005369 and CCF-0830481, and Wang was partially supported by the National Science Foundation of China under Grant 60970134 and by 973 Project under Grant 2007CB311201.

idea of signatures and rewriting rules that can detect many useless S-polynomials hence saving a significant amount of time that would be used in reducing them. In fact, for a regular sequence of polynomials, his algorithm F5 detects all useless reductions. By computer experiments, Faugère showed that his algorithm F5 is many times faster than previous algorithms. In fact, Faugère and Joux (2003, [9]) solved the first Hidden Field Equation (HFE) Cryptosystem Challenge which involves a system of 80 polynomial equations with 80 variables over the binary field (1996, [14]). Since F5 seems difficult to both understand and implement, there have been several papers trying to simplify and improve F5; see Eder and Perry (2009, [6]), Sun and Wang (2009, [15]), and Hashemi and Ars (2010, [11]).

In another direction of research, one tries to speed up the reduction step. Lazard (1983, [12]) pointed out the connection between Gröbner bases and linear algebra, that is, a Gröbner basis can be computed by Gauss elimination of a Sylvester matrix. The XL algorithm of Courtois et al. (2000, [4]) is an implementation of this Sylvester matrix, which is recently improved by Ding et al. (2008, [5]). A more clever approach is the F4 algorithm of Faugère (1999, [7]). F4 is an efficient method for reducing several S-polynomials simultaneously where the basic idea is to apply fast linear algebra methods to the submatrix of the Sylvester matrix consisting of only those rows that are needed for the reductions of a given list of S-polynomials. This method benefits from the efficiency of fast linear algebra algorithms. The main problem with this approach, however, is that the memory usage grows too quickly, even for medium systems of polynomials.

Our contribution in this paper is the presentation of a new algorithm that matches Buchberger's algorithm in simplicity yet is much faster than F5. The main difference between our algorithm and F5 is that we don't use rewriting rules. In our previous work [10], we presented an algorithm (G2V) that is incremental in the same fashion as F5 but is much simpler and 2 to 10 times faster than F5 and F5C on benchmark systems. This paper describes an extension of the G2V algorithm. We extend the notion of a signature so that G2V is considered a special case. We remove the necessity of running the algorithm incrementally and allow a Gröbner basis for $\langle g_1, \ldots, g_m \rangle$ to be computed in one-shot, incrementally, or as a hybrid of the two. And just as reductions to zero in G2V provided more information about the colon ideal (and thus saving further computations), the present algorithm uses reductions to zero to discover more information about the syzygy module in order to prevent reductions to zero. Our new algorithm provides an extra layer of flexibility. Unless one needs to compute the syzygy module under a certain term order, one is free to choose the term order of the syzygy module to maximize performance. We show that under the POT order, this new algorithm is exactly the same as G2V. But under certain other orders, this new algorithm performs 2 to 10 times faster than G2V and thus something like 4 to 20 times faster than F5 and F5C.

The paper is organized as follows. In Section 2, we introduce the basic concepts and theory for our algorithm. In particular, we define signatures, regular top-reductions, super top-reductions, and the concept of eventually super top-reducible. We also introduce J-pairs that are in some sense similar to S-polynomials. Then we characterize Gröbner bases in term of J-pairs in a similar fashion as Buchberger's characterization in term of S-polynomials. Our characterization goes a step further, that is, it also tells us when we have a Gröbner basis for the corresponding

syzygy module. In Section 3, we present our algorithm and prove its correctness. The problem of finite termination of our algorithm is still open. We present computer experiments of our algorithm that shows how the algorithm perform under different term orders for the syzygy module. Finally, in Section 4, we show how our algorithm can be adapted to compute Gröbner bases for modules and for polynomials over quotient rings, which would allow one to design more flexible incremental algorithms.

2. Theory

Let $R = \mathbb{F}[x_1, \ldots, x_n]$ be a polynomial ring over a field $\mathbb{F}$ with n variables. Given polynomials $g_1, \ldots, g_m \in R$, we wish to compute a Gröbner basis for the ideal

$$I = \langle g_1, \ldots, g_m \rangle = \{u_1 g_1 + \cdots + u_m g_m : u_1, \ldots, u_m \in R\} \subseteq R \tag{1}$$

with respect to some term order on R. Define

$$\mathbf{H} = \{(u_1, \ldots u_m) \in R^m : u_1 g_1 + \cdots + u_m g_m = 0\}, \tag{2}$$

called **the syzygy module** of $\mathbf{g} = (g_1, \ldots, g_m)$. We would like to develop an algorithm that computes Gröbner bases for both I and $\mathbf{H}$. Note that elements of R^m are viewed as row vectors and are denoted by bold letters say $\mathbf{g}, \mathbf{u}$ etc. We consider the following R-submodule of $R^m \times R$:

$$M = \left\{(\mathbf{u}, v) \in R^m \times R : \mathbf{u}\mathbf{g}^t = v\right\}. \tag{3}$$

We define $\mathbf{E}_i \in R^m$ to be the i^{th} unit vector, that is $(\mathbf{E}_i)_j = \delta_{ij}$. Note that a monomial (or a term) in R is of the form

$$x^\alpha = \prod_{i=1}^{n} x_i^{a_i}$$

where $\alpha = (a_1, \ldots, a_n) \in \mathbb{N}^n$ is any vector of non-negative integers, and a term in R^m is of the form

$$x^\alpha \mathbf{E}_i$$

where $1 \leq i \leq m$ and $\alpha \in \mathbb{N}^n$. The R-module M is generated by

$$(\mathbf{E}_1, g_1), (\mathbf{E}_2, g_2), \ldots, (\mathbf{E}_m, g_m). \tag{4}$$

Fix any term order $\prec_1$ on R and any term order $\prec_2$ on R^m. We emphasize that the order $\prec_2$ may or may not be related to $\prec_1$ in the theory below, though $\prec_2$ is usually an extension of $\prec_1$ to R^m in implementation. For the sake of convenience, we shall use the following convention for leading terms:

$$\mathrm{lm}(v) = \mathrm{lm}_{\prec_1}(v), \quad \mathrm{lm}(\mathbf{u}) = \mathrm{lm}_{\prec_2}(\mathbf{u})$$

for any $v \in R$ and $\mathbf{u} \in R^m$. Note that, for $v \in R$, $\mathrm{lm}(v)$ is a monomial x^α, while, for $\mathbf{u} \in R^m$, $\mathrm{lm}(\mathbf{u})$ is a term $x^\alpha \mathbf{E}_i$ for some $\alpha \in \mathbb{N}^n$ and $1 \leq i \leq m$. We make the convention that if $v = 0$ then $\mathrm{lm}(v) = 0$; similarly for $\mathrm{lm}(\mathbf{u})$. This should not cause any confusion, but the reader should keep the two different orders in mind.

For any $(\mathbf{u}, v) \in R^m \times R$, we call $\mathrm{lm}(\mathbf{u})$ the **signature** of $(\mathbf{u}, v)$. This is similar to the signature used in F5. Suppose $(\mathbf{u}_1, v_1), (\mathbf{u}_2, v_2) \in R^m \times R$ are two pairs with v_1 and v_2 both nonzero. Let

$$t = \mathrm{lcm}(\mathrm{lm}(v_1), \mathrm{lm}(v_2)), \quad t_1 = \frac{t}{\mathrm{lm}(v_1)}, \quad t_2 = \frac{t}{\mathrm{lm}(v_2)}.$$

Suppose $\max(t_1 \mathrm{lm}(\mathbf{u}_1), t_2 \mathrm{lm}(\mathbf{u}_2)) = t_i \mathrm{lm}(\mathbf{u}_i)$ where $i = 1$ or 2. Then

- $t_i(\mathbf{u}_i, v_i) = (t_i\mathbf{u}_i, t_i v_i)$ is called a **J-pair** of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$;
- $t_i \mathrm{lm}(\mathbf{u}_i)$ is called the **J-signature** of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$;

where J means "joint" of the two pairs. When $t_1\mathrm{lm}(\mathbf{u}_1) = t_2\mathrm{lm}(\mathbf{u}_2)$, we can pick either $t_1(\mathbf{u}_1, v_1)$ or $t_2(\mathbf{u}_2, v_2)$ as a J-pair. However, the J-signature of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$ is unique.

We should mention that the S-polynomial of v_1 and v_2 is $t_1 v_1 - c t_2 v_2$ where $c = \mathrm{lc}(v_1)/\mathrm{lc}(v_2)$. Hence the monomials t_1 and t_2 used in our J-pair is the same as those used in the S-polynomial. In the case of S-polynomials, the goal is to cancel the leading terms of v's. In our J-pairs, the leading terms of v's are not cancelled, but will be cancelled in later top-reductions (for most cases). In our algorithm below, we may produce many J-pairs that have the same J-signature, but we only keep one per distinct signature. Note that we never calculate the J-pair of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$ when either v_1 or v_2 is zero.

Next we define top-reductions in $R^m \times R$. Let $(\mathbf{u}_1, v_1), (\mathbf{u}_2, v_2) \in R^m \times R$ be any two pairs. When v_2 is nonzero, we say $(\mathbf{u}_1, v_1)$ is top-reducible by $(\mathbf{u}_2, v_2)$ if the following two conditions are satisfied:

(i) v_1 is nonzero and $\mathrm{lm}(v_2)$ divides $\mathrm{lm}(v_1)$; and
(ii) $\mathrm{lm}(t\mathbf{u}_2) \preceq \mathrm{lm}(\mathbf{u}_1)$ where $t = \mathrm{lm}(v_1)/\mathrm{lm}(v_2)$.

The corresponding **top-reduction** is then

$$(\mathbf{u}_1, v_1) - ct(\mathbf{u}_2, v_2) = (\mathbf{u}_1 - ct\mathbf{u}_2, v_1 - ctv_2),$$

where $c = \mathrm{lc}(v_1)/\mathrm{lc}(v_2)$. The effect of a top-reduction is that the leading monomial in the v-part is canceled without increasing the signature of $(\mathbf{u}_1, v_1)$. Such a top-reduction is called **regular**, if

$$\mathrm{lm}(\mathbf{u}_1 - ct\mathbf{u}_2) = \mathrm{lm}(\mathbf{u}_1),$$

and **super** otherwise. So the signature of $(\mathbf{u}_1, v_1)$ remains the same under a regular top-reduction but becomes smaller under a super top-reduction. A super top-reduction happens if

$$\mathrm{lm}(t\mathbf{u}_2) = \mathrm{lm}(\mathbf{u}_1) \text{ and } \frac{\mathrm{lc}(\mathbf{u}_1)}{\mathrm{lc}(\mathbf{u}_2)} = \frac{\mathrm{lc}(v_1)}{\mathrm{lc}(v_2)}.$$

When $v_2 = 0$, we say that $(\mathbf{u}_1, v_1)$ is **top-reducible** by $(\mathbf{u}_2, 0)$ if $\mathbf{u}_1$ and $\mathbf{u}_2$ are both nonzero and $\mathrm{lm}(\mathbf{u}_2)$ divides $\mathrm{lm}(\mathbf{u}_1)$. In this case, we could use $(\mathbf{u}_2, 0)$ to reduce the signature of $(\mathbf{u}_1, v_1)$ without increasing the leading term of v_1 (even if $v_1 = 0$); such a top-reduction is always called super. We note that a pair $(\mathbf{u}_1, 0)$ is never top-reducible by $(\mathbf{u}_2, v_2)$ for $v_2 \neq 0$. In our algorithm below, we only detect super top-reductions of the two kinds defined here, but never actually perform super top-reductions. We should mention that the top-reductions used in F5 correspond to regular top-reductions in our sense, but some of our regular top-reductions are not allowed in F_5 (e.g. when $\mathrm{lm}(\mathbf{u}_1) = t\mathrm{lm}(\mathbf{u}2)$).

Lemma 1. *Let t be a monomial in R. If a pair $t(\mathbf{u}_1, v_1)$ is (regular) top-reducible by $(\mathbf{u}_2, v_2)$, where both v_1 and v_2 are nonzero, then $t_1(\mathbf{u}_1, v_1)$ is a J-pair of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$ where*

$$t_1 = \frac{lcm(lm(v_1), lm(v_2))}{lm(v_1)} = \frac{lm(v_2)}{\gcd(lm(v_1), lm(v_2))}$$

and t_1 is a divisor of t. Furthermore, $t_1(\mathbf{u}_1, v_1)$ is (regular) top-reducible by $(\mathbf{u}_2, v_2)$.

Proof. Since $t(\mathbf{u}_1, v_1)$ is top-reducible by $(\mathbf{u}_2, v_2)$ and both v_1 and v_2 are nonzero, there is a monomial s such that

$$t\,\mathrm{lm}(v_1) = s\,\mathrm{lm}(v_2), \quad t\,\mathrm{lm}(\mathbf{u}_1) \succeq s\,\mathrm{lm}(\mathbf{u}_2). \tag{5}$$

Let

$$t_2 = \frac{\mathrm{lcm}(\mathrm{lm}(v_1), \mathrm{lm}(v_2))}{\mathrm{lm}(v_2)} = \frac{\mathrm{lm}(v_1)}{\gcd(\mathrm{lm}(v_1), \mathrm{lm}(v_2))}.$$

Then (5) implies that, for some monomial w,

$$\begin{aligned} t &= \frac{\mathrm{lm}(v_2)}{\gcd(\mathrm{lm}(v_1), \mathrm{lm}(v_2))} w = t_1 w, \text{ and} \\ s &= \frac{\mathrm{lm}(v_1)}{\gcd(\mathrm{lm}(v_1), \mathrm{lm}(v_2))} w = t_2 w. \end{aligned}$$

Hence (5) implies that $t_2\mathrm{lm}(\mathbf{u}_2) \preceq t_1\mathrm{lm}(\mathbf{u}_1)$. So $\max(t_2\mathrm{lm}(\mathbf{u}_2), t_1\mathrm{lm}(\mathbf{u}_1)) = t_1\mathrm{lm}(\mathbf{u}_1)$, thus $t_1(\mathbf{u}_1, v_1)$ is a J-pair of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$. Note that by (5), we have that $t_1(\mathbf{u}_1, v_1)$ is regular top-reducible by $(\mathbf{u}_2, v_2)$ whenever $t(\mathbf{u}_1, v_1)$ is regular top-reducible by $(\mathbf{u}_2, v_2)$. □

Now let

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_k, v_k) \tag{6}$$

be a list of pairs in M as defined in (3). The list (6) is called a **strong Gröbner basis for** M if every pair $(\mathbf{u}, v) \in M$ is top-reducible by some pair in (6).

Proposition 1. *Suppose that the list of pairs in* (6) *is a strong Gröbner basis for* M*. Then*

(1) $\mathbf{G}_0 = \{\mathbf{u}_i : v_i = 0, 1 \le i \le k\}$ *is a Gröbner basis for the syzygy module of* $\mathbf{g} = (g_1, \ldots, g_m)$*, and*
(2) $G_1 = \{v_i : 1 \le i \le k\}$ *is a Gröbner basis for* $I = \langle g_1, \ldots, g_m \rangle$*.*

Proof. For any $\mathbf{u} = (u_1, \ldots, u_m)$ in the syzygy module of $\mathbf{g}$, we have $(\mathbf{u}, 0) \in M$. By our assumption, $(\mathbf{u}, 0)$ is top-reducible by some pair $(\mathbf{u}_i, v_i)$ in (6). Then we must have $v_i = 0$, thus $\mathbf{u}_i \in G_0$ and $\mathrm{lm}(\mathbf{u})$ is reducible by $\mathrm{lm}(\mathbf{u}_i)$. This proves that G_0 is a Gröbner basis for the syzygy module of $\mathbf{g}$.

Now suppose $v \in I$ and is nonzero. Then there exists $\mathbf{u} = (u_1, \ldots, u_m) \in R^m$ so that $\mathbf{u}\mathbf{g}^t = v$, hence $(\mathbf{u}, v) \in M$. Among all such $\mathbf{u}$, we pick one so that $\mathrm{lm}(\mathbf{u})$ is minimum. Since $(\mathbf{u}, v) \in M$, it is top-reducible by some $(\mathbf{u}_i, v_i)$ where $1 \le i \le k$. If $v_i = 0$, then we could use $(\mathbf{u}_i, 0)$ to reduce $(\mathbf{u}, v)$ to get a $\mathbf{u}'$ so that $\mathbf{u}'\mathbf{g}^t = v$ and $\mathrm{lm}(\mathbf{u}')$ is smaller than $\mathrm{lm}(\mathbf{u})$, contradicting to the minimality of $\mathrm{lm}(\mathbf{u})$. So $v_i \neq 0$ and $\mathrm{lm}(v_i)$ divides $\mathrm{lm}(v)$. Hence G_1 is a Gröbner basis for I. □

Remark. Note that $M \subset R^m \times R$ has a Gröbner basis in the usual sense as a submodule of R^{m+1} where the leading term of $(\mathbf{u}, v)$ is $\mathrm{lm}(v)\mathbf{E}_{m+1}$ if $v \neq 0$ and $\mathrm{lm}(\mathbf{u})$ if $v = 0$. The above proposition implies that a strong Gröbner basis for M is a Gröbner basis for M as a submodule of R^{m+1}, but the converse may not be true for an arbitrary submodule M of R^{m+1}. This is why we call our basis a strong Gröbner basis.

Let S be any set of pairs in $R^m \times R$. We say that a pair $(\mathbf{u}, v) \in R^m \times R$ is regular top-reducible by S if it is regular top-reducible by at least one pair in S. We call $(\mathbf{u}, v)$ **eventually super top-reducible** by S if there is a sequence of regular top-reductions of $(\mathbf{u}, v)$ by pairs in S that reduce $(\mathbf{u}, v)$ to a pair $(\mathbf{u}', v')$ that is no

longer regular top-reducible by S but is super top-reducible by at least one pair in S.

Theorem 1. *Suppose the list* (6) *satisfies the following: for any term $T \in R^m$, there is a pair $(\mathbf{u}_i, v_i)$, $1 \le i \le k$, and a monomial t such that $T = t\, lm(\mathbf{u}_i)$. Then* (6) *is a strong Gröbner basis for M if and only if every J-pair of the pairs from* (6) *is eventually super top-reducible by* (6).

Proof. The forward implication is immediate from the definition of a strong Gröbner basis. To show the reverse, we assume that every J-pair of the pairs in (6) is eventually super top-reducible by (6). Assume that there is a pair $(\mathbf{u}, v) \in M$ that is not top-reducible by any pair in (6). We want to get a contradiction. Among all such pairs $(\mathbf{u}, v)$ we pick one with minimal signature $T = \mathrm{lm}(\mathbf{u})$. Note that $T \neq \mathbf{0}$. Next, we select a pair $(\mathbf{u}_i, v_i)$ from (6) such that

(a) $T = t\,\mathrm{lm}(\mathbf{u}_i)$ for some monomial t, and
(b) $t\,\mathrm{lm}(v_i)$ is minimal among all $1 \le i \le k$ satisfying (a).

We claim that $t(\mathbf{u}_i, v_i)$ is not regular top-reducible by (6). To prove this claim, we suppose that $t(\mathbf{u}_i, v_i)$ is regular top-reducible by some $(\mathbf{u}_j, v_j)$, $j \neq i$, so both v_i and v_j are nonzero. We want to derive a contradiction to the condition (b). By Lemma 1, the J-pair of $(\mathbf{u}_i, v_i)$ and $(\mathbf{u}_j, v_j)$ is $t_1(\mathbf{u}_i, v_i)$ and that $t_1(\mathbf{u}_i, v_i)$ is still regular top-reducible by $(\mathbf{u}_j, v_j)$, where

$$t_1 = \frac{\mathrm{lcm}(\mathrm{lm}(v_i), \mathrm{lm}(v_j))}{\mathrm{lm}(v_i)}, \text{ and } t = t_1 w$$

for some monomial w. As $t_1(\mathbf{u}_i, v_i)$ is a J-pair of two pairs from (6), $t_1(\mathbf{u}_i, v_i)$ is eventually super top-reducible by (6), say

$$t_1(\mathbf{u}_i, v_i) = \sum_{r=1}^{d} m_r(\mathbf{u}_{i_r}, v_{i_r}) + (\mathbf{u}', v'), \tag{7}$$

where the first part of the sum represents a sequence of regular top-reductions of $t_1(\mathbf{u}_i, v_i)$ by (6), and $(\mathbf{u}', v')$ is not regular top-reducible by any pair in (6) but is super top-reducible by some pair in (6). Note that $d \ge 1$ as $t_1(\mathbf{u}_i, v_i)$ is regular top-reducible by $(\mathbf{u}_j, v_j)$. Also, each regular top-reduction strictly reduces the leading monomial of v_i, but the leading monomial of $\mathbf{u}_i$ remains unchanged. Thus we have $\mathrm{lm}(\mathbf{u}') = \mathrm{lm}(t_1\mathbf{u}_i)$ but $\mathrm{lm}(v') \prec t_1\mathrm{lm}(v_i)$. Let $1 \le \ell \le k$ be such that $(\mathbf{u}', v')$ is super top-reducible by $(\mathbf{u}_\ell, v_\ell)$. If $v_\ell = 0$, then $\mathrm{lm}(\mathbf{u}_\ell)$ divides $\mathrm{lm}(\mathbf{u}') = \mathrm{lm}(t_1\mathbf{u}_i)$, thus divides $\mathrm{lm}(\mathbf{u})$. Hence $(\mathbf{u}, v)$ is top-reducible by $(\mathbf{u}_\ell, 0)$, contradicting our assumption that $(\mathbf{u}, v)$ is not top-reducible by (6). We may thus assume that $v_\ell \neq 0$, hence $v' \neq 0$. Then

$$(\mathrm{lm}(\mathbf{u}'), \mathrm{lm}(v')) = t_3(\mathrm{lm}(\mathbf{u}_\ell), \mathrm{lm}(v_\ell)),$$

where $t_3 = \mathrm{lm}(v')/\mathrm{lm}(v_\ell)$. Let $\bar{t} = t_3 w$. Then

$$\bar{t}\,\mathrm{lm}(\mathbf{u}_\ell) = w\,\mathrm{lm}(\mathbf{u}') = t\,\mathrm{lm}(\mathbf{u}_i) = T$$

and

$$\bar{t}\,\mathrm{lm}(v_\ell) = w\,\mathrm{lm}(v') \prec wt_1\mathrm{lm}(v_i) = t\,\mathrm{lm}(v_i).$$

Thus $(\mathbf{u}_i, v_i)$ satisfies (a) but violates (b). Hence $t(\mathbf{u}_i, v_i)$ is not regular top-reducible by (6) as claimed.

Returning to the main proof, we perform the cancellation

$$(\overline{\mathbf{u}}, \overline{v}) = (\mathbf{u}, v) - ct(\mathbf{u}_i, v_i), \tag{8}$$

where $c = \mathrm{lc}(\mathbf{u})/\mathrm{lc}(\mathbf{u}_i)$ so that $\mathrm{lm}(\overline{\mathbf{u}}) \prec \mathrm{lm}(\mathbf{u}) = T$. Note that $\mathrm{lm}(v) \neq t\,\mathrm{lm}(v_i)$, since otherwise $(\mathbf{u}, v)$ would be top-reducible by $(\mathbf{u}_i, v_i)$ which contradicts the minimality assumption of $T = \mathrm{lm}(\mathbf{u})$. Hence $\overline{v} \neq 0$. Also, as $(\overline{\mathbf{u}}, \overline{v}) \in M$ and $\mathrm{lm}(\overline{\mathbf{u}}) \prec T$, we have that $(\overline{\mathbf{u}}, \overline{v})$ is top-reducible by (6). If $(\overline{\mathbf{u}}, \overline{v})$ is top-reducible by some pair $(\mathbf{u}_\ell, v_\ell)$ from (6) with $v_\ell = 0$, then we can reduce $(\overline{\mathbf{u}}, \overline{v})$ repeatedly by such pairs to get a new pair $(\tilde{\mathbf{u}}, \overline{v})$ that is not top-reducible by any pair in (6) with v-part being zero. Note that $(\tilde{\mathbf{u}}, \overline{v})$ is still in M and $\mathrm{lm}(\tilde{\mathbf{u}}) \prec T$. Hence $(\tilde{\mathbf{u}}, \overline{v})$ is top-reducible by some pair $(\mathbf{u}_\ell, v_\ell)$ from (6) with $v_\ell \neq 0$. As $\mathrm{lm}(v) \neq t\,\mathrm{lm}(v_i)$, we consider two cases:

(i) $\mathrm{lm}(v) \prec t\,\mathrm{lm}(v_i)$. Then $\mathrm{lm}(\overline{v}) = t\,\mathrm{lm}(v_i)$, hence $t(\mathbf{u}_i, v_i)$ is regular top-reducible by $(\mathbf{u}_\ell, v_\ell)$ (as $\mathrm{lm}(\tilde{\mathbf{u}}) \prec t\,\mathrm{lm}(\mathbf{u}_i)$). Since $t(\mathbf{u}_i, v_i)$ is not regular top-reducible by any pair in (6), this case is impossible.

(ii) $\mathrm{lm}(v) \succ t\,\mathrm{lm}(v_i)$. Then $\mathrm{lm}(\overline{v}) = \mathrm{lm}(v)$, and $(\mathbf{u}, v)$ is regular top-reducible by $(\mathbf{u}_\ell, v_\ell)$, contradicting the fact that $(\mathbf{u}, v)$ is not top-reducible by any pair in (6).

Therefore such a pair $(\mathbf{u}, v)$ does not exist in M, thus every pair in M is top-reducible by (6). □

Theorem 2. *In Theorem 1, the condition "every J-pair of pairs from* (6) *is eventually super top-reducible by* (6)*" can be replaced by "for every distinct J-signature from* (6) *there is at least one J-pair from* (6) *with the same J-signature that is eventually super top-reducible by* (6)*."*

Proof. The proof is the same as that of Theorem 1, except that in the equation (7) we can still assume that the pair $(\mathbf{u}', v')$ is not regular top-reducible by (6) but we need to prove that it is super top-reducible by (6), which is used in the subsequent proof. By our assumption, however, there is a J-pair, say $t_2(\mathbf{u}_\ell, v_\ell)$, that has the same signature as that of $t_1(\mathbf{u}_i, v_i)$ and is eventually super top-reducible by (6). Suppose that

$$t_2(\mathbf{u}_\ell, v_\ell) = \sum_{r=1}^{s} n_r(\mathbf{u}_{\ell_r}, v_{\ell_r}) + (\mathbf{u}'', v''),$$

where the first part of the sum represents a sequence of regular top-reductions of $t_2(\mathbf{u}_\ell, v_\ell)$ by (6), and $(\mathbf{u}'', v'')$ is not regular top-reducible by any pair in (6) but is super top-reducible by some pair $(\mathbf{u}_e, v_e)$ in (6). Note that

$$\mathrm{lm}(\mathbf{u}'') = t_2\mathrm{lm}(\mathbf{u}_\ell) = t_1\mathrm{lm}(\mathbf{u}_i) = \mathrm{lm}(\mathbf{u}') = T,$$

and $\mathrm{lm}(v'') \preceq t_2\mathrm{lm}(v_\ell)$. We may assume that both $\mathbf{u}'$ and $\mathbf{u}''$ are monic. We claim that $\mathrm{lm}(v') = \mathrm{lm}(v'')$ and their coefficients are also equal. This implies the desired property that $(\mathbf{u}', v')$ is super top-reducible by $(\mathbf{u}_e, v_e)$, since $\mathrm{lm}(v') = \mathrm{lm}(v'')$, $\mathrm{lm}(\mathbf{u}') = \mathrm{lm}(\mathbf{u}'')$ and $(\mathbf{u}'', v'')$ is super top-reducible by $(\mathbf{u}_e, v_e)$.

To prove the claim, suppose it is not true, that is, either $\mathrm{lm}(v') \neq \mathrm{lm}(v'')$ or $\mathrm{lm}(v') = \mathrm{lm}(v'')$ but their coefficients are not equal. Then $\mathrm{lm}(v' - v'') = \mathrm{lm}(v')$ or $\mathrm{lm}(v'')$. Note that the signature of $(\mathbf{u}' - \mathbf{u}'', v' - v'')$ is strictly smaller than $\mathrm{lm}(\mathbf{u}') = \mathrm{lm}(t_1\mathbf{u}_i) \preceq T$. By the hypothesis on the minimality of T, the pair $(\mathbf{u}' - \mathbf{u}'', v' - v'')$ is top-reducible by (6). It follows that either $(\mathbf{u}', v')$ is regular top-reducible by (6) if $\mathrm{lm}(v' - v'') = \mathrm{lm}(v')$ or $(\mathbf{u}'', v'')$ is regular top-reducible by (6) if $\mathrm{lm}(v' - v'') = \mathrm{lm}(v'')$. Both are contradicting to our assumption that they are not regular top-reducible. Hence the claim, and thus the theorem is proved. □

Remark. Suppose a final strong Gröbner basis for M is

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_k, v_k). \tag{9}$$

At any intermediate step of computation, we only know

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_p, v_p) \tag{10}$$

for some $p < k$. In general, a pair $(\mathbf{u}, v)$ may be eventually super top-reducible by (10) but not by (9). How can one decide whether $(\mathbf{u}, v)$ is eventually super top-reducible by (9) when only (10) is known? Our strategy is to always pick the J-pair with minimal signature to reduce. Then a pair that is eventually super top-reducible by an intermediate basis is always eventually super top-reducible by the final basis. A more detailed argument will be given in the next section.

3. Algorithm, Term Orderings and Time Comparison

Algorithm and Its Correctness. Our algorithm is based on Theorems 1 and 2. The basic idea of our algorithm is as follows. Initially, we have the pairs in (4) in our Gröbner basis. So the first condition of the theorem is satisfied. From these pairs, we form all J-pairs, keeping only one J-pair for each J-signature. We then take the smallest J-pair among them and repeatedly perform regular top-reductions until it is no longer regular top-reducible, say to get $(\mathbf{u}, v)$. If the v part of the resulting pair is zero, then the $\mathbf{u}$ part is a syzygy in $\mathbf{H}$, and we store this vector. If the v part is nonzero, then we check if $(\mathbf{u}, v)$ is super top-reducible. If so, then we discard this J-pair; otherwise, we add this $(\mathbf{u}, v)$ pair to the current Gröbner basis, and form new J-pairs. Repeat this process until all J-pairs are eventually super top-reducible.

We make two improvements on this basic algorithm. First, storing and updating syzygies $\mathbf{u} \in \mathbf{H}$ are expensive. In our computation, we shall make all pairs $(\mathbf{u}, v)$ monic, namely, the leading coefficient of $\mathbf{u}$ is 1. Now suppose $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$ are any two monic pairs. Then a top-reduction (regular or super) is determined only by $\mathrm{lm}(\mathbf{u}_1)$, $\mathrm{lm}(\mathbf{u}_2)$, v_1 and v_2. The other terms of $\mathbf{u}_1$ and $\mathbf{u}_2$ are not used at all. Let $T_1 = \mathrm{lm}(\mathbf{u}_1)$ and $T_2 = \mathrm{lm}(\mathbf{u}_2)$, the signatures of $(\mathbf{u}_1, v_1)$ and $(\mathbf{u}_2, v_2)$, respectively. Suppose we store only (T_1, v_1) and (T_2, v_2). Then (T_1, v_1) is regular top-reducible by (T_2, v_2) when $v_2 \neq 0$, $\mathrm{lm}(v_1)$ is divisible by $\mathrm{lm}(v_2)$, $tT_2 \preceq T_1$, and $\mathrm{lc}(v_1) \neq \mathrm{lc}(v_2)$ if $tT_2 = T_1$. The corresponding top-reduction is

$$v := v_1 - ctv_2$$

where $t = \mathrm{lm}(v_1)/\mathrm{lm}(v_2)$ and $c = \mathrm{lc}(v_1)/\mathrm{lc}(v_2)$, and furthermore, if $tT_2 = T_1$ then we update v as

$$v := v/(1 - c).$$

Then (T_1, v) is the resulting pair of the reduction, and it replaces (T_1, v_1). Our algorithm below will performs regular top-reductions in this fashion.

Another improvement is to use trivial syzygies. We will store the leading terms of known syzygies in a list called H. Let (T_1, v_1) and (T_2, v_2) be any two pairs from the Gröbner basis computed so far, where v_1 and v_2 are both nonzero. Then, for $1 \leq i \leq 2$, there are $\mathbf{u}_i \in R^m$ such that $\mathrm{lm}(\mathbf{u}_i) = T_i$ and $(\mathbf{u}_i, v_i) \in M$. Then we have

$$v_2(\mathbf{u}_1, v_1) - v_1(\mathbf{u}_2, v_2) = (v_2\mathbf{u}_1 - v_1\mathbf{u}_2, 0) \in M.$$

Hence $v_2\mathbf{u}_1 - v_1\mathbf{u}_2$ is a syzygy of $(g_1, \ldots, g_m)$. Its leading term is

$$T = \max(T_1\mathrm{lm}(v_2), T_2\mathrm{lm}(v_1)),$$

provided that $T_1\mathrm{lm}(v_2) \neq T_2\mathrm{lm}(v_1)$, or $T_1\mathrm{lm}(v_2) = T_2\mathrm{lm}(v_1)$ but $\mathrm{lc}(v_1) \neq \mathrm{lc}(v_2)$. When $T_1\mathrm{lm}(v_2) = T_2\mathrm{lm}(v_1)$ and $\mathrm{lc}(v_1) = \mathrm{lc}(v_2)$, the leading terms in $v_2(\mathbf{u}_1, v_1)$ and $v_1(\mathbf{u}_2, v_2)$ cancel each other. In that case, we don't know the leading term of the syzygy, so we just ignore such a syzygy. In all other cases, our algorithm will add T to the list H. The benefit of H is in detecting useless reductions. That is, whenever a J-pair has a signature that is divisible by a term in H, it is always eventually super top-reducible and hence discarded, thus saving time.

The algorithm is described more precisely in Figure 1 below. As mentioned above, we use H to record leading terms of syzygies. In addition to H, our algorithm uses two more lists to store the pairs $(T_1, v_1), (T_2, v_2), \ldots, (T_k, v_k)$ with $v_i \neq 0$ for $1 \leq i \leq k$. This list will be stored as

$$U = [T_1, T_2, \ldots, T_k], \quad V = [v_1, v_2, \ldots, v_k].$$

Then $[U, V]$ represents the whole list $(T_1, v_1), (T_2, v_2), \ldots, (T_k, v_k)$.

Theorem 3. *If the algorithm in Figure 1 terminates, then V is a Gröbner basis for $I = \langle g_1, g_2, \ldots, g_m\rangle$ and H is a Gröbner basis for the leading terms of the syzygy module of $(g_1, g_2, \ldots, g_m)$.*

Proof. To prove the correctness of the algorithm, we need to show the following:

(i) One can delete J-pairs in Steps 0, 3a, and 3b whose signatures are divisible by $\mathrm{lm}(\mathbf{u})$, where $\mathbf{u} \in H$.
(ii) A pair that is eventually super top-reducible by an intermediate basis is always eventually super top-reducible by the final basis.
(iii) One just needs to keep one J-pair for each signature, which follows directly from Theorem 2.

Our current basis consists of pairs in (U, V) and $(H, 0)$. For (i), let $(\mathbf{u}, v)$ be any pair whose signature $\mathrm{lm}(\mathbf{u})$ is divisible by $\mathrm{lm}(\mathbf{u}')$ for some $\mathbf{u}' \in H$. Then $(\mathbf{u}, v)$ is top-reducible by $(\mathbf{u}', 0)$. Any regular top-reduction of $(\mathbf{u}, v)$ won't change $\mathrm{lm}(\mathbf{u})$, so the pair obtained from $(\mathbf{u}, v)$ by any sequence of regular top-reductions will be super top-reducible by $(\mathbf{u}', 0)$. Hence $(\mathbf{u}, v)$ is eventual super top-reducible by the current basis. This means that we don't need to reduce $(\mathbf{u}, v)$, and so we simply discard it.

To see (ii), suppose the final Gröbner basis computed for M is

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_k, v_k), \tag{11}$$

while at any intermediate step, we only know

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_p, v_p) \tag{12}$$

for some $p < k$. Suppose that the smallest J-pair from JP is (t, i) (i.e., $t(u_i, v_i)$). If $t(u_i, v_i)$ is eventually super top-reducible by (12), then $t(u_i, v_i)$ remains eventually super top-reducible by (11), as all $(\mathbf{u}_j, v_j)$, $j > p$, have strictly larger signature than $t(u_i, v_i)$. If $t(u_i, v_i)$ is not eventually super top-reducible by (12), then the basis (12) is augmented by a new pair $(\mathbf{u}_{p+1}, v_{p+1})$, which is obtained from $t(u_i, v_i)$ via regular top-reductions by (12). Hence the J-pair $t(u_i, v_i)$ is eventually super top-reducible by the new basis

$$(\mathbf{u}_1, v_1), \ldots, (\mathbf{u}_p, v_p), (\mathbf{u}_{p+1}, v_{p+1}). \tag{13}$$

	Algorithm for computing Gröbner bases
Input:	$g_1, \ldots, g_m \in R = \mathbb{F}[x_1, \ldots, x_n]$, a term order for R, and a term order on R^m
Output:	A Gröbner basis for $I = \langle g_1, \ldots, g_m \rangle$, and a Gröbner basis for $\mathrm{lm}(\mathbf{H})$, the leading terms of the syzygy module
Variables:	U a list of terms $T_i = \mathrm{lm}(\mathbf{u}_i)$, representing signatures for $(\mathbf{u}_i, v_i) \in M$, V a list of polynomials for v_i for $(\mathbf{u}_i, v_i) \in M$, H a list for $\mathrm{lm}(\mathbf{u})$ were $\mathbf{u} \in R^m$ is a syzygy found so far, JP a list of pairs (t, i), where t is a monomial so that $t(\mathbf{u}_i, v_i)$ is the J-pair of $(\mathbf{u}_i, v_i)$ and $(\mathbf{u}_j, v_j)$ for some $j \neq i$. We shall refer (t, i) as a J-pair of $(\mathbf{u}_i, v_i)$ and $(\mathbf{u}_j, v_j)$.
Step 0.	$U = [\mathbf{E}_1, \ldots, \mathbf{E}_m]$, and $V = [g_1, \ldots, g_m]$. Find the leading terms of the principle syzygies $g_j\mathbf{E}_i - g_i\mathbf{E}_j$ for $1 \leq i < j \leq m$, and add them in H. Compute all the J-pairs of $(\mathbf{E}_1, g_1), \ldots, (\mathbf{E}_m, g_m)$ storing into JP all such J-pairs whose signatures are not reducible by H (storing only one J-pair for each distinct signature).
Step 1.	Take a minimal (in signature) pair (t, i) from JP, and delete it from JP.
Step 2.	Reduce the pair $t(T_i, v_i)$ repeatedly by the pairs in (U, V), using regular top-reductions until it is not regular top-reducible, say to get (T, v).
Step 3a.	If $v = 0$, then append T to H, and delete every J-pair (t, j) in JP whose signature tT_j is divisible by T.
Step 3b.	If $v \neq 0$ and (T, v) is not super top-reducible by (U, V), then *i)* Append T to U and v to V, *ii)* Form new J-pairs of (T, v) and (T_j, v_j), $1 \leq j \leq \lvert U\rvert - 1$, and *iii)* Insert into JP all such J-pairs whose signatures are not reducible by H (storing only one J-pair for each distinct signature). *iv)* Add the leading terms of the principle syzygies, $vT_j - v_jT$ for $1 \leq$ $j \leq \lvert U\rvert - 1$, to H.
Step 4.	While JP is not empty, go to step 1.
Return:	V and H.

FIGURE 1

Note that $(\mathbf{u}_{p+1}, v_{p+1})$ has the same signature as the J-pair $t(\mathbf{u}_i, v_i)$. All new J-pairs formed using $(\mathbf{u}_{p+1}, v_{p+1})$ will have strictly greater signature than that of $(\mathbf{u}_{p+1}, v_{p+1})$ (we never keep any future J-pair that has the same signature as $(\mathbf{u}_{p+1}, v_{p+1})$). Hence $(\mathbf{u}_{p+1}, v_{p+1})$ can not be top-reducible by any pair $(\mathbf{u}_j, v_j)$, $j > p + 1$, so the J-pair $t(u_i, v_i)$ remains eventually super top-reducible by (11). Therefore, any pair that is eventually super top-reducible by our current basis remains so by the final basis. □

Remarks on finite termination and Gröbner bases for the syzygy module. Presently, we have been unable to show that the algorithm actually terminates in finitely many steps. We believe that it has finite termination, but leave its proof as an open problem. We also mention that we believe that the proof of finite termination of F5 given in [11] is not correct. The proof of finite termination seems to be still open.

We next explain that once we have U, V and H from the above algorithm, we can compute a Gröbner basis for the syzygy module as follows. For the first m terms of U and V (i.e., $(\mathbf{E}_i, g_i),\ 1 \leq i \leq m$) are already in M. Other terms are of the form (T_i, v_i) where $T_i = \mathrm{lm}(\mathbf{u}_i)$ for some $\mathbf{u}_i \in R^m$ such that $(\mathbf{u}_i, v_i) \in M$. We need to find these $\mathbf{u}_i$. Suppose that (U, V) is sorted so that the terms T_i are in increasing order (as done by the above algorithm). Suppose (T_i, v_i) is the smallest pair in (U, V) that is not in M. Find a pair $(\mathbf{u}_j, v_j)$ from (U, V) that is already in M (i.e. $j < i$) and a monomial t so that $T = t\,\mathrm{lm}(\mathbf{u}_j)$ (take j so that t is minimum), perform regular top-reductions of $t(\mathbf{u}_j, v_j)$ by (U, V) using only those pairs that are in M (both parts of the pair are updated as in Section 2) until it's not regular top-reducible, say to get a pair $(\mathbf{u}, v)$. Then $(\mathbf{u}, v)$ must be super top-reducible by (T_i, v_i), hence we replace (T_i, v_i) by $(\mathbf{u}, v)$ in (U, V), and continue with the next smallest (T_j, v_j) $(j > i)$ if any that is not in M. After all T_i's in U are processed, (U, V) represents a list $(\mathbf{u}_1, v_1), (\mathbf{u}_2, v_2), \ldots, (\mathbf{u}_m, v_m)$ so that $\mathrm{lm}(\mathbf{u}_i) = T_i$ and $(\mathbf{u}_i, v_i) \in M$ for all i. To get a Gröbner basis for the syzygy module, just do the following. For each term T in H, first find a minimum monomial t so that there is a pair $(\mathbf{u}_i, v_i)$ in (U, V) so that $T = t\,\mathrm{lm}(\mathbf{u}_i)$, next perform regular top-reductions of $t(\mathbf{u}_i, v_i)$ by (U, V) until the v-part is zero, then the $\mathbf{u}$-part is a syzygy with leading term equal to T. All these syzygies form a Gröbner basis for the syzygy module. This gives an algorithm for computing a Gröbner basis for syzygy module under any term order.

Term Orders. Now we discuss choices of term orders. We use $\prec_1$ to represent a term ordering on R and $\prec_2$ to represent a term ordering on R^m. While computing Gröbner bases for both $\langle g_1, \ldots, g_m \rangle$ and $\mathbf{H}$, one should set $\prec_1$ and $\prec_2$ to the appropriate term orderings for the Gröbner bases desired. Often, however, the Gröbner basis for $\mathbf{H}$ is not needed. Then we only need the leading terms of $\mathbf{H}$ to speed up the computation of $\langle g_1, \ldots, g_m \rangle$. In this case, we have tremendous freedom in the choice of $\prec_2$.

There are many ways that we can construct a term ordering on R^m. We consider four extreme cases below. Let $\prec$ be some term order on R. We extend $\prec$ to R^m as follows.

(POT) The first is called position over term ordering (POT). We say that $x^\alpha \mathbf{E}_i \prec x^\beta \mathbf{E}_j$ if $i < j$ or $i = j$ and $x^\alpha \prec x^\beta$.

(TOP) The second is the term over position ordering (TOP). We say that $x^\alpha \mathbf{E}_i \prec x^\beta \mathbf{E}_j$ if $x^\alpha \prec x^\beta$ or $x^\alpha = x^\beta$ and $i < j$.

(**g**1) Next is the **g**-weighted degree followed by TOP. We say that $x^\alpha \mathbf{E}_i \prec x^\beta \mathbf{E}_j$ if $\deg(x^\alpha g_i) < \deg(x^\beta g_j)$ or $\deg(x^\alpha g_i) = \deg(x^\beta g_j)$ and $x^\alpha \mathbf{E}_i \prec_{top} x^\beta \mathbf{E}_j$ where deg is for total degree.

(**g**2) Finally, we have **g**-weighted $\prec$ followed by POT. We say that $x^\alpha \mathbf{E}_i \prec x^\beta \mathbf{E}_j$ if $\mathrm{lm}(x^\alpha g_i) \prec \mathrm{lm}(x^\beta g_j)$ or $\mathrm{lm}(x^\alpha g_i) = \mathrm{lm}(x^\beta g_j)$ and $x^\alpha \mathbf{E}_i \prec_{pot} x^\beta \mathbf{E}_j$.

We remark that, under the POT order, our new algorithm corresponds with the G2V algorithm presented in [10]. The reason is that this new algorithm always first picks J-pairs with signatures containing $\mathbf{E}_1$, then those with $\mathbf{E}_2$, etc. This means that it computes Gröbner bases for $\langle g_1 \rangle$, $\langle g_1, g_2 \rangle$, ..., $\langle g_1, g_2, \ldots, g_m \rangle$, just like G2V and F5. The only difference is that the intermediate bases may not be reduced and non-leading terms are not reduced as in the computing of normal forms. Because of this fact and other implementation choices, the running times under POT reported here are much slower than those in [10].

Test Case (#generators)	F5	F5C	G2V
Katsura5 (22)	1.48	0.93	0.36
Katsura6 (41)	2.79	2.34	0.37
Katsura7 (74)	30.27	22.76	4.64
Katsura8 (143)	290.97	177.74	29.88
Schrans-Troost (128)	1180.08	299.65	21.34
F633 (76)	30.93	29.87	2.06
Cyclic6 (99)	28.44	22.06	5.65
Cyclic7 (443)	4591.20	2284.05	732.33

TABLE 1. Run-times in seconds comparing F5, F5C and G2V (GVW under POT ordering) for various test cases in Singular 3110 on an Intel Core 2 Quad 2.66 GHz. This table is reproduced from [10].

Another remark is that our algorithm under the **g**1 order corresponds with an improved version of the XL algorithm [4]. In the XL algorithm, one performs row reductions on a matrix whose rows correspond to all polynomials $x^\alpha g_i$, $1 \leq i \leq m$, with total degree of $x^\alpha g_i$ smaller than some bound. Our algorithm basically works with only some of those rows that correspond to J-signatures. So our algorithm needs much less storage.

Comparison. For ease of exposition, we refer to our algorithm as GVW. We implemented GVW in Singular CAS and C++ so that $\langle g_1, \ldots, g_m \rangle$ is computed in one-shot, that is, non-incrementally. The Singular implementation is very similar to the that provided in [10] except that it no longer uses Singular's "reduce" function. Without this use of Singular's kernel, GVW and G2V are not very comparable in terms of runtimes. For the exact same reason, we did not compare GVW to F5 or F5C as we did in [10] (see Table 1, reproduced here for comparison purpose). But as mentioned earlier, GVW under POT is the G2V algorithm.

Just as in [10], various benchmark examples (from [6]) were run for comparison. We collected data from each example under each term ordering for comparison. Tables 2 and 3 list the runtimes in seconds of GVW for each of the four term orderings. One might notice that the Singular runtimes are surprisingly large (especially compared to G2V in [10]), but that is most likely the result of relying on Singular's kernel routines less. In examining the timings, we find that **g**2 seems to be a clear winner among the four term orders.

A more computer independent measure would be a count of J-pairs processed and the number of extraneous generators produced. Table 4 lists the total number of J-pairs processed for each term ordering. It's analogous to counting the number of S-pairs processed in F5 or Buchberger's algorithm. As with the timings, **g**2 seems to be the most efficient. Finally, Table 5 lists the size of the Gröbner bases produced by GVW with each term ordering. These are the Gröbner bases produced by the algorithm before any interreduction occurs to produce a reduced Gröbner basis. We believe this measure to be significant since fewer extraneous generators means quicker reductions. Again, we see that **g**2 produces less redundancy than the other orderings.

Test Case (# gen)	POT/G2V	TOP	**g**1	**g**2
Katsura5 (22)	1.12	1.01	1.17	0.70
Katsura6 (41)	1.60	3.25	3.76	1.92
Katsura7 (74)	24.03	18.00	19.94	9.22
Katsura8 (143)	167.40	107.97	115.89	52.45
Schrans-Troost (128)	80.08	62.19	66.34	66.26
F633 (76)	10.57	41.90	38.43	11.13
Cyclic 6 (99)	27.09	1043.36	1129.20	20.63
Cyclic 7 (443)	4194.24	-	-	1835.63

TABLE 2. Runtime in seconds using Singular 3110 on an Intel Core 2 Quad 2.66 GHz processor

Test Case (# gen)	POT/G2V	TOP	**g**1	**g**2
Katsura5 (22)	0.00	0.01	0.01	0.01
Katsura6 (41)	0.02	0.04	0.04	0.04
Katsura7 (74)	0.34	0.37	0.36	0.37
Katsura8 (143)	3.26	2.92	2.97	3.16
Schrans-Troost (128)	1.78	3.65	3.64	3.81
F633 (76)	0.08	0.44	0.36	0.09
Cyclic 6 (99)	0.34	3.30	3.24	0.15
Cyclic 7 (443)	139.56	21417.40	20800.60	35.75
Cyclic 8 (1182)	107684.35	-	-	5737.41

TABLE 3. Runtime in seconds using our C++ implementation on an Intel Core 2 Quad 2.66 GHz processor

Test Case (# gen)	POT/G2V	TOP	**g**1	**g**2
Katsura5 (22)	64	61	61	36
Katsura6 (41)	72	90	90	50
Katsura7 (74)	216	181	181	93
Katsura8 (143)	439	359	359	182
Schrans-Troost (128)	475	204	204	214
F633 (76)	313	378	346	276
Cyclic 6 (99)	441	4388	4388	368
Cyclic 7 (443)	3562	69502	69502	2375
Cyclic 8 (1182)	37757	-	-	12245

TABLE 4. A count of the J-pairs processed

One might make the observation that in [10] (or Table 1), G2V outperformed F5 and F5C by runtimes of 2 to 10 times, while with the present algorithm, GVW under the **g**2 ordering outperforms G2V (GVW under the POT ordering) by another factor of 2 to 10 times. This comparison shows that if GVW under **g**2 were implemented comparably to F5 or F5C, it would compute Gröbner bases around 4 to 20 times faster.

Test Case (# gen)	POT/G2V	TOP	**g**1	**g**2
Katsura5 (22)	67	64	64	27
Katsura6 (41)	75	91	91	46
Katsura7 (74)	224	175	175	80
Katsura8 (143)	448	343	343	151
Schrans-Troost (128)	402	137	137	136
F633 (76)	138	185	171	109
Cyclic 6 (99)	160	1189	1189	193
Cyclic 7 (443)	755	9237	9237	852
Cyclic 8 (1182)	3872	-	-	3647

TABLE 5. Size of GB before any interreduction

Test Case (# gen)	POT/G2V	TOP	**g**1	**g**2
Katsura5 (22)	1.26	1.26	1.26	0.75
Katsura6 (41)	1.26	1.76	1.76	1.26
Katsura7 (74)	6.28	5.77	5.77	2.76
Katsura8 (143)	25.83	22.81	22.84	8.78
Schrans-Troost (128)	39.82	6.28	6.28	6.28
F633 (76)	2.28	4.28	3.28	1.78
Cyclic 6 (99)	1.28	13.50	13.56	1.79
Cyclic 7 (443)	22.24	-	-	26.62

TABLE 6. Maximal amount of memory used (MiB) by Singular

4. Algorithm for Quotient Rings and Modules

Quotient Rings. Let $\mathbb{F}$ be any field and $R = \mathbb{F}[x_1, \ldots, x_n]$ be a polynomial ring. Let J be an ideal of R with Gröbner basis $G = \{f_1, \ldots, f_k\}$. Suppose I is an ideal of R/J generated by $\{g_1, \ldots, g_m\}$ where each g_i is already in normal form with respect to G. We wish to compute a Gröbner basis for $I = \langle g_1, \ldots, g_m \rangle$, and the $(g_1, \ldots, g_m)$-syzygy module.

We represent polynomials in R/J in normal form modulo G. This means that, for any $g \in R$, we have

$$g \equiv \sum_{i=1}^{\ell} c_i x^{\alpha_i} \pmod{G} \tag{14}$$

where no term x^{α_i} is divisible by any leading term of G. This expression can be obtained from g by long division via G. When $g \in R$ is viewed as a polynomial in R/J, the leading term of g is the maximal x^{α_i} that appears in the normal form (14) of g. So the leading term of g is never divisible by any leading term of G.

We begin by defining a Gröbner basis for an ideal $I \subset R/J$. We say that a generating set $\{g_1, \ldots, g_m\} \subset R/J$ is a Gröbner basis for $I = \langle g_1, \ldots, g_m \rangle$ if for any $h \in I$, the leading monomial of h is divisible by the leading monomial of one of the generators, that is

$$\mathrm{lm}(g_i) \mid \mathrm{lm}\,(h) \quad \text{for some } 1 \le i \le m.$$

In other words, if $\{g_1, \ldots, g_m\}$ is a Gröbner basis for $I \subset R/J$ and $\{f_1, \ldots, f_k\}$ is a Gröbner basis for $J \subset R$, then $\{g_1, \ldots, g_m, f_1, \ldots, f_k\}$ is a Gröbner basis for $\langle g_1, \ldots, g_m, f_1, \ldots, f_k\rangle \subset R$.

The syzygy module for $\mathbf{g} = (g_1, \ldots, g_m) \in (R/J)^m$ is defined as

$$\mathbf{H} = \{(u_1, \ldots, u_m) \in (R/J)^m : u_1 g_1 + \cdots + u_m g_m = 0 \text{ in } R/J\}.$$

If viewed in the original ring R, every $(g_1, \ldots, g_m)$-syzygy in $(R/J)^m$ can be extended to an $(g_1, \ldots, g_m, f_1, \ldots, f_k)$-syzygy in R^{m+k}, which may vary depending on how $u_1 g_1 + \cdots + u_m g_m$ is reduced to 0 by G. In our computation, we only need to store the leading term of $(u_1, \ldots, u_m) \in \mathbf{H}$ where no terms in the u_i's are divisible by the leading term of G.

Figure 2 describes a slight modification to the GVW algorithm that produces a Gröbner basis for $\langle g_1, \ldots, g_m\rangle \subset R/J$ and a Gröbner basis for the leading terms of the syzygy module $\mathbf{H}$, which can be used to calculate an actual Gröbner basis for $\mathbf{H}$. Figure 2 needs little explanation beyond Figure 1. By saying a J-signature $x^\alpha \mathbf{E}_i$ "is not reducible by G or H", we mean the following. Not being reducible by G means that for any $x^\beta \in \mathrm{lm}(G)$, we require that $x^\beta \nmid x^\alpha$. While not being reducible by H means for any $x^\beta \mathbf{E}_j \in H$, then either $i \neq j$ or $x^\beta \nmid x^\alpha$.

This version of GVW can be used to compute Gröbner incrementally, each time adding m polynomials. For example, to compute a Gröbner basis for an ideal $I = \langle g_1, \ldots, g_t\rangle \subset R$, one can first compute a Gröbner basis G for $J = \langle g_1, \ldots, g_k\rangle \subset R$ where $k < t$. Then compute a Gröbner basis G_1 for $\langle g_{k+1}, \ldots, g_t\rangle$ in the quotient ring R/J. Then $G \cup G_1$ is a Gröbner basis for I. And in the process, G is used in the reduction of many polynomials (e.g., the v part of every J-pair). By interpreting any polynomial in R/J as having already been reduced to normal form modulo G, we keep the number of terms in each polynomial to a minimum, thus reducing computational and storage requirements. Also, as the choice for k and m are arbitrary, one can design an algorithm that can compute Gröbner bases in one-shot, incrementally, or some hybrid of the two. This provides a flexible strategy for computing Gröbner bases for large systems of polynomials.

Modules. Let $\mathbb{F}$ be a field and $R = \mathbb{F}[x_1, \ldots, x_n]$ be a polynomial ring. Let $\mathbf{g}_1, \ldots, \mathbf{g}_m$ be elements in R^s. We define an R-linear operator $T : R^m \to R^s$, uniquely determined by $\mathbf{g}_1, \ldots, \mathbf{g}_m$, given by

$$(f_1, \ldots, f_m) \longmapsto (f_1, \ldots, f_m) \begin{bmatrix} \mathbf{g}_1 \\ \mathbf{g}_2 \\ \vdots \\ \mathbf{g}_m \end{bmatrix}.$$

We wish to determine the image space and kernel of T. Note that the image is the R-submodule $\mathbf{I}$ generated by $\{\mathbf{g}_1, \ldots, \mathbf{g}_m\}$ in R^s while the kernel of T corresponds to the $(\mathbf{g}_1, \ldots, \mathbf{g}_m)$-syzygy module $\mathbf{H}$ in R^m.

We fix term orders $\prec_1$ on R^s and $\prec_2$ on R^m, and let $\mathbf{u} = (f_1, \ldots, f_m) \in R^m$ and $\mathbf{v} = T(\mathbf{u}) \in R^s$. We redefine M as an R-submodule of $R^m \times R^s$ so that

$$M = \{(\mathbf{u}, \mathbf{v}) \in R^m \times R^s : T(\mathbf{u}) = \mathbf{v}\}.$$

We continue to use $\mathbf{E}_i$, $1 \leq i \leq m$ as the i^{th} unit vector in R^m, but to avoid confusion we use $\mathbf{F}_j$, $1 \leq j \leq s$ as the j^{th} unit vector in R^s. And now, the

	Algorithm for computing Gröbner bases in quotient rings
Input:	$G = [f_1, \ldots, f_k]$, a Gröbner basis for an ideal $J \subset R$, $g_1, \ldots, g_m$ polynomials in R in normal form modulo G, and term orders for R and R^m.
Output:	A Gröbner basis for $\langle g_1, \ldots, g_m \rangle \in R/J$ and a Gröbner basis for $\mathrm{lm}(\mathbf{H})$, the leading terms of the syzygy module.
Variables:	U a list of terms $T_i = \mathrm{lm}(\mathbf{u}_i)$, representing signatures for $(\mathbf{u}_i, v_i) \in M$. V a list of polynomials for v_i for $(\mathbf{u}_i, v_i) \in M$; H a list for $\mathrm{lm}(\mathbf{u})$ where $\mathbf{u} \in R^m$ is a syzygy found so far, JP a list of pairs (t, i), where t is a monomial so that $t(\mathbf{u}_i, v_i)$ is a J-pair of $(\mathbf{u}_i, v_i)$ and $(\mathbf{u}_j, v_j)$ for some $j \neq i$. We shall refer (t, i) as the J-pair of $(\mathbf{u}_i, v_i)$ and $(\mathbf{u}_j, v_j)$.
Step 0.	$U = [\mathbf{0}, \ldots, \mathbf{0}]$ with length k, and $V = [f_1, \ldots, f_k]$ (so that $(u_i, v_i) = (\mathbf{0}, f_i)$, $1 \leq i \leq k$); $JP = [\,]$ and $H = [\,]$, empty lists; Add $\mathbf{E}_i$ to U and g_i to V for $1 \leq i \leq m$ (so that $(u_{k+i}, v_{k+i}) = (\mathbf{E}_i, g_i)$); Find the leading terms of the principle syzygies $g_j\mathbf{E}_i - g_i\mathbf{E}_j$ for $1 \leq i < j \leq m$, and add them to H; For each $1 \leq i \leq m$, and $1 \leq j < k + i$ compute the J-pair of the two pairs $(\mathbf{u}_{k+i}, v_{k+i}) = (\mathbf{E}_i, g_i)$ and $(\mathbf{u}_j, v_j)$, inserting it into JP whenever the J-signature is not reducible reducible by G or H (storing only one J-pair for each distinct J-signature).
Step 1.	Take a minimal (in signature) pair (t, i) from JP, and delete it from JP.
Step 2.	Reduce the pair $t(T_i, v_i)$ repeatedly by the pairs in (U, V), using regular top-reductions, until it is not regular top-reducible, say to get (T, v)
Step 3a.	If $v = 0$, then append T to H and delete every J-pair (t, j) in JP whose signature tT_j is divisible by T.
Step 3b.	If $v \neq 0$ and (T, v) is not super top-reducible by (U, V), then *i)* append T to U and v to V, *ii)* form new J-pairs of (T, v) and (T_j, v_j), $1 \leq j \leq \lvert U\rvert - 1$, and *iii)* insert into JP all such J-pairs whose signatures are not reducible by G or H (storing only one J-pair for each distinct J-signature). *iv)* Add the leading terms of the principle syzygies $vT_j - v_jT$ for $1 \leq j \leq \lvert U\rvert - 1$ to H.
Step 4.	While JP is not empty, go to step 1.
Return:	V and H.

FIGURE 2. the GVW algorithm applied to quotient rings

R-module M is generated by

$$(\mathbf{E}_1, \mathbf{g}_1), (\mathbf{E}_2, \mathbf{g}_2), \ldots, (\mathbf{E}_m, \mathbf{g}_m).$$

By now it should be clear that the GVW algorithm is a special case of this situation where $s = 1$ and is immediately applicable. The only differences that arise in this general case are in dealing with the leading monomials of the $\mathbf{v}$ part. Suppose $(\mathbf{u}_1, \mathbf{v}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2)$ are two pairs in $R^m \times R^s$, with $x^\alpha \mathbf{F}_j = \mathrm{lm}(\mathbf{v}_1)$ and $x^\beta \mathbf{F}_k = \mathrm{lm}(\mathbf{v}_2)$. We consider $(\mathbf{u}_2, \mathbf{v}_2)$ as a candidate to top-reduce $(\mathbf{u}_1, \mathbf{v}_1)$ only if $j = k$. Also, we only calculate the J-pair between $(\mathbf{u}_1, \mathbf{v}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2)$ if $j = k$. In

this case, assuming $\mathbf{v}_1, \mathbf{v}_2 \neq 0$, we have

$$t = \operatorname{lcm}(x^\alpha, x^\beta), \quad t_1 = \frac{t}{x^\alpha}, \quad t_2 = \frac{t}{x^\beta},$$

and if $t_i \mathbf{u}_i = \max\{t_1\mathbf{u}_1, t_2\mathbf{u}_2\}$, then $t_i(\mathbf{u}_i, \mathbf{v}_i)$ is a J-pair. Everything else proceeds as before.

5. Conclusions

We have presented a simple and fast algorithm for computing Gröbner bases for ideals and modules (including syzygy modules). Our algorithm is more flexible than F5 and our previous algorithm G2V [10] in that we allow a Gröbner basis to be computed incrementally, in one-shot, or a hybrid of the two. It is in this flexibility that we achieve an efficiency boost over G2V as some monomial orderings perform better than others.

In terms of simplicity, GVW is as simple as Buchberger's algorithm making implementation an easy matter. In terms of speed, we have shown that GVW derives its efficiency from the use of the syzygy module in preventing future reductions to zero and allowing GVW to outperform other known algorithms by at least a factor of 4 to 20 times. We believe that F4 style fast reductions are possible within the context of our algorithm, but the question remains as to how to implement it efficiently.

References

[1] Buchberger, B. *Ein Algorithmus zum Auffinden der Basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal.* PhD thesis, Leopold-Franzens University, 1965.

[2] Buchberger, B. A criterion for detecting unnecessary reductions in the construction of Gröbner bases. In *EUROSAM '79: Proceedings of the International Symposiumon on Symbolic and Algebraic Computation* (London, UK, 1979), Springer-Verlag, pp. 3–21.

[3] Buchberger, B. *Gröbner-Bases: An Algorithmic Method in Polynomial Ideal Theory.* Reidel Publishing Company, Dodrecht - Boston - Lancaster, 1985.

[4] Courtois, N., Klimov, E., Patarin, J., and Shamir, A. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In *In Advances in Cryptology, Eurocrypt2000, LNCS 1807* (2000), Springer-Verlag, pp. 392–407.

[5] Ding, J., Buchmann, J., Mohamed, M. S. E., Mohamed, W. S. A. E., and Weinmann, R.-P. MutantXL. In *First International Conference on Symbolic Computation and Cryptography* (2008), Springer-Verlag.

[6] Eder, C., and Perry, J. F5C: A variant of Faugère's F5 algorithm with reduced Gröbner bases. *Journal of Symbolic Computation 45*, 12 (2010), 1442 – 1458. MEGA'2009.

[7] Faugère, J. C. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra 139*, 1-3 (1999), 61 – 88.

[8] Faugère, J. C. A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In *ISSAC '02: Proceedings of the 2002 international symposium on Symbolic and algebraic computation* (New York, NY, USA, 2002), ACM, pp. 75–83.

[9] Faugère, J. C., and Joux, A. Algebraic cryptanalysis of hidden field equation (HFE) cryptosystems using gröbner bases. In *In Advances in Cryptology CRYPTO 2003* (2003), Springer, pp. 44–60.

[10] Gao, S., Guan, Y., and Volny IV, F. A new incremental algorithm for computing Gröbner bases. In *ISSAC'10: Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation* (Munich, Germany, 2010), ACM, pp. 13–19.

[11] Hashemi, A., and Ars, G. Extended F5 criteria. *Journal of Symbolic Computation 45*, 12 (2010), 1330 – 1340. MEGA'2009.

[12] Lazard, D. Gröbner-bases, Gaussian elimination and resolution of systems of algebraic equations. In *EUROCAL '83: Proceedings of the European Computer Algebra Conference on Computer Algebra* (London, UK, 1983), Springer-Verlag, pp. 146–156.

[13] Möller, H. M., Mora, T., and Traverso, C. Gröbner bases computation using syzygies. In *ISSAC '92: Papers from the international symposium on Symbolic and algebraic computation* (New York, NY, USA, 1992), ACM, pp. 320–328.

[14] Patarin, J. Hidden fields equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. In *EUROCRYPT'96: Proceedings of the 15th annual international conference on Theory and application of cryptographic techniques* (Berlin, Heidelberg, 1996), Springer-Verlag, pp. 33–48.

[15] Sun, Y., and Wang, D. A new proof of the F5 algorithm, December 2009.

Department of Mathematical Sciences, Clemson University, Clemson, SC 29634-0975 USA *E-mail address*: sgao@clemson.edu

Department of Mathematical Sciences, Clemson University, Clemson, SC 29634-0975 USA *E-mail address*: fvolny@clemson.edu

Information Security Lab, Institute of Software, Chinese Academy of Sciences, P.O. Box 8718, Beijing 100080, P. R. China *E-mail address*: mingsheng_wang@yahoo.com.cn