A new efficient algorithm for computing Gröbner bases (F_4) ☆

Jean-Charles Faugère

LIP6/CNRS Université Paris VI, case 168, 4 pl. Jussieu, F-75252 Paris Cedex 05, France

Abstract

This paper introduces a new efficient algorithm for computing Gröbner bases. To avoid as much intermediate computation as possible, the algorithm computes successive truncated Gröbner bases and it replaces the classical polynomial reduction found in the Buchberger algorithm by the simultaneous reduction of several polynomials. This powerful reduction mechanism is achieved by means of a symbolic precomputation and by extensive use of sparse linear algebra methods. Current techniques in linear algebra used in Computer Algebra are reviewed together with other methods coming from the numerical field. Some previously untractable problems (Cyclic 9) are presented as well as an empirical comparison of a first implementation of this algorithm with other well known programs. This comparison pays careful attention to methodology issues. All the benchmarks and CPU times used in this paper are frequently updated and available on a Web page. Even though the new algorithm does not improve the worst case complexity it is several times faster than previous implementations both for integers and modulo p computations.
ⓒ 1999 Published by Elsevier Science B.V. All rights reserved.

MSC: 14Q20; 14-04; 13P10; 15A36

1. Introduction

In view of the progress already achieved and the promising potential of current and planned algorithms, polynomial solving could become one of the more attractive application of Computer Algebra: practical problems can be solved, the algorithms are competitive with numerical methods. The main conclusion to be drawn from practice and experience of solving polynomial systems coming from various fields (industrial [19] problems, pure mathematics [20]) is the following: first of all, even though the computation of a Gröbner basis is a crucial point it must be emphasized that it is

only *one* step in the full solving process (change of ordering, triangular systems, real or numerical roots are complementary tools); secondly classical Buchberger algorithms [6,7,9] must be improved since even the best implementations often do not succeed to compute Gröbner bases from big problems. This paper is concerned with describing a new algorithm (whose name is F_4) for computing Gröbner basis. Even if the algorithm works for any admissible ordering, the algorithm F_4 has been designed to be efficient for a degree reverse lexicographical ordering (DRL); computing efficiently a lexico-graphical Gröbner basis from an already computed Gröbner basis being the task of another algorithm. Paradoxically, if the Buchberger algorithm without optimizations is very simple to describe it becomes much harder to understand how to improve a real implementation. By and large, however, it may eventually be possible to suggest two improvements: since 90% of the times is spent computing zero it would be useful to have more powerful criterion to remove useless critical pairs [8] (a powerful theoretical criteria exists but it is too costly); this crucial aspect of the problem is *not studied in this paper*, but is implemented in another algorithm (F_5 [18]). The second improve-ment is concerned with strategies: during a Gröbner computation, several choices can be made (select a critical pair, choose a reductor) and even if strategies have been proposed ([25] or even [24]) the heuristics which they rely on could not be satisfac-torily explained. So it is difficult to be convinced that they are optimal optimizations. Another bad consequence is that it is very difficult to (massively) parallelize the Buch-berger algorithm because the sugar (for instance) strategy imposes a strong sequential ordering. The primary objective of this paper is to propose a more powerful reduction algorithm. For that purpose we will reduce *simultaneously* several polynomials by a list of polynomials by using linear algebra techniques which ensure a global view of the process.

The plan of the paper is as follows. The main Section 2 is devoted to presenting the new algorithm. This section has been divided into several parts: in Section 2.1, we review the necessary mathematical notations (we make the choice to use the same notations as in the book [4]) and in Section 2.2 we establish the link between linear algebra (matrices) and polynomial algebra. Then we present in Section 2.3 a basic version of the algorithm without any criteria to eliminate useless pairs. An improved version of the algorithm including the Buchberger criteria is then given in Section 2.4. We close this discussion in Section 2.5 by motivating the choice of a good selection strategy (it seems that selecting all critical pairs with a minimal exponent is a good strategy). Since the algorithm relies heavily on linear algebra, Section 3 contains a short survey of linear algebra techniques we have implemented. A first version of this algorithm has been implemented in C in a new small system called FGb (for Fast Gb). In Section 4 we report an experimental evaluation of this first implementation. The best Gröbner bases programs are compared on a set of well known benchmarks and industrial problems. Finally, in Section 5 we outline the main features of the algorithm along with a list of possible related works and open issues.

The name of this algorithm is simply algorithm number 4. In the rest of this paper F_4 stands for this algorithm.

2. Description of the F_4 algorithm

2.1. Standard notations

We use the notations of [4] for basic definitions: R is the ground ring, $R[\underline{x}] = R[x_1, \ldots, x_n]$ is the polynomial ring. We denote by $T(x_1, \ldots, x_n)$, or simply by T, the set of all terms in these variables. We choose $<$ an admissible ordering on T. If $t = x_1^{\alpha_1} \cdots x_n^{\alpha_n} \in T$, then the *total degree* of t is defined as $\deg(t) = \sum_{i=1}^{n} \alpha_i$. Now let $f \in R[\underline{x}]$, $f \neq 0$, so that $f = \sum c(\alpha_1, \ldots, \alpha_n) x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ (where $c(\alpha_1, \ldots, \alpha_n)$ are elements of R). Then we define the set $M(f)$ of *monomials of* f as $M(f) = \{c(\alpha_1, \ldots, \alpha_n) x_1^{\alpha_1} \cdots x_n^{\alpha_n} \mid c(\alpha_1, \ldots, \alpha_n) \neq 0\}$. The set $T(f)$ of *terms of* f is $T(f) = \{x_1^{\alpha_1} \cdots x_n^{\alpha_n} \mid c(\alpha_1, \ldots, \alpha_n) \neq 0\}$. The *total degree* of $f \neq 0$ is defined as $\deg(f) = \max\{\deg(t) \mid t \in T(f)\}$. We define the *head term* $\mathrm{HT}(f)$, the *head monomial* $\mathrm{HM}(f)$, and the *head coefficient* $\mathrm{HC}(f)$ of f w.r.t. $<$ as follows: $\mathrm{HT}(f) = \max(T(f))$, $\mathrm{HM}(f) = \max(M(f))$, and $\mathrm{HC}(f) =$ the coefficient of $\mathrm{HM}(f)$. If F is a subset of $R[\underline{x}]$ we can extend the definition $\mathrm{HT}(F) = \{\mathrm{HT}(f) \mid f \in F\}$, $\mathrm{HM}(F) = \{\mathrm{HM}(f) \mid f \in F\}$ and $T(F) = \{T(f) \mid f \in F\}$; $Id(F)$ denotes the ideal generated by F.

Let $f, g, p \in R[\underline{x}]$ with $p \neq 0$, and let F be a finite subset of $R[\underline{x}]$. Then we say that

- f *reduces to* g *modulo* p (notation $f \xrightarrow{p} g$), if $\exists t \in T(f)$, $\exists s \in T$ such that $s * \mathrm{HT}(p) = t$ and $g = f - \frac{a}{\mathrm{HC}(p)} * s * p$ where a is the coefficient of t in p.
- f *reduces to* g *modulo* P (notation $f \xrightarrow{P} g$), if $f \xrightarrow{p} g$ for some $p \in P$.
- f *is reducible modulo* p if there exists $g \in R[\underline{x}]$ such that $f \xrightarrow{P} g$.
- f *is reducible modulo* P if there exists $g \in R[\underline{x}]$ such that $f \xrightarrow{P} g$.
- f *is top reducible modulo* P if there exists $g \in R[\underline{x}]$ such that $f \xrightarrow{P} g$ and $\mathrm{HT}(g) < \mathrm{HT}(f)$.
- $f \xrightarrow{*}{P} g$ is the reflexive-transitive closure of $\xrightarrow{P}$.
- The *S-polynomial* of f and g is defined as

$$spol(f, g) = \mathrm{HC}(g) \frac{lcm(\mathrm{HT}(f), \mathrm{HT}(g))}{\mathrm{HT}(f)} f - \mathrm{HC}(f) \frac{lcm(\mathrm{HT}(f), \mathrm{HT}(g))}{\mathrm{HT}(g)} g$$

2.2. Linear algebra and polynomials

Definition 2.1. By convention if M is a $s \times m$ matrix, $M_{i,j}$ is the jth element of the ith row of M. If $T_M = [t_1, \ldots, t_m]$ an ordered set of terms, let $(\varepsilon_i)_{i=1,\ldots,m}$ be the canonical basis of R^m, we consider the linear map $\varphi_{T_M} : V_{T_M} \to R^m$ (where V_{T_M} is the submodule of $R[\underline{x}]$ generated by T_M) such that $\varphi_{T_M}(t_i) = \varepsilon_i$. The reciprocal function will be denoted by ψ_{T_M}. The application ψ_{T_M} allows to interpret vectos of R^n as polynomials. We note by (M, T_M) a matrix with such an interpretation.

Definition 2.2. If (M, T_M) is a $s \times m$ matrix with such an interpretation, then we can construct the set of polynomials:

$$Rows(M, T_M) := \{\psi_{T_M}(row(M, i)) \mid i = 1, \dots, s\} \setminus \{0\}$$

where $row(M, i)$ is the ith row of M (an element of R^m). Conversely, if l is a list of polynomials and T_l an ordered set of terms we can construct an $s \times m$ matrix A (where $s = \text{size}(l)$, $m = \text{size}(T_l)$):

$$A_{i,j} := \text{coeff}(l[i], T_l[j]), \quad i = 1, \dots, s, \ j = 1, \dots, m.$$

We note $A^{(l, T_l)}$ the matrix $(A_{i,j})$.

$$
M =
\begin{array}{c}
\\
f_{i_1} \\
f_{i_2} \\
\vdots \\
f_{i_{2k}} \\
f_{i_{2k+1}} \\
f_{i_s}
\end{array}
\begin{array}{cccc}
t_1 & t_2 & t_3 & \dots \\
\hline
\times & \times & \times & \cdots \\
\times & \times & \times & \cdots \\
\vdots & \vdots & \vdots & \dots \\
\times & \times & \times & \cdots \\
\times & \times & \times & \cdots \\
\times & \times & \times & \cdots
\end{array}
\tag{1}
$$

Definition 2.3. Let M be a $s \times m$ matrix, and $\underline{Y} = [Y_1, \dots, Y_m]$ new variables. Then $F = Rows(M, \underline{Y})$ is a set of equations, so we can compute $\tilde{F}$ a reduced Gröbner basis of F for a lexicographical ordering such that $Y_1 > \cdots > Y_m$. From this basis we can reconstruct a matrix $\tilde{M} = A^{(\tilde{F}, \underline{Y})}$. We called $\tilde{M}$ the (unique) row echelon form[1] of M. We say also that $\tilde{F}$ is a row echelon basis of F.

$$
M =
\begin{array}{c}
\\
f_{j_1} \\
f_{j_2} \\
\\
f_{jl} \\
f_{jl+1} \\
\\
f_{j_m}
\end{array}
\begin{array}{ccccccc}
t_1 & t_2 & & t_k & t_{k+1} & & t_m \\
\hline
1 & 0 & \dots & 0 & \times & \dots & \times \\
0 & 1 & \dots & 0 & \times & \dots & \times \\
 & & \ddots & & \times & \dots & \times \\
0 & 0 & \dots & 1 & \times & \dots & \times \\
0 & 0 & \dots & 0 & 0 & \dots & 0 \\
 & & \vdots & & 0 & \dots & 0 \\
0 & 0 & \dots & 0 & 0 & \dots & 0
\end{array}
\tag{2}
$$

where $\times$ denotes a possibly non-zero element.

In the case of polynomials we have a similar definition:

Definition 2.4. Let F be a finite subset of $R[\underline{x}]$ and $<$ an admissible ordering. We define $T_<(F)$ to be $Sort(\{T(f) \mid f \in F\}, <)$, $A := A^{(F, T_<(F))}$ and $\tilde{A}$ the row echelon form of A. We say that $\tilde{F} = Rows(\tilde{A}, T_<(F))$ is the row echelon form of F w.r.t. $<$.

---

[1] In some computer algebra system, this is the only way to compute a row echelon form !

Elementary properties of row echelon matrices are summarized by the following theorem:

Theorem 2.1. *Let M be a $s \times m$ matrix, and $\underline{Y} = [Y_1, \ldots, Y_m]$ new variables, $F = \text{Rows}(M, \underline{Y})$, $\tilde{M}$ the row echelon form of M, $\tilde{F} = \text{Rows}(\tilde{M}, \underline{Y})$. We define*

$$\tilde{F}^+ = \{g \in \tilde{F} \mid \text{HT}(g) \notin \text{HT}(F)\},$$
$$\tilde{F}^- = \tilde{F} \backslash \tilde{F}^+.$$

For any subset F_- of F such that $\text{size}(F_-) = \text{size}(\text{HT}(F))$ and $\text{HT}(F_-) = \text{HT}(F)$, then $G = \tilde{F}^+ \cup F_-$ is a triangular basis of the R-module V_M generated by F. That is to say, for all $f \in V_M$ there exist $(\lambda_k)_k$ elements of R and $(g_k)_k$ elements of G such that $f = \sum_k \lambda_k g_k$, $\text{HT}(g_1) = \text{HT}(f)$ and $\text{HT}(g_k) > \text{HT}(g_{k+1})$.

Proof. Since the head terms of G are pairwise distinct, G is linearly independent. We claim that it is also a generating system of V_M. Suppose for a contradiction that there exists $f \in V_M$ such that $f \xrightarrow[G]{*} f' \neq 0$. By definition of a Gröbner basis, $f' \xrightarrow[\tilde{F}]{*} 0$, consequently f' is top reducible modulo $\text{HT}(\tilde{F}) = \text{HT}(\tilde{F}^+) \cup \text{HT}(\tilde{F}^-) = \text{HT}(\tilde{F}^+) \cup \text{HT}(F_-) = \text{HT}(G)$, so that f' is top-reducible modulo G. This is a contradiction. □

We can transpose immediately the theorem for polynomials:

Corollary 2.1. *Let F be a finite subset of E and $<$ an admissible ordering, and $\tilde{F}$ the row echelon form of F w.r.t. $<$. We define*

$$\tilde{F}^+ = \{g \in \tilde{F} \mid \text{HT}(g) \notin \text{HT}(F)\}$$

For all subset F_- of F such that $\text{size}(F_-) = \text{size}(\text{HT}(F))$ and $\text{HT}(F_-) = \text{HT}(F)$, then $G = \tilde{F}^+ \cup F_-$ is a triangular basis of V_M the R-module generated by F. For all $f \in V_M$ there exist $(\lambda_k)_k$ elements of R and $(g_k)_k$ elements of G such that $f = \sum_k \lambda_k g_k$, $\text{HT}(g_1) = \text{HT}(f)$ and $\text{HT}(g_k) > \text{HT}(g_{k+1})$.

2.3. The F_4 algorithm

It is well known that during the execution of the Buchberger algorithm, one has a lot of choices:
- select a critical pair in the list of critical pairs;
- choose one reductor among a list of reductors when reducing a polynomial by a list of polynomials.

Buchberger [6] proves that these choices are not important for the correctness of the algorithm, but it is well known that these choices are crucial for the total time computation. Moreover, the best strategies [25] inspect only the leading terms of the polynomials to make a choice. Consider the case where all the input polynomials have the same leading term. In that case, all the critical pairs are equal and it is not possible to take a decision. In some sense this problem can be corrected in a simple

and surprising way: *we make no choice.* More precisely instead of choosing *one* critical pair at each step, we select a *subset* of critical pairs at the same time. So, in fact, we are delaying the necessary choices in a second step of the algorithm, the linear algebra part of the algorithm.

Definition 2.5. A critical pair of two polynomials (f_i, f_j) is an element of $T^2 \times R[\underline{x}] \times T \times R[\underline{x}]$, $Pair(f_i, f_j) := (lcm_{ij}, t_i, f_i, t_j, f_j)$ such that

$$lcm(Pair(f_i, f_j)) = lcm_{ij} = \mathrm{HT}(t_i f_i) = \mathrm{HT}(t_j f_j) = lcm(\mathrm{HT}(f_i), \mathrm{HT}(f_j)).$$

Definition 2.6. We say that the degree of the critical pair $p_{i,j} = Pair(f_i, f_j)$, $deg(p_{i,j})$, is $deg(lcm_{i,j})$. We define the two projections $Left(p_{i,j}) := (t_i, f_i)$ and $Right(p_{i,j}) := (t_j, f_j)$. If $(t, p) \in T \times R[\underline{x}]$ then we note $mult((t, p))$ the evaluated product $t * p$.

We have now the tools needed to present the basic version of our algorithm. All the matrices occurring in following algorithms are the representation of a list of polynomials through the set of all their terms, as explained in Definition 2.2.

Algorithm $F4$

Input : $\begin{cases} F \text{ a finite subset of } R[\underline{x}] \\ \mathscr{S}el \text{ a function } List(Pairs) \rightarrow List(Pairs) \text{ such that } \mathscr{S}el(l) \neq \emptyset \text{ if } l \neq \emptyset \end{cases}$
Output : a finite subset of $R[\underline{x}]$.
$G := F, \tilde{F}_0^+ := F$ and $d := 0$
$P := \{Pair(f, g) \mid f, g \in G \text{ with } f \neq g\}$
while $P \neq \emptyset$ **do**
  $d := d + 1$
  $Pd := \mathscr{S}el(P)$
  $P := P \backslash P_d$
  $L_d := Left(P_d) \cup Right(P_d)$
  $\tilde{F}_d^+ := Reduction(L_d, G)$
  for $h \in \tilde{F}_d^+$ **do**
    $P := P \cup \{Pair(h, g) \mid g \in G\}$
    $G := G \cup \{h\}$
return G

We have to extend the reduction of a polynomial modulo a subset of $R[\underline{x}]$, to the reduction of a subset of $R[\underline{x}]$ modulo another subset of $R[\underline{x}]$:

Reduction

Input : $\begin{cases} L \text{ a finite subset of } T \times R[\underline{x}] \\ G \text{ a finite subset of } R[\underline{x}] \end{cases}$
Output : a finite subset of $R[\underline{x}]$ (possibly an empty set).
  $F := \text{Symbolic Preprocessing}(L, G)$

$\tilde{F}$:=Reduction to Row Echelon Form of F w.r.t. $<$

$\tilde{F}^{+}$:=$\{f \in \tilde{F} \mid HT(f) \notin HT(F)\}$

return $\tilde{F}^{-}$

Remark 2.1. By Lemma 2.1, we will see that an equivalent (but slower) definition of $\tilde{F}^{+}$ could be $\tilde{F}^{+}$:=$\{f \in \tilde{F} \mid f \text{ top irreducible by } G\}$.

We have now to describe the main function of our algorithm, that is to say the consruction of the "matrix" F. This subalgorithm can be viewed as an usual reduction of all the considered polynomials if we replace the standard arithmetic by: let $0 \neq x$, $0 \neq y \in R$, then $x + y = 1$, $x * y = 1$, $x * 0 = 0$ and $x + 0 = 1$. So this is really a *symbolic* preprocessing.

Symbolic Preprocessing

Input: $\begin{cases} L \text{ a finite subset of } T \times R[\underline{x}] \\ G \text{ a finite subset of } R[\underline{x}] \end{cases}$

Output: a finite subset of $R[\underline{x}]$.

F:=$\{t * f \mid (t, f) \in L\}$

$Done$:=HT(F)

while $T(F) \neq Done$ **do**

   m an element of $T(F) \backslash Done$

   $Done$:=$Done \cup \{m\}$

   if m top reducible modulo G **then**

      $m = m' * HT(f)$ for some $f \in G$ and some $m' \in T$

      F:=$F \cup \{m' * f\}$

return F

Remark 2.2. It seems that the initial values of Done should be $\emptyset$ but in all application of this function the result is in fact the same with less iterations.

Remark 2.3. The symbolic preprocessing is very efficient since its complexity is linear in the size of its outout if $size(G)$ is smaller than the final size of tf $T(F)$ which is usually the case.

Lemma 2.1. *Let G be finite subset of $R[\underline{x}]$, L be the image by* mult *of a finite subset of $T \times G$ and $\tilde{F}^{+} = Reduction(L, G)$. Then for all $h \in \tilde{F}^{+}$, $HT(h) \notin Id(HT(G))$.*

Proof. Let F the set computed by the algorithm Symbolic Preprocessing (L, G). Assume for a contradiction that $\exists h \in \tilde{F}^{+}$ such that $t = HT(h) \in Id(HT(G))$. Hence $HT(g)$ divides t for some $g \in G$. So t is in $T(\tilde{F}^{+}) \subset T(\tilde{F}) \subset T(F)$ and is top reducible by g, hence $\frac{t}{HT(g)} * g$ is inserted in F by Symbolic Preprocessing (or another product with the same head term). This contradicts the fact that $HT(h) \notin HT(F)$. $\square$

The following lemma is useful to proof the correctness of the algorithm.

Lemma 2.2. *Let G be finite subset of $R[\underline{x}]$, L be the image by* mult *of a finite subset of $T \times G$ and $\tilde{F}^+ = Reduction(L, G)$. Then $\tilde{F}^+$ is a subset of $Id(G)$. Moreover for all f in the R-module generated by L, $f \xrightarrow[G \cup \tilde{F}^+]{*} 0$.*

Proof. Apply the Corollary 2.1 to F the set generated by Symbolic Preprocessing(L, G). Clearly F is a subset of $F \cup Id(G)$, but it is obvious that L is a subset of $Id(G)$, so that F is a subset of $Id(G)$. Hence any F_- fulfilling the hypothesis of Theorem 2.1 is a subset of $Id(G)$. This conclude the proof of the lemma since the R-module generated by L is a submodule of the R-module generated by F. $\square$

Remark 2.4. Let G be a finite subset of $R[\underline{x}]$. It is possible that $f \xrightarrow[G]{*} 0$ but that *NormalForm* $(f, G) \neq 0$ where *NormalForm* is the reduction which is used in Buchberger algorithm. The reason for that is that the result of *NormalForm* depends on many choices (strategies).

Theorem 2.2. *The algorithm F_4 computes a Gröbner basis G in $R[\underline{x}]$ such that $F \subseteq G$ and $Id(G) = Id(F)$.*

Proof. *Termination*: Assume for a contradiction that the **while**-loop does not terminate. We see that there exists an ascending sequence (d_i) of natural numbers such that $\tilde{F}_{d_i}^+ \neq \emptyset$ for all i. Let say that $q_i \in \tilde{F}_{d_i}^+$ (hence q_i can be any element in $\tilde{F}_{d_i}^+$). Let U_i be $U_{i-1} + Id(\mathrm{HT}(q_i))$ for $i > 1$ and $U_0 = (0)$. By Lemma 2.1 we have $U_{i-1} \subsetneqq U_i$. This contradicts the fact that $R[\underline{x}]$ is noetherian.

Correctness: We have $G = \bigcup_{d \geq 0} \tilde{F}_d^+$. We claim that the following are loop invariants of the **while**-loop: G is a finite subset of $R[\underline{x}]$ such that $F \subset G \subset Id(F)$, and $spol(g_1, g_2) \xrightarrow[G]{*} 0$ for all $g_1, g_2 \in G$ such that $\{g_1, g_2\} \notin P$. The first claim is an immediate consequence of the first part of Lemma 2.2. For the second one, if $\{g_1, g_2\} \notin P$, this means that $Pair(g_1, g_2) = (lcm_{1,2}, t_1, g_1, t_2, g_2)$ has been selected in a previous step (say d) by the function $\mathscr{S}el$. Hence $t_1 * g_1$ and $t_2 * g_2$ are in L_d, so $spol(g_1, g_2)$ is an element of the R-module generated by L_d hence by Lemma 2.2 $spol(g_1, g_2) \xrightarrow[G]{*} 0$. $\square$

Remark 2.5. If $size(\mathscr{S}el(l)) = 1$ for all $l \neq \emptyset$ then the algorithm F_4 is the Buchberger algorithm. In that case the $\mathscr{S}el$ function corresponds to the selection strategy in the Buchberger algorithm.

Example 2.1. One might wonder why in the proof of the termination of the algorithm we consider only one element of $\tilde{F}_d^+$ and not the whole $\tilde{F}_d^+$.

If $x > y > z$ for a lexicographical ordering, $F = [f_1 = xy^2 + 1, f_2 = xz^2 + 1, f_3 = y^3 + y^2]$ and $\mathscr{S}el =$ identity, we find $P_1 = \{Pair(f_1, f_2), Pair(f_2, f_3), Pair(f_1, f_3)\}$ and $\tilde{F}_1^+ = \{y^2 - z^2, y + 1\}$ so that $Id(\mathrm{HT}(\tilde{F}_1^+)) = \{y\}$. So contrarily to Buchberger Algorithm this not true that after each operation $G' := G \cup \{h\}$, we have $Id(\mathrm{HT}(G')) \supsetneqq Id(\mathrm{HT}(G))$.

2.4. Buchberger criteria. Improved F_4 algorithms

In order to obtain an efficient algorithm we need to insert into the previous algorithm the Buchberger Criteria. Since it is not the subject of this paper to improve the Buchberger Criteria we will use a standard implementation of these criteria such as the Gebauer and Moller installation [21]:

Buchberger Criteria

Specification : $\begin{cases} (G_{new}, P_{new}) := Update(G_{old}, P_{old}, h) \\ \text{Update of critical pair list and ideal basis (see [4], p. 230)} \end{cases}$

Input : $\begin{cases} \text{a finite subset } G_{old} \text{ of } R[\underline{x}] \\ \text{a finite set } P_{old} \text{ of critical pairs of } R[\underline{x}] \\ 0 \neq h \in R[\underline{x}] \end{cases}$

Output : a finite subset of $R[\underline{x}]$ and a list of critical pairs.

In the previous version of the algorithm we used only *some* rows of the reduced matrix (the sets $\tilde{F}^+$), rejecting the rows which were already in the original matrix (the sets TF). In the new version of the algorithm we keep these useless rows, and we try to replace some products $m * f$ occurring in the rows of the "matrix" F by a new "equivalent" product $m' * f'$ with $m \geq m'$. This is the task of the function Simplify: $T \times R[\underline{x}] \times List(Subset(R[\underline{x}])) \to T \times R[\underline{x}]$. The third argument of Simplify is the list of all the already computed matrices. A complete description of this function will be given below.

Improved Algorithm $F4$

Input : $\begin{cases} F \text{ a finite subset of } R[\underline{x}] \\ \mathscr{S}el \text{ a function } List(Pairs) \to List(Pairs) \text{ such that } \mathscr{S}el(l) \neq \emptyset \text{ if } l \neq \emptyset \\ \text{Update the part of Buchberger algorithm which select the pairs to} \\ \text{compute, using the criteria like the algorithm of p. 230 in [4].} \end{cases}$

Output : a finite subset of $R[\underline{x}]$.

$G := \emptyset$ and $P := \emptyset$ and $d := 0$

while $F \neq \emptyset$ **do**

   $f := first(F)$

   $F := F \setminus \{f\}$

   $(G, P) := Update(G, P, f)$

while $P \neq \emptyset$ **do**

   $d := d + 1$

   $P_d := \mathscr{S}el(P)$

   $P := P \setminus P_d$

   $L_d := Left(P_d) \cup Right(P_d)$

   $(\tilde{F}_d^+, F_d) := Reduction(L_d, G, (F_i)_{d=1,\ldots,(d-1)})$

   for $h \in \tilde{F}_d^+$ **do**

      $(G, P) := Update(G, P, h)$

return G

The new Reduction function is identical to the previous version except that there is a new argument and that it returns also the result of Symbolic Preprocessing:

Reduction

$$\textbf{Input}: \begin{cases} L \text{ a finite subset of } T \times R[\underline{x}] \\ G \text{ a finite subset of } R[\underline{x}] \\ \mathscr{F} = (F_k)_{k=1,\dots,(d-1)}, \text{ where } F_k \text{ is finite subset of } R[\underline{x}] \end{cases}$$

Output: two finite subsets of $R[\underline{x}]$.

F:=Symbolic Preprocessing $(L, G, \mathscr{F})$

$\tilde{F}$:= Reduction to Row Echelon Form of F w.r.t. $<$

$\tilde{F}^+ := \{f \in \tilde{F} \mid HT(f) \notin HT(F)\}$

return $(\tilde{F}^+, F)$

Symbolic Preprocessing

$$\textbf{Input}: \begin{cases} L \text{ a finite subset of } T \times R[\underline{x}] \\ G \text{ a finite subset of } R[\underline{x}] \\ \mathscr{F} = (F_k)_{k=1,\dots,(d-1)}, \text{ where } F_k \text{ is finite subset of } R[\underline{x}] \end{cases}$$

Output: a finite subset of $R[\underline{x}]$.

$F := \{\text{mult}(\text{Simplify}(m, f, \mathscr{F})) \mid (m, f) \in L\}$

$Done$:=HT(F)

while T$(F) \neq Done$ **do**

   m an element of T$(F) \backslash Done$

   $Done$:=$Done \cup \{m\}$

   if m top reducible modulo G **then**

   $m = m' * \text{HT}(f)$ for some $f \in G$ and some $m' \in T$

   F:=$F \cup \{\text{mult}(\text{Simplify}(m', f, \mathscr{F}))\}$

return F

Simplify

$$\textbf{Input}: \begin{cases} t \in T \text{ a term} \\ f \in R[\underline{x}] \text{ a polynomial} \\ \mathscr{F} = (F_k)_{k=1,\dots,(d-1)}, \text{ where } F_k \text{ is finite subset of } R[\underline{x}] \end{cases}$$

Output: a non evaluated product, i.e. an element of $T \times R[\underline{x}]$.

for $u \in$ list of divisors of t **do**

   if $\exists j (1 \leq j < d)$ such that $(u * f) \in F_j$

      then

      $\tilde{F}_j$ is the row echelon form of F_j w.r.t. $<$

         there exists a (unique) $p \in \tilde{F}_j^+$ such that $\text{HT}(p) = \text{HT}(u * f)$

         if $u \neq t$ **then**

            return Simplify$(\dfrac{t}{u}, p, \mathscr{F})$

         else

            return $(1, p)$

return (t, f)

Lemma 2.3. *If* (t', f') *is the result of Simplify* $(t, f, \mathscr{F})$ *then* $\mathrm{HT}(t' * f') = \mathrm{HT}(t * f)$. *Moreover if* $\tilde{\mathscr{F}}$ *denotes* $(\tilde{F}_k^+)_{k=1,\dots,(d-1)}$ *there exist* $0 \neq \lambda \in R$, *and* $r \in R$*-module* $(\tilde{\mathscr{F}} \cup \mathscr{F})$ *such that* $tf = \lambda t'f' + r$ *with* $\mathrm{HT}(r) < \mathrm{HT}(t * f)$.

Proof. *Termination*: Simplify constructs a sequence (t_k, f_k) such that $t_0 = t$, $f_0 = f$ and $t_{k+1} < t_k$ except perhaps for the last step. T is noetherian, this implies that the algorithm stops after r_k steps.

Correctness: The first part is obvious since $\mathrm{HT}(u_k f_k) = \mathrm{HT}(f_{k-1})$ so that $\mathrm{HT}(t_k f_k) = \mathrm{HT}(\frac{t_k}{u_k} f_{k+1}) = \mathrm{HT}(t_{k+1} f_{k+1})$. The proof is by induction on the step number. So we suppose $r_k = 1, t' = \frac{t}{u}$ and $u * f \in F_j, f' \in \tilde{F}_j$ for some j with $\mathrm{HT}(f') = \mathrm{HT}(u * f)$. The set $F_- = \{uf\}$ can be supplemented by other elements of F_j such that $\mathrm{HT}(F_-) = \mathrm{HT}(F)$ and $size(F_-) = size(\mathrm{HT}(F))$. We can apply Corollary 2.1 we find $(\alpha_k) \in R$, $g_k \in F_- \cup (\tilde{F}_j)_+$, such that $f' = \sum_k \alpha_k g_k$ and $\mathrm{HT}(g_1) = \mathrm{HT}(f')$ and $\mathrm{HT}(f') > \mathrm{HT}(g_k)$ for $k > 2$. By construction of $F_-, g_1 = u * f$. Hence $f' = \alpha_1 uf + r$ with $\mathrm{HT}(r) < \mathrm{HT}(f')$, consequently $\alpha_1 \neq 0$ and we have $tf = \frac{1}{\alpha_1} t'f' - \frac{1}{\alpha_1} t'r$. $\square$

Remark 2.6. Experimental observation establishes that the effect of *Simplify* is to return, in 95% of the cases, a product (x_i, p) where x_i is a variable (and frequently the product (x_n, p)). This technique is very similar to the FGLM algorithm for computing normal forms by using matrix multiplications.

For the verification of the improved version of the algorithm we recall the following definition and theorem ([4], p. 219):

Definition 2.7. Let P be a finite subset of $R[\underline{x}]$, $0 \neq f \in R[\underline{x}]$, and $t \in T$. Suppose $f = \sum_{i=1}^k m_i p_i$ with monomials $0 \neq m_i \in R[\underline{x}]$ and $p_i \in P$ not necessarily pairwise different $(1 \leq i \leq k)$. Then we say that this is a *t-representation* of f w.r.t. P, if $\mathrm{HT}(m_i p_i) \leq t$ for all $i = 1, 2, \dots, k$.

Theorem 2.3. *Let* G *be a finite subset of* $R[\underline{x}]$ *with* $0 \in G$. *Assume that for all* $g_1, g_2 \in G$, $spol(g_1, g_2)$ *either equals zero or it has a t-representation w.r.t.* G *for some* $t < lcm(\mathrm{HT}(g_1), \mathrm{HT}(g_2))$. *Then* G *is a Gröbner basis.*

Theorem 2.4. *Let* F *be a finite subset of* $R[\underline{x}]$, $\mathscr{F} = (F_k)_{k=1,\dots,(d-1)}$, *where* F_k *is finite subset of* $R[\underline{x}]$, $Pair(g_1, g_2) = (lcm_{1,2}, t_1, g_1, t_2, g_2)$ *with* $lcm_{1,2}, t_1, t_2 \in T$ *such that the following hold*:

   (i) F_k *is the image by* mult *of a finite subset of* $T \times F$

   (ii) $(\tilde{F}_k)_- \subset G$ *for* $k = 1, \dots, (d-1)$ ($\tilde{F}_k$ *being as usual the row echelon of* F_k)

   (iii) $f_i = \mathrm{mult}(\mathrm{Simplify}(t_i, g_i, \mathscr{F}))$ *for* $i = 1, 2$.

   (iv) $spol(f_1, f_2)$ *has a t-representation w.r.t.* F *with* $t < lcm(\mathrm{HT}(f_1), \mathrm{HT}(f_2))$.

Then the S-polynomial $spol(g_1, g_2)$ *has a* t'-*representation w.r.t.* F *with* $t' < lcm_{1,2}$.

Proof. Let (t'_i, g'_i) be *Simplify*$(t_i, g_i, \mathscr{F})$. By Lemma 2.3 we have $\mathrm{HT}(t'_1 g'_1) = \mathrm{HT}(t_1 g_1) = lcm_{1,2} = \mathrm{HT}(t_2 g_2) = \mathrm{HT}(t'_2 g'_2)$ so that (we suppose that all the polynomial are monics):

$$spol(g'_1, g'_2) = t'_1 g'_1 - t'_2 g'_2$$
$$= (t'_1 g'_1 - t_1 g_1) + (t_1 g_1 - t_2 g_2) + (t_2 g_2 - t'_2 g'_2)$$
$$= r + spol(f_1, f_2) + r'$$

with $r, r' \in Id(\tilde{\mathscr{F}} \cup \mathscr{F}) \subset Id(G)$ such that $max(\mathrm{HT}(r), \mathrm{HT}(r')) < lcm_{1,2}$. Hence $spol(g_1, g_2)$ has a t'-representation for $t' = max(\mathrm{HT}(r), \mathrm{HT}(r'), t) < lcm_{1,2}$. $\square$

2.5. Selection strategy

The choice of a good selection strategy, that is to say the choice of the function $\mathscr{S}el$, is very important for the performance of the algorithm.

Computing Gröbner bases for a degree ordering is very frequently the most difficult step in the solving process (other steps are elimination or decomposition of the ideal). One reason for that is that the input of the algorithm is only a subset of $R[\underline{x}]$ with no mathematical structure. We want to give some structure to these polynomials at the beginning of the computation: we use the concept of d-Gröbner bases:

Definition 2.8. If G_d is the result of the Buchberger algorithm truncated to the degree d (that is to say we reject all critical pairs whose total degree (Definition 2.6) are $> d$), then we call G_d a (truncated) d-Gröbner basis of I.

The following theorem give a structure to this list when the polynomials are homogeneous.

Theorem 2.5 ([34,4, p. 471]). *For homogeneous polynomials $f_1, \ldots, f_l$, G_d is a Gröbner basis "up to degree d" that is to say:*

- $\xrightarrow[G_d]{*}$ *is well defined for polynomials f such that* $\deg(f) \leq d$.
- $\forall p \in I$ *s.t.* $\deg(p) \leq d \Rightarrow p \xrightarrow[G_d]{*} 0$.
- $spol(f, g) \xrightarrow[G_d]{*} 0$ *for f, g in G_d such that* $\deg(lcm(\mathrm{HT}(f), \mathrm{HT}(g))) \leq d$.

Moreover, there exists a D_0 such that for all $d \geq D_0$, $G_d = G_{D_0}$ is the Gröbner basis of I. We note by ∞ the number D_0.

An effective Nullstellensatz may give an estimate of D_0; from a theoretical point of view such an explicit bound for the degrees reduces the problem of finding the polynomials G_∞ to the resolution of a system of linear equations. This reduction to linear algebra of the computation Gröbner bases has been used for a long time for analysing the complexity of Buchberger algorithm [34].

For practical computation this does not work well since:

- D_0 is often over estimated.

- The linear system is huge: the matrix which is generated is frequently larger that really needed.
- The matrix of the linear system has a generalized sylvester structure and solving efficiently such a system is not a well known task.

Other algorithms are also closely related to linear algebra [3,24,32,33,35].

In fact the Buchberger algorithm and the F_4 algorithm give *incremental* methods to solve this systems. The new algorithm will compute G_{d+1} from G_d. Thus the algorithm transforms a mathematical object (G_d is unique) into another object with a stronger structure. In fact the Buchberger algorithm is also incremental since it computes one polynomial after another but in our case we compute a whole new truncated basis.

$$G_0 \to G_1 \to G_2 \to \cdots \to G_d \to G_{d+1} \to \cdots$$

Unfortunately Theorem 2.5 is false for non homogeneous polynomials. One solution to overcome this difficulty is to homogenize the list of polynomials but it is not efficient for big systems since parasite solutions are also computed (solutions at infinity can be of greater dimension); a better method is to consider the sugar degree [25] of each polynomial: we add a "phantom" homogenization variable and the sugar degree $\deg_S$ of a polynomial is the degree of the polynomial if all the computations were performed with the additional variable:

Definition 2.9. For the initial polynomials: $\deg_S(f_i) = \deg(f_i)$, for all $i = 1, \ldots, m$. The polynomials occurring in the computation have the following forms $p + q$ or $t * p$ where $t \in T$ is a term. We define $\deg_S(p + q) = \max(\deg_S(p), \deg_S(q))$ and $\deg_S(m * p) = \deg(m) + \deg_S(p)$. We say that $\deg_S(q)$ is the "sugar" of q.

Definition 2.10. $G_d^{(S)}$ is the result of the Buchberger algorithm when we reject all critical pairs whose $\deg_S$ is $> d$. (We replace deg by $\deg_S$ is the Definition 2.8.)

The weak point of this approach is that $G_{d+1}^{(S)} \setminus G_d^{(S)}$ contains polynomials with various degrees and that near to the end of the computation $\deg_S(p) \gg \deg(p)$.

We give now some possible implementation of $\mathcal{S}el$. These results are not discussed in detail as they will been reported in a more technical paper.

- The easiest way to implement $\mathcal{S}el$ is to take the identity! In that case we really reduce all the critical pairs at the same time.
- The best function we have tested is to take all the critical pairs with a minimal total degree:

Sel(**P**)
    Input: P a list of critical pairs
    Output: a list of critical pairs
       $d := min\{\deg(lcm(p)), \; p \in P\}$
       $P_d := \{ p \in P \mid \deg(lcm(p)) = d \}$
       return P_d

We call this strategy the *normal strategy for F_4*. If the input polynomials are homogeneous, we already have a Gröbner basis up to degree $d - 1$ and $\mathscr{S}el$ selects exactly all the critical which are needed for computing a Gröbner basis up to degree d.

- We can also change the $\deg(Pair(f_i, f_j))$ to be the sugar degree, $\deg_S(lcm_{ij})$ (see Definition 2.6). In our experiments, this variant of the algorithm was less efficient.

2.6. Example

The reader should be aware that it is impossible to fully appreciate the efficiency of the algorithm for small examples. We consider the cyclic 4 problem. We choose a total degree reverse lexicographical ordering and the normal strategy.

$$F = [f_4 = abcd - 1, f_3 = abc + abd + acd + bcd, f_2 = ab + bc + ad + cd,$$
$$f_1 = a + b + c + d].$$

At the beginning $G = \{f_4\}$ and $P_1 = \{Pair(f_3, f_4)\}$ so that $L_1 = \{(1, f_3), (b, f_4)\}$. We enter now in Symbolic Preprocessing$(L_1, G, \emptyset)$; $F_1 = L_1$, $Done = \mathrm{HT}(F_1) = \{ab\}$ and $T(F_1) = \{ad, ab, b^2, bc, bd, cd\}$, we pick an element in $T(F_1)\backslash Done$, for instance ad, but ad is top reducible by G; we have $Done = \{ab, ad\}$, $F_1 = F_1 \cup \{df_4\}$ and $T(F_1) = T(F_1) \cup \{d^2\}$. Since the other elements of $T(F_1)$ are not top reducible by G, Symbollic Preprocessing returns $\{f_3, bf_4, df_4\}$. Or in matrix form:

$$A_1 = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

the row echelon form of A_1 is

$$\tilde{A}_1 = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & -1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 & 2 & 0 & 1 & 0 \end{bmatrix}$$

that is to say $\tilde{F}_1 = \{f_5 = ad + bd + cd + d^2, f_6 = ab + bc - bd - d^2, f_7 = b^2 + 2bd + d^2\}$ and since $ab, ad \in \mathrm{HT}(F_1)$ we have $\tilde{F}_{1_\_} = \{f_7\}$ and now $G = \{f_4, f_7\}$.

In the next step we have to study $P_2 = \{Pair(f_2, f_4)\}$, thus $L_2 = \{(1, f_2), (bc, f_4)\}$ and $\mathscr{F} = \{F_1\}$. In Symbollic Preprocessing we first try to simplify $(1, f_2)$ and (bc, f_4) with $\mathscr{F}$. We see that $bf_4 \in F_1$ and that f_6 is the unique polynomial in $\tilde{F}_1$ such that $\mathrm{HT}(f_6) = \mathrm{HT}(bf_4) = ab$, hence Simplify $(bc, f_4, \mathscr{F}) = (c, f_6)$. Now $F_2 = \{f_2, cf_6\}$ and $T(F_2) = \{abc, bc^2, abd, acd, bcd, cd^2\}$. We pick abd which is reducible by bdf_4 but again we can replace this product by bf_5. After several steps we find $F_2 = \{cf_5, df_7, bf_5, f_2, cf_6\}$

$$A_2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 \end{bmatrix},$$

$$\tilde{A}_2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & -1 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0 & 0 & -1 & -1 & 1 & -1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & -1 & 0 \end{bmatrix}$$

$\tilde{F}_2 = \{f_9 = acd + bcd + c^2d + cd^2, f_{10} = b^2d + 2bd^2 + d^3, \; f_{11} = abd + bcd - bd^2 - d^3, f_{12} = abc - bcd - c^2d + bd^2 - cd^2 + d^3, f_{13} = bc^2 + c^2d - bd^2 - d^3\}$ and $G = \{f_4, f_7, f_{13}\}$.

In the next step we have $L_3 = \{(1, f_1), (bcd, f_4), (c^2, f_7), (b, f_{13})\}$, and we call recursively the function Simplify: $Simplify(bcd, f_4) = Simplify(cd, f_6) = Simplify(d, f_{12}) = (d, f_{12})$. We have $F_3 = \{f_1, df_{12}, c^2f_7, bf_{13}\}$. Notice that c^2f_7 cannot be simplified, but very often we have only a polynomial multiplied by a variable. After several steps in Symbollic Preprocessing we find $F_3 = \{f_1, df_{12}, c^2f_7, bf_{13}, df_{13}, df_{10}\}$ and $\tilde{F}_3 = \{f_{15} = c^2b^2 - c^2d^2 + 2bd^3 + 2d^4, f_{16} = abcd - 1, f_{17} = -bcd^2 - c^2d^2 + bd^3 - cd^3 + d^4 + 1, f_{18} = c^2bd + c^2d^2 - bd^3 - d^4, f_{19} = b^2d^2 + 2bd^3 + d^4\}$. Note that the rank of F_3 is only 5. This means that there is a reduction to zero.

2.7. Conclusion

So we have transformed the degree of freedom in the Buchberger algorithm into strategies for efficiently solving linear algebra systems. This is easier because we have constructed a matrix A (the number of rows in A is a little overestimated by the symbolic prediction) and we can decide to begin the reduction of one row before another with a "*good reason*". For integer coefficients it is a major advantage to be able to apply an iterative algorithm on the whole matrix (p-adic method). Some negative aspects are that the matrix A is singular and that A is often huge. A good compression algorithm for the matrix reduces the storage requirements by a factor greater than 10 (see Section 3.4).

3. Solving sparse linear systems

Compared with naive linear algebra approach we have reduced dramatically the size of matrices which are involved. Despite this reduction, the matrices that we have to solve during the program execution are very big sparse matrices (see p. 22): say $50\,000 \times 50\,000$ to give an idea (the record is $750\,000 \times 750\,000$ for a very sparse matrix). To give a comprehensive review of all useful linear techniques is far beyond the scope of this paper and we give only some references to the original papers. It should be observed that we have to solve *sparse* matrices in a *computer algebra* system. It is therefore not surprising that we have to merge techniques coming from sparse linear algebra (possibly designed initially for floating point coefficients) and techniques coming from computer algebra (for big integer computations for instance). But first we establish the link with the main algorithm:

3.1. Solving a matrix and reduction to row echelon

In the main algorithm we have to *reduce* (find a basis of the image of the correspon-fing linear map) sparse matrices which are singulars and not square. On the other hand, linear algebra techniques are often described for *solving* linear systems of equations $Ax=b$. One way to connect the two approaches is to first extract a square sub matrix and to put the remaining columns into the right-hand side. For instance if we want to reduce:

$$
A = \begin{array}{c} \\ f_1 \\ f_2 \\ f_3 \end{array}
\begin{array}{cccc} a^2 & ab & b^2 & bc \\ \end{array}
\left(\begin{array}{cccc} 1 & 2 & 1 & 1 \\ 1 & 2 & 1 & 2 \\ 1 & 2 & 2 & 2 \end{array} \right)
$$

We first try to find the rank of the matrix using a fast algorithm mod p and we find that the second column is deffective. So we extract a square matrix and a right-hand side:

$$
\tilde{A} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 2 \\ 1 & 2 & 2 \end{bmatrix} \quad r = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}
$$

Hence the system of equations can be rewritten:

$$
\tilde{A} \begin{bmatrix} a^2 \\ b^2 \\ bc \end{bmatrix} = -(ab)r \quad or \quad \begin{bmatrix} a^2 \\ b^2 \\ bc \end{bmatrix} = -(ab)(\tilde{A}^{-1}r)
$$

so we have to *solve* the linear system

$$
\begin{bmatrix} x_0 \\ z_0 \\ t_0 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix}
$$

and the reduced form of the matrix

$$
\begin{array}{c} \\ f_1 \\ f_2 \\ f_3 \end{array}
\begin{array}{cccc} a^2 & ab & b^2 & bc \\ \end{array}
\left(\begin{array}{cccc} 1 & x_0 & 0 & 0 \\ 1 & z_0 & 1 & 0 \\ 1 & t_0 & 0 & 1 \end{array} \right)
=
\begin{array}{c} \\ f_1 \\ f_2 \\ f_3 \end{array}
\begin{array}{cccc} a^2 & ab & b^2 & bc \\ \end{array}
\left(\begin{array}{cccc} 1 & 2 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{array} \right)
$$

3.2. Solving sparse linear systems

Solving sparse linear systems has been studied for a very long time for numerical computations, there are mainly two types of methods: iterative methods (computing successively $y_{i+1} = Ay_i$ if A is an $n \times n$ matrix) and decomposition methods. The methods which are useful in our context are the following:

3.2.1. Iterative solvers

After $O(n)$ steps we obtain an exact solution (up to rounding errors). This is the case of conjugate gradient method or the Lanczos algorithm [39] since after n steps

the result is exact. Another well known iterative algorithm is the Wiedemann algorithm [29,48], which uses the efficient Berlekamp and Massey algorithm [38] to recover the solution. Note that there exists a more efficient version of the Wiedemann algorithm: the Wiedemann algorithm by blocks (but we have not implemented this version yet).

The key operation in these algorithms is the multiplication $A*y$. It is very easy to implement this operation efficiently to take advantage of the sparsity of A and to obtain a complexity of $O(|A|)$ for computing $A*y$ where $|A|$ is the number of non zero elements in A; thus the global complexity for solving the system is $O(n|A|)$ instead of $O(n^3)$.

When the matrix has a regular pattern it is possible to apply even more efficient techniques (for Toeplitz matrices for instance [5]). These methods cannot be applied in our case since the pattern of the generated matrices is not regular enough. A significant drawback of these methods is that there is no speedup if we try to solve simultaneously several linear systems with the same left-hand side.

3.2.2. Factorization methods

The classical LU decomposition tries to decompose the input matrix A into a product LU where L (resp. U) is a lower (resp. upper) triangular matrix. For *sparse* LU decomposition [15,22,23,42,43] there is an additional constraint: the number of non-zero elements in L and U minus the number of non-zero elements in A must be as smallest as possible (this number is called the number of "fill-ins"). The sparse LU decomposition starts with a symbolic preprocessing very similar to ours. It aims is to avoid to spend time for computing coefficients the vanishing of which is predictable. The interesting feature of this method is that both preprocessing may be done simultaneously.

A large number of implementations for these methods are available (mainly in C/C++ or Fortran) [12] or even in Matlab. Unfortunately, these programs and some algorithms are not very robust: very often the input matrix A must be non-singular or square or positive definite. On the other hand, parallel algorithms and parallel implementations exist [2,14,28,41,47]. We have modified the smms [1] program in order to work with modulo p coefficients in order to evaluate the costs of different algorithms.

Solving large linear equations modulo a small prime number has been studied in cryptology [31,39] for factoring integers or computing discrete logarithms (very often $p=2$ in that case). These authors use a combination of structured Gaussian elimination and other iterative algorithms. In the current implementation of the F_4 algorithm we use by default the structured Gaussian elimination. Note that in an ideal implementation the program should first analyse the shape of the matrix and then decide which algorithm should be applied.

3.3. Computer Algebra methods for solving linear systems with integer coefficients

Very specific methods have been developed in Computer Algebra for solving linear equations when the coefficients are integers. First we recall that the Bareiss algorithm is better than the classical Gauss algorithm but is very inefficient for big systems.

The best way of solving linear systems $Ax = b$ with integer coefficients is to use p-adic computations: we choose a prime p, and we compute $C = A^{-1} \bmod p$ (very often a sparse LU decomposition is more appropriate). Then we define

$$b^{(0)} = b, \qquad y^{(m)} = Cb^{(m)} \bmod p,$$

$$b^{(m+1)} = \frac{b^{(m)} - Ay^{(m)}}{p}, \qquad x^{(m)} = \sum_{j=0}^{m-1} y^{(j)} p^j.$$

Theorem 3.1 (Dixon [13]).

- $Ax^{(m)} - b = 0 \bmod p^m$
- *If B is upper bound of the coefficients of A and b it is sufficient to compute $x^{(m)}$ for $m \geq 2n \log_p(nB)$ and then to apply the extended Euclidean algorithm to $x^{(m)}$ and p^m to recover the value in $\mathbf{Q}$.*

In fact the iteration may be stopped when $x^{(m)}$ becomes stable (see [40] for multi-modular methods which may be generalized to p-adic method).

The global complexity [30] of this algorithms is

$O(n^5 \log(nB)^2)$	Bareiss method
$O(n^3(n + \log(nB))\log(nB))$	multi modular method
$O(n^3 \log(nB)^2)$	p-adic method

Note that the p-adic method is also an iterative algorithm (in fact this is a Newton algorithm) and that we have previously noticed that this kind of algorithm is less efficient for solving simultaneous systems $Ax = b_1, \ldots, Ax_k = b_k$. If $k \geq n$ it is probably better to use a multi-modular approach: we compute $A^{-1}b \bmod p_i$ for different primes p_i and we apply the Chinese remainder theorem to find the solution modulo $\prod_i p_i$.

3.4. Matrix compression

When the matrices are big it is necessary to adopt fairly complicated storage schemes to compress the matrix in memory: consider a $(5 \times 10^4) \times (5 \times 10^4)$ matrix with 10% non-zero elements (this is the case in Cyclic 9 for instance); even if only one byte is allocated for each coefficients (an optimistic assumption since some coefficients have hundred digits); the matrix requires 250×10^6 bytes to be stored! In our implementations we do not duplicate coefficients (relying on the fact that some rows in the matrix are just term multiples of others); thus we have only to compress the position of the non-zero elements; we have experimented the following techniques:

(i) No compression: inefficient both for CPU time and memory usage.

(ii) Bitmap compression: each row of the matrix contains zero and non-zero elements:

$$x \quad 0 \quad 0 \quad x \quad x \quad 0 \quad 0 \quad 0 \quad x \quad \ldots$$

$j_1, j_2, \ldots$ denote the indexes of the non-zero elements (here $j_1 = 1$, $j_2 = 4, \ldots$), then $\sum_k 2^{j_k - 1}$ is the bitmap compression (for the example $1 + 2^3 + 2^8 = 281$). This method is efficient but consumes a substantial amount of memory. This is the prefered way to implement the compression when the coefficients are in $GF(2)$

(iii) Distributed representation: the row is represented by the array

$$\boxed{j_1 \;|\; j_2 - j_1 \;|\; j_3 - j_2}\; \cdots$$

using bytes when $j_k - j_{k-1} < 255$ (this occurs very frequently). This method is most efficient in memory than the previous one and a little slower (10%).

(iv) Apply a standard compression tool (gzip for instance) to one of the previous representation. Very efficient in memory but rather slow.

For the moment our prefered methods are (ii) and (iii) (depending on the ground field), but the compression algorithms should be seriously improved in future versions.

4. Experiments

The quality of the computer implementation of Gröbner bases computations may have a dramatic effect on their performance. The theoretical results on complexity (even for homogeneous systems, the complexity is d^{3n} in most cases and 2^{2^n} in some very pathological cases) cannot throw light on the practical behavior of effective implementation. Thus, while "paper and pencil" analysis of polynomial solving algorithms are important, they are not enough. Our view is that studying and using programs which implement these algorithms is an essential component in a good understanding of this area of Computer Algebra. To this end, we provide some experiments and comparison with similar programs. This section should be considered as the validation of our algorithm. Thus it plays the same role as its proof for theorems.

4.1. Methodology

Empirical analysis means that we have to pay particular attention to the development of methodologies in the field of benchmark for linear system solving. We adopt the following points:

1. We compare the new algorithm with *state of the art* Gröbner bases implementation (namely Magma [10], PoSSo/Frisco Gröbner engine [46], Macaulay 2 [26,44], Singular [27], Asir [45], Cocoa, Axiom, Maple [11] and Gb [17]). It is also crucial to compare the implementation of the new algorithm with the Buchberger algorithm implemented by *the same person* (in this case the author). In our opinion it is also important to compare low-level implementation of the algorithms to avoid parasite interactions with a high level language/compiler. F_4 has been implemented in C and most of the competitors are implemented in C/C++.

2. The list of examples is also a crucial issue: the examples can be *easily accessed* [16] (the web site contains pointers to the Frisco test suite). The list is composed of classical benchmarks (Cyclic n, Katsura n) but also of industrial examples (coming from signal theory, robotics). We have removed from the list all the *toy examples*

since nothing can be concluded with them. Of course the toy concept is relative to current computers. For us, an example is toy if it takes less than 1 sec on a PC.

3. This section contains two timing tables: the first one corresponds to modular computations, the second one corresponds to big integers computations. All the computations are done several times on equivalent computers to prevent as much as possible interactions with other programs. For each timing the program was run several times. This was necessary to eliminate fluctuations in the measurements stemming from some other programs running on the same computer. Of course the timings are rounded.

4. We rigorously use the last version of all the programs and use an appropriate computer to execute them.

5. An even more convincing proof of the efficiency of a new algorithm is to solve previously untractable problems. So we should test the algorithm on very difficult problems. In our case, F_4 solves Cyclic8 and Cyclic9.

6. The same strategy is used for all the programs. For instance if we homogenize the input polynomials for one program we try the same strategy with all other programs. This is the reason why we give the timing for cyclic n and homogeneous cyclic n. Some systems (Singular for instance) allow the user to customize the internal strategy; we try several parameters and we retain the best method.

7. The outputs of the different programs are checked to be equal.

8. Last, and it is the more difficult, we compare the algorithm with other methods: triangular sets (Moreno/Aubry), homotopies (J. Verschelde), Bezoutian (Mourrain) and sparse resultant (Emiris). The task is more difficult since the outputs are very different (quasi component, floating point approximations).

8.1. Triangular methods (Wu, Wang, Lazard, Kalkebrenner). On one hand triangular methods seem to be less efficient than lexico Gröbner bases computation (the current limit is Cyclic 6 and Katsura 5) but on the other hand the quality of the output is better. So we think that these methods are useful to simplify lexicographical Gröbner bases.

8.2. Homotopy methods. J. Vershelde reports timings (Cyclic 8 on a Sparc-Server 1000 in $4h35m$) which are less efficient than Gröbner bases techniques. It is difficult to handle over constrained systems.

8.3. Mixed volume is extremely fast to estimate the number of isolated roots (up to Cyclic' 12) but with our experience it is not so efficient for other systems (and the number of solution is often over estimated).

4.2. Modular computations

Modular computations are very useful in Computer Algebra because they give a fast result (with a very high probability) and information on the number of solutions (Hilbert function).

Moreover, since big integer computations could be done by means of p-adic or multi modular arithmetics it means that the cost of an integer computation is roughly

time of modular computation $*$ size of the output coeffs

So it is also very important to have an efficient solver modulo a prime p.

The computer is a PC Pentium Pro 200 MHz with 128 Meg of RAM.

We consider only non homogeneous systems, but in the following we give $\tau = \min(\tau_{homog}, \tau_{non\ homog})$ because Singular and Macaulay are often very slow for non homogeneous system.

Note that the PoSSo library uses generic coefficient to implement Z/pZ but the other softwares implement inlined modular arithmetic. In other words the overhead of function calls is heavier in the PoSSo library. With a better implementation one can probably divide timings by a factor between 2 and 4.
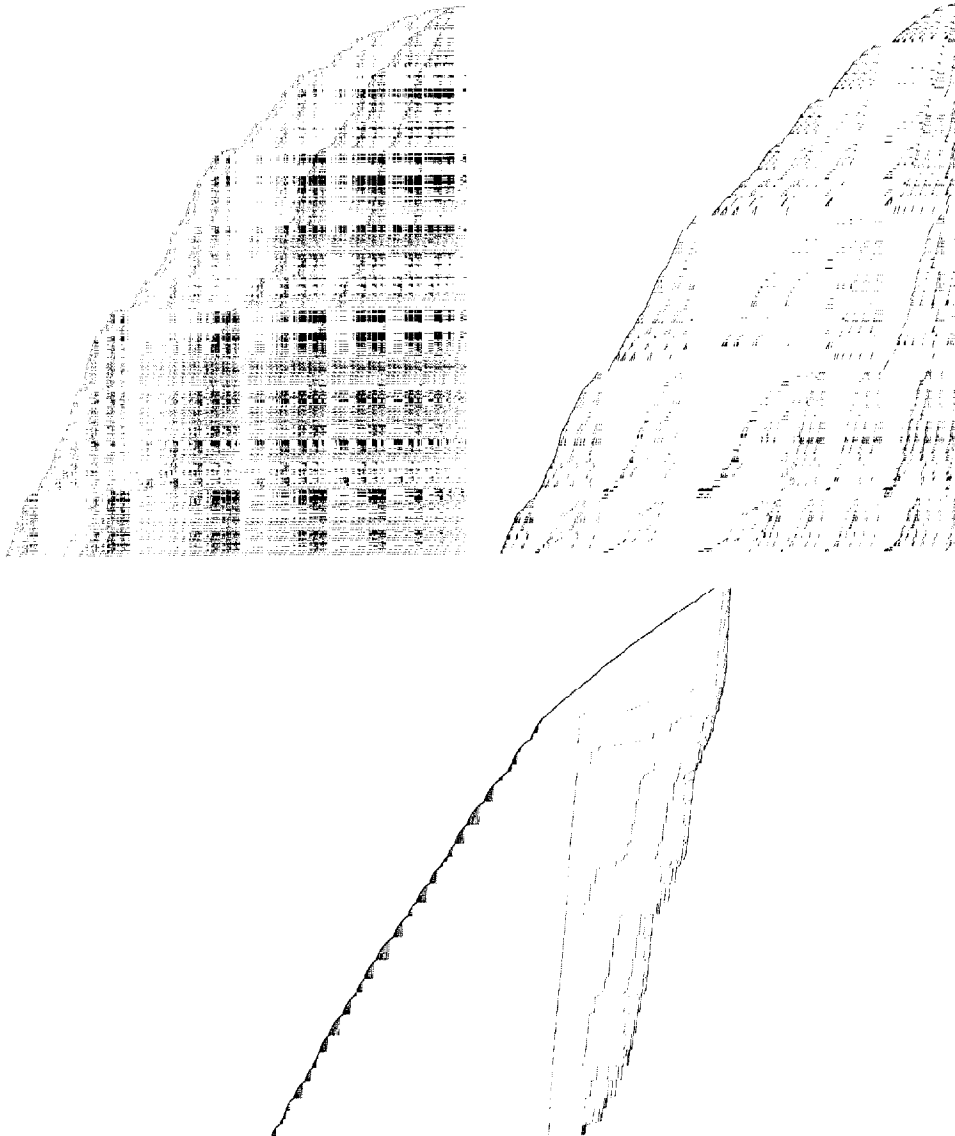
Table 1
Academic examples mod p

	FGb	Gb	PoSSo	Singular	Mac 2	Asir
K_7	0.8 s	3.2 s	46 s	10 s	12 s	56 s
K_8	4.1 s	29 s	9 min 58 s	1 min 47 s	1 min 55 s	
K_9	30 s	3 min 48 s	1 h25 min	17 min 41 s	17 min 11 s	
K_{10}	4 min 13 s	36 min 23 s		3 h6 min		∞
K_{11}	30 min 29 s		∞	∞	∞	∞
Cyclic 7	4.6 s	1 min 15 s	9 min 3 s	2 min 34 s	2 min 0 s	6 min 51 s
Hom C_7	5.2 s	1 min 02 s				
Cyclic 8	1 min 55 s	26 min 17 s	4 h 38 min	1 h 56 min	1 h 33 min	3 h 54 min
Hom C_8	184.4 s	39 min 17 s				
Cyclic 9	4 h 32 min	∞	∞	∞	∞	∞
Hom C_9	11 h 10 min[a]	∞	∞	∞	∞	∞

[a] 11 h 9 min 40 s on a 500 MHz Alpha workstation with 1 Go op RAM.

We give the shape of some generated matrices. The matrices can have a very different structure and the number of non-zero elements varies greatly.

Upper left figure:

From $Cyclic_7$

Matrix A_9 475 × 786 (13.8% non-zero).

Upper right figure:

application example: from $f_{8,5,5}$ [19] matrix A_5 1425 × 2561 (0.47% non-zero).

Lower left figure:

Example: engineering problem (Nuclear) matrix A_7 1557 × 3838 (0.2% non-zero).

The conclusion of Table 1 is that *the old Gb is still faster than other systems, and that* FGb ≫ Gb.

The following figure shows the performance results for the homogenization strategy and the affine strategy (the quotient "time to solve homog Cyclic n" divided "time to solve non-homog Cyclic n" is plotted).

Comparison Homegeneous/ Non homogeneous
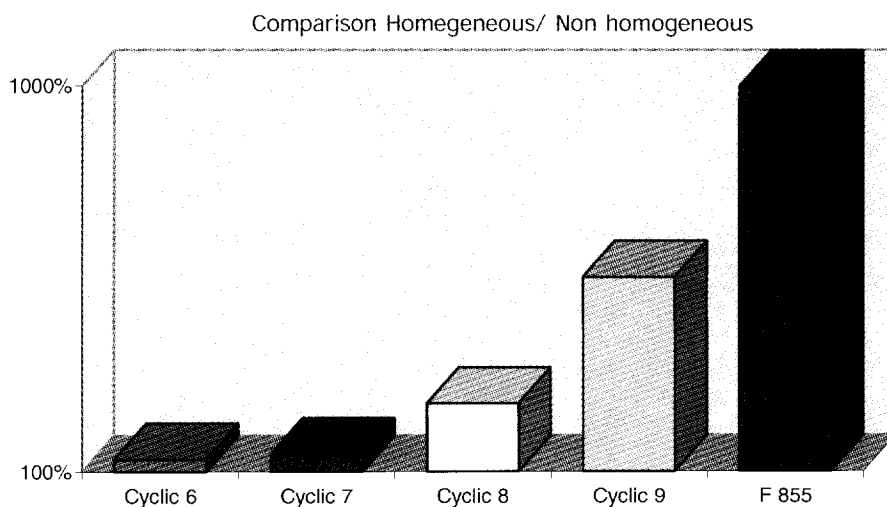
We conclude that for small systems the difference is small but for big systems like Cyclic 9 there is a huge difference (you add several components of dimension 3 and 4). For Cyclic 9 the algorithm generates a 292051×317850 matrix and it is 3.2 times slower!

Cyclic n - Modulo p - DRL Groebner basis
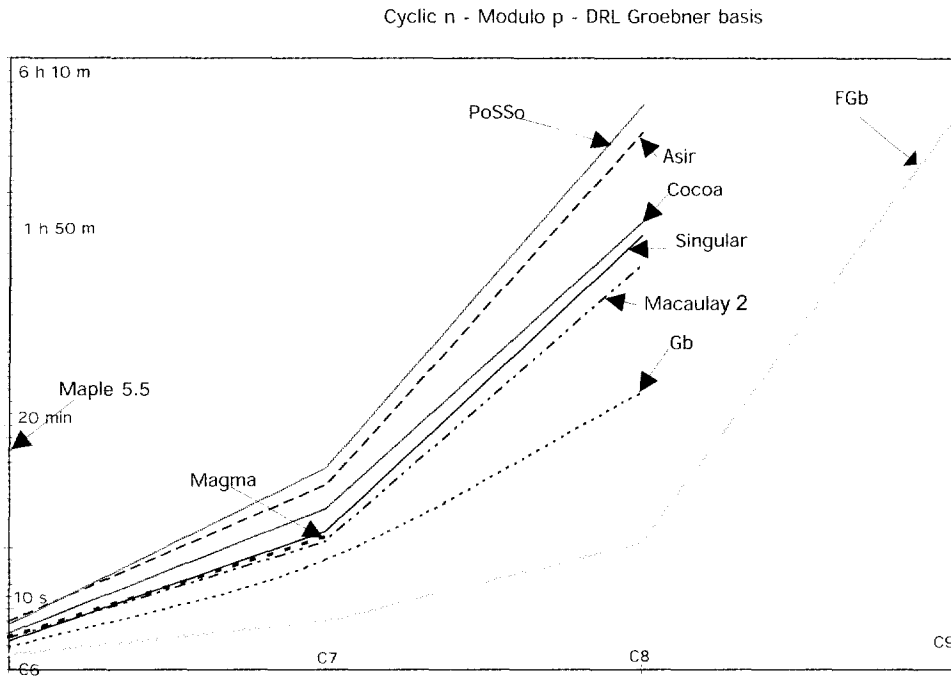
4.3. Integers

In Table 2, we have included a special version of the the classical PoSSo Gröbner engine called "Rouil", this version has been optimized by Fabrice Rouillier.[2] p > In the table we remove the Singular entry because Singular and Macaulay2 seem very close and very slow for big integers.

Size is the size in digits of the biggest coefficient in the output.

---

[2] The algorithm uses a prime p to avoid syzygies. Then the algorithm checks that the result was correct. At the present state the implementation sometimes does not detect bad primes.

Table 2
Academic examples $\mathbf{Z}$

|	FGb	Gb	Singular	Rouil	Asir	Mac 2	Magma	Size
K_7	2.9 s	42 s		50 s	3 min 28 s	10 min 56 s		53
K_8	23 s	10 min 21 s		8 min 23 s	8 min 37 s	2 h 17 min		102
K_9	3 min 3 s	5 h 35 min		1 h 31 min		∞		133
K_{10}	31 min 24 s	∞		∞	∞	∞		192
Cylic 6	0.3 s	3.2 s	5.3 s	5.4 s	8 s	19.5 s	2.6 s	96
Hom C_7	54.2 s	1 h 32 min	10 h 35 min	>25 min[a]	2 h 50 min		2202 s	96
Cylic 7	39.7 s	5 h 17 min[b]	∞	∞	>2 h[c]	∞	=	96
Cylic 8	24 min 4 s	∞	∞	∞	∞	∞	∞	202
Cylic 9	18 days[d]	∞	∞	∞	∞	∞	∞	800

[a] 134 Mbyte of data when stopped.
[b] Estimated from an original time of 24 h 26 min 40 s on Sparc 10 with Lazard s algorithm.
[c] 162 Mbyte of data when stopped.
[d] The size of the result in binary is 1 660 358 088 bytes. Run on 512 Meg RAM PC.

We observe that Magma is very efficient for big integers (in fact the Magma version for alpha workstations is even 7 times faster).
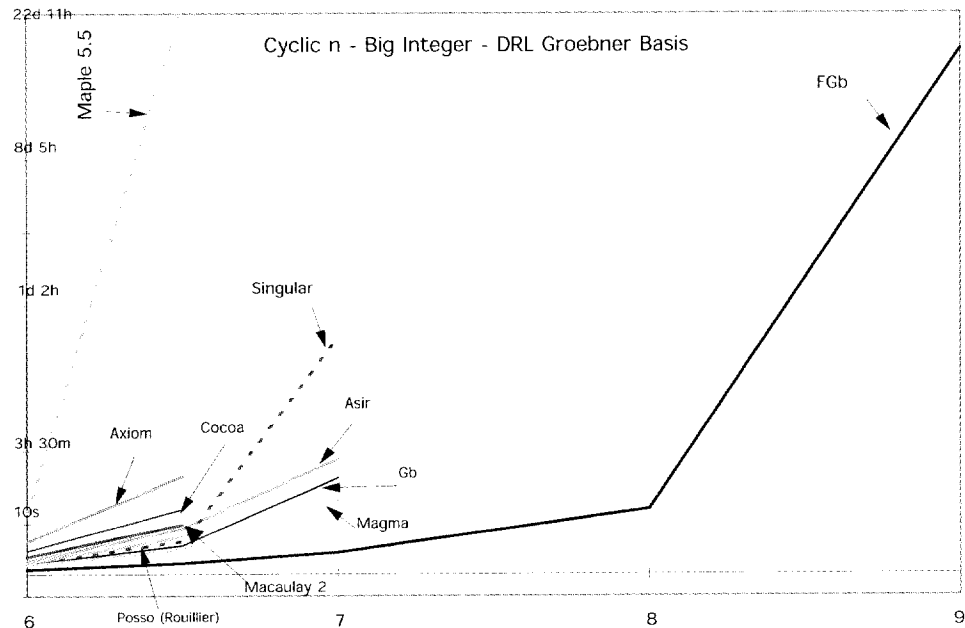
Cyclic-9 for big integers is an example of *huge* computation; we use:

- 3 Processors PPro 200 MHz 512 Meg RAM + 1 Alpha 400 MHz 570 Meg.
- Total sequential CPU time: 18 days.
- Size of the file of coefficients in the output (binary): 1660 Meg.
- The result contains: 1344 polynomials with 1000 monomials and 800 digits numbers !

This success is also a failure in some sense: the size of the output is so big that we cannot do anything with this result. That is to say we are now near to the intrinsic complexity of Gröbner bases. On the other side, the output is very big because the co-efficients are big and floating point computation would not suffer from this exponential growth.

We conclude that for all these examples FGb≫Gb and that it faster than other systems.

5. Final remarks

The conclusion is that F_4 is at least one order of magnitude faster than all previously implemented algorithms.

We recall the main features of this algorithm are:

- Replace all the previous strategies (sugar([25]),...) by algorithms for (sparse) linear algebra. It explains why the usual strategies in Buchberger algorithm could not be optimal.
- Faster for all kind of arithmetic (modular computation, integers, "generic" computation)
- In some sense it is as fast as possible for big integers coefficients or coefficients with parameters ($k(y_1, \ldots, y_r)$) because the practical complexity is almost linear in the size of the output coefficients: in the case of homogeneous polynomials the complexity of F_4 is $s^{1+o(1)} + M$ where s is the size of othe result and M is the time needed for modular computation which is generally much smaller.
- For homogeneous systems the algorithm generate reduction to zero or non-zero poly-nomials (completely reduced) which are all in the final Gröbner basis. So the algo-rithm does not generate "parasite" intermediate data.
- Very good experimental behavior for non-homogeneous systems (several times faster than the corresponding homogeneous system).
- Parallel versions of the algorithm can be implemented (we have done a first imple-mentation).
- The algorithm is easier to implement (no polynomial arithmetic is required, do not need an efficient power product (exponents) implementation).
- It can solve previously intractable problems: we are now able to compute easily three new records: Cyclic9 for modular coefficients (4h30), Cyclic8 for big integers coefficients (25 mins) and the very challenging Cyclic9 for bignum coefficients (18 days of CPU time).

A lot of work remains to do in the linear algebra part to apply less naive algorithms to the generated matrices and to compress more efficiently those matrices. Considerable works must also be done to compare the algorithm with different possible strategies (sugar and homogenizing, multi-weight in the case of multi-homogeneous ideals might reduce the size ot the matrices (as suggest by D. Lazard and one anonymous referee)). How to use the symmetry of the problems to handle more efficiently the matrices is also

an open problem. Even if the algorithm presented in this paper is heavily connected with the Buchberger algoritm (use the same criteria for useless pairs), we think that an interesting work would be to use the technique of Mandache [36,37] to check that F_4 *is not* a Buchberger algorithm in the sense that the Buchberger algorithm cannot simulate the new algorithm for any strategy. When the normal strategy is used, we can plot the function $d \rightarrow deg(first(P_d))$; we obtain an increasing function for homogeneous systems but in the affine case we obtain different curves:

Cyclic 8 – Total degree of critical pairs

E1 50 – Total degree of critical pairs

An open issue is to understand deeply the shape of this curves.

Acknowledgements

References

[1] F.L. Alvarado, The Sparse Matrix Manipulation System User and Reference Manual, University of Wisconsin, Wisconsin, May 1993.

[2] F.L. Alvarado, A. Pothen, R. Schreiber, Highly parallel sparse triangular solution, Preprint.

[3] G. Attardi, C. Traverso, Homogeneous parallel algorithm, J. Symbolic Comput. (1996).

[4] T. Becker, V. Weispfenning, Groebner Bases, a Computationnal Approach to Commutative Algebra, Graduate Texts in Mathematics, Springer, Berlin, 1993.

[5] D. Bini, V. Pan, Polynomial and Matrix Computations, Progress in Theoretical Computer Science, vol. 1, Fundamental Algorithms, Birkäuser, Basel, 1994.

[6] B. Buchberger, Ein Algorithmus zum Auffinden der Basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal, PhD thesis, Innsbruck, 1965.

[7] B. Buchberger, An algorithmical criterion for the solvability of algebraic systems, Aequationes Mathematicae 4 (3) (1970) 374–383 (in German).

[8] B.A. Buchberger, Criterion for detecting unnecessary reductions in the construction of Gröbner basis, in: Proc. EUROSAM 79, Lecture Notes in Computer Science, vol. 72, Springer, Berlin, 1979, pp. 3–21.

[9] B. Buchberger, Gröbner bases: an algorithmic method in polynomial ideal theory, in: Reidel (Ed.), Recent Trends in Multidimensional System Theory, Bose, 1985.

[10] J. Cannon, The Magma Computational Algebra System 2.20-7, February 1998. http://www.maths.usyd.edu.au:8000/u/magma/.

[11] B. Char, K. Geddes, G. Gonnet, B. Leong, M. Monagan, S. Watt, Maple V Library Reference Manual, Springer, Berlin, 1991, Third Printing, 1993.

[12] T. Davis, Users' Guide for the Unsymmetric-pattern MultiFrontal Package, Computer and Information Sciences Department, University of Florida, January 1995. Technical Report TR-95-004.

[13] J.D. Dixon, Exact solution of linear equations using p-adic expansions, Numer. Math. 40 (1982) 137–141.

[14] J. Dongarra, A. Lumsdaine, X. Niu, R. Pozo, K. Remington, A sparse matrix library in C+— for high performance architectures, Preprint.

[15] I.S. Duff, J.K. Reid, The multifrontal solution of unsymmetric sets of linear equations, SIAM J. Sci. Statist. Comput. 5 (3) (1984) 633–641.

[16] J.C. Faugère, Benchmarks for polynomial solvers, Available on the WEB http://posso.lip6.fr/~jcf/benchs.html.

[17] J.C. Faugère, On line documentation of Gb, Available on the WEB *http://posso.lip6.fr/~jcf*.

[18] J.C. Faugère, Computing Gröbner Basis without reduction to zero (F5), Technical report, LIP6 report, 1998, in preparation.

[19] J.C. Faugère, F. Moreau de Saint-Martin, F. Rouillier, Design of nonseparable bidimensional wavelets and filter banks using Gröbner bases techniques, IEEE SP Trans. Signal Processing 46 (4) 1998. Special Issue on Theory and Applications of Filter Banks and Wavelets.

[20] A. Fronville, Singularité résiduelle et Problème du centre, PhD thesis, Université Paris 6, Décember 1996.

[21] R. Gebauer, H.M. Möller, On an installation of Buchberger's algorithm, J. Symbolic Comput. 6 (2,3) (1988) 275–286.

[22] A. George, J.W.H. Liu, Computer Solution of Large Sparse Positive Definite Systems, Prentice-Hall, Englewood Cliffs, NJ, 1981.

[23] A. George, J.R. Gilbert, J.W.H. Liu, Graph Theory and Sparse Matrix Computation, Springer, New York, 1993.

[24] V.P. Gerdt, Involutive Polynomial Bases, in: PoSSo on Software Paris, F, 1995.

[25] A. Giovini, T. Mora, G. Niesi, L. Robbiano, C. Traverso, One sugar cube, please, or Selection strategies in the Buchberger Algorithm, in: S.M. Watt (Ed.), Proc. 1991 Internat. Symp. on Symbolic and Algebraic Computation, ISSAC' 91, ACM, New York, 1991.

[26] D. Grayson, M. Stillman, Macaulay 2 User Manual, 1997. http://www.math.uiuc.edu/Macaulay2/.

[27] G.-M. Greuel, G. Pfister, H. Schoenemann, SINGULAR 1.0.2, July 1997. http://www.mathematik.unikl.de/~zca/Singular/Welcome.html.

[28] M.T. Heath, P. Raghavan, Distributed solution of sparse linear systems, Preprint.

[29] E. Kaltofen, On Wiedemann's method of solving sparse linear systems, Lecture Notes in Computer Science, vol. 539, Springer, Berlin, 1991, pp. 29–38, AAECC-9.

[30] D.E. Knuth, The Art of Computer Programming, vol. 2, 2 ed., Addison-Wesley, Reading, MA, 1981.

[31] B.A. LaMacchia, A.M. Odlyzko, Solving large sparse linear systems over finite fields, Lecture Notes in Computer Science, vol. 537, Springer, Berlin, 1991. Advances in Cryptology, CRYPTO'90.

[32] D. Lazard, Systems of algebraic equations, in: EUROSAM 79, 1979, 88–94.

[33] D. Lazard, Resolution des systemes d'equations algebriques, Theoret. Comput. Sci. 15 (1981) 77–110.

[34] D. Lazard, Gaussian elimination and resolution of systems of algebraic equations, in: Proc. EUROCAL 83, Lecture Notes in Computer Science, vol. 162, Springer, Berlin, 1983, pp. 146–157.

[35] H. Lombardi, Un nouvel algorithme de calcul d'une Base de Gröbner, Technical report, Equipe de Mathématiques de Besançon, Janvier 1998.

[36] A.M. Mandache, The Gröbner basis algorithm and subresultant theory, in: Proc. ISSAC'94, Oxford, UK, ACM, New York, 1994, pp. 123–128.

[37] A.M. Mandache, A note on the relationship between involutive bases and Gröbner bases, Mega 96, 1996.

[38] J.L. Massey, Shift-register synthesis and bch decoding, IEEE Trans. Inform. Theory 15 (1) (1969) 122–127.

[39] P. Montgomery, A block Lanczos algorithm for finding dependencies over GF(2), Lecture Notes in Computer Science, Springer, Berlin, 1995, pp. 106–120.

[40] S.R. Petrick (Ed.), Proc. 2nd Symp. Symbolic and Algebraic Manipulation, ACM Press, New York, 1971.

[41] B.W. Peyton, A. Pothen, X. Yuan, Partitioning a choral graph into transitive subgraphs for parallel sparse triangular solution, Preprint.

[42] J.K. Reid, Large Sparse Sets of Linear Equations, Academic Press, London, 1971.

[43] D. Rose, R.A. Willoughby, Sparse Matrices and their Applications, Plenum Press, New York, 1972.

[44] M. Stillman, D. Bayer, Macaulay User Manual, 1989, Available via anonymous ftp on `zariski.harvard.edu`.

[45] T. Takeshima, Risa/Asir. Fujitsu Laboratories Limited, 1996, Version 950831, Available from ftp://endeavor.fujitsu.co.jp/pub/isis/asir.

[46] C. Traverso, Groebner Engine. Rel. 2.0, 1997, `http://janet.dm.unipi.it/posso\_demo.html`.

[47] A.F. Van der Stappen, R.H. Bisseling, J.G.G. Van de Vorst, Parallel sparse LU decomposition on a mesh network of transputers, SIAM J. Matrix Anal. Appl. 14 (3) (1993) 853–879.

[48] D. Wiedemann, Solving sparse linear equation over finite fields, IEEE Trans. Inform. Theory IT-32 (1) (1986) 54–62.