UNIVERSITÄT KAISERSLAUTERN
Zentrum für Computeralgebra

REPORTS ON COMPUTER ALGEBRA
NO. 18

Monomial Representations for Gröbner Bases Computations

by

O. Bachmann and H. Schönemann

January 1998

The Zentrum für Computeralgebra (Centre for Computer Algebra) at the University of Kaiserslautern was founded in June 1993 by the Ministerium für Wissenschaft und Weiterbildung in Rheinland-Pfalz (Ministry of Science and Education of the state of Rheinland-Pfalz). The centre is a scientific institution of the departments of **Mathematics, Computer Science, and Electrical Engineering** at the University of Kaiserslautern.
The goals of the centre are to advance and to support the use of Computer Algebra in industry, research, and teaching. More concrete goals of the centre include

- the development, integration, and use of software for Computer Algebra
- the development of curricula in Computer Algebra under special consideration of interdisciplinary aspects
- the realisation of seminars about Computer Algebra
- the cooperation with other centres and institutions which have similar goals

The present coordinator of the Reports on Computer Algebra is:
Olaf Bachmann (email: `obachman@mathematik.uni-kl.de`)

Zentrum für Computeralgebra
c/o Prof. Dr. G.-M. Greuel, FB Mathematik
Erwin-Schrödinger-Strasse
D-67663 Kaiserslautern; Germany
Phone: 49 - 631/205-2850 Fax: 49 - 631/205-5052
email: `greuel@mathematik.uni-kl.de`
URL: `http://www.mathematik.uni-kl.de/~zca/`

Monomial Representations for Gröbner Bases Computations*

Olaf Bachmann Hans Schönemann
Centre for Computer Algebra
Department of Mathematics
University of Kaiserslautern
Kaiserslautern, Germany
[obachman,hannes]@mathematik.uni-kl.de

Abstract

Monomial representations and operations for Gröbner bases computations are investigated from an implementation point of view. The technique of vectorized monomial operations is introduced and it is shown how it expedites computations of Gröbner bases. Furthermore, a rank-based monomial representation and comparison technique is examined and it is concluded that this technique does not yield an additional speedup over vectorized comparisons. Extensive benchmark tests with the Computer Algebra System SINGULAR are used to evaluate these concepts.

1 Introduction

The method of Gröbner bases (GB) (see, for example, [9] for an introduction) is undoubtly one of the most important and prominent success stories of the field of Computer Algebra. Starting in the 1960's, an unsolved problem has developed into an essential computational tool with a great variety of applications and more and more powerful implementations. The heart of the GB method are computations of Gröbner or Standard bases with the Buchberger algorithm or descended variants thereof (we call such computations GB computations, for short).

Unfortunately, the Buchberger algorithm and its variants have a worst case exponential time and space complexity [2]. Consequently, GB computations have limited practicality and tend to tremendously long running times and consumptions of huge amounts of memory. For these reasons, a large amount of work has been done to try to improve the (time and space) efficiency of GB computations (see, for example, [7, 8, 10]). Most of this work resulted in algorithmic improvements which led to more manageable computations for many classes of problems. Our approach to improving the efficiency of GB computations is somewhat different: instead of trying to improve algorithmic aspects of GB computations we take the algorithms more or less as they are and study efficient ways to implement them.

Analyzing GB computations from an implementation point of view leads to many interesting questions. For example: What is the best way to implement coefficient arithmetic? How should polynomials and monomials be represented and their operations be implemented? How should memory management be realized?

*Work reported herein has been supported in part by the Deutsche Forschungsgemeinschaft (DFG).

Apart from some exceptions (for example, [1]), these questions have not yet received a lot of systematic research and literature attention. Nevertheless, work in this direction can be very fruitful, as we report in this article. We concentrate on monomials – the most frequently used primitive data in GB computations – and show that improvements in the representation of and operations on monomials can lead to significant efficiency gains of GB computations.

In Section 2 we discuss the basics of monomial operations and representations. In Section 3 we investigate vectorized monomial operations, which handle more than one exponent at a time. In section 4 we examine a rank-based representation of monomials, which encodes the exponent vector as one single integer for the purpose of expediting monomial comparisons. In both of these sections, we give extended sets of timings which show the effects of the respective concepts. Finally, we summarize our results in Section 5 and give more details about our used benchmark examples in the Appendix.

We used the Computer Algebra System (CAS) SINGULAR [11] to implement, examine and benchmark the techniques discussed in this paper. SINGULAR is CAS for polynomial computations with special emphasize on the needs of Commutative Algebra, Algebraic Geometry, and Singularity Theory. It features, among others, one of the fastest and most general implementation of GB computations and was therefore an ideal test-bed for our experiments.

2 Basic monomial operations and representations

Given an integer $n > 0$ we define the set of exponent vectors M_n by $\{\alpha = (\alpha_1, \dots, \alpha_n) | \alpha \in \mathbf{N}^n\}$. Notice that monomials usually denote terms of the form $c x_1^{\alpha_1} \dots x_n^{\alpha_n}$. However, in this paper we do only consider the exponent vector of a monomial and shall therefore use the words exponent vector and monomial interchangeably (i.e., we identify a monomial with its exponent vector).

We furthermore use Greek letters to denote monomials and the letter n to denote the a–priory given length of monomials (which is the number of variables in the corresponding polynomial ring).

Monomials play a central role in GB computations. In this section, we describe the basic monomial operations and discuss basic facts about monomial (resp. polynomial) representations for GB computations.

2.1 monomial operations

The basic monomial operations within GB computations are:

1. Computations of the degree (resp. weighted degree): the *degree* (resp. *weighted degree*) of a monomial α is the sum of the exponents $deg(\alpha) := \sum_{i=1}^{n} \alpha_i$ (resp. the weighted sum with respect to a weight vector w: $deg(\alpha) := \sum_{i=1}^{n} \alpha_i \, w_i$).

2. Test for divisibility: $\alpha | \beta \Leftrightarrow \forall i \in \{1..n\} : \alpha_i \leq \beta_i$.

3. Addition of two monomials: $\gamma := \alpha + \beta$ with $\forall i \in \{1..n\} : \gamma_i = \alpha_i + \beta_i$.

4. Comparison of two monomials with respect to a monomial ordering.

A *monomial ordering* $>$ (term ordering) on the set of monomials M_n is a total ordering on M_n which is compatible with the natural semigroup structure, i.e., $\alpha > \beta$ implies $\gamma + \alpha > \gamma + \beta$ for any $\gamma \in M_n$. A monomial ordering is a well-ordering if $(0, \ldots, 0)$ is the smallest monomial. We furthermore call an ordering negative if $(0, \ldots, 0)$ is the largest monomial.

Robbiano (cf.[15]) proved that any monomial ordering $>$ can be defined by a matrix $A \in GL(n, \mathbf{R})$: $\alpha > \beta \Leftrightarrow A\alpha >_{lex} A\beta$. Matrix–based descriptions of monomial orderings are very general, but have the disadvantage that their realization in an actual implementation is usually rather time–consuming. Therefore, they are not very widely used in practice.

Instead, the most frequently used descriptions of orderings have at most two defining conditions: a (possibly weighted) degree and a (normal or reverse) lexicographical comparison. We call such orderings *simple orderings*. The most important simple orderings (and their SINGULAR abbreviations) are:

- lexicographical (lp): used to eliminate variables and solve equations,
- (weighted) degree reverse lexicographical (dp): in general the most efficient one to compute a GB (as shown in [3]),
- (weighted) degree lexicographical (Dp),
- negative lexicographical ordering (ls),
- (weighted) negative degree reverse lexicographical (ds): also called tangent cone ordering,
- (weighted) negative degree lexicographical (Ds).

For monomials $\alpha, \beta \in M_n$ let

$$\mathrm{Lex}(\alpha,\beta) = \begin{cases} 1, \text{if } \exists i : \alpha_1 = \beta_1, \ldots, \alpha_{i-1} = \beta_{i-1}, \alpha_i > \beta_i \\ 0, \text{if } \alpha = \beta \\ -1, \text{otherwise,} \end{cases}$$

$$\mathrm{RevLex}(\alpha,\beta) = \begin{cases} 1, \text{if } \exists i : \alpha_n = \beta_n, \ldots, \alpha_{i-1} = \beta_{i-1}, \alpha_i < \beta_i \\ 0, \text{if } \alpha = \beta \\ -1, \text{otherwise,} \end{cases}$$

$$\mathrm{Deg}(\alpha,\beta) = \begin{cases} 1, \text{if } \deg(\alpha) > \deg(\beta) \\ 0, \text{if } \deg(\alpha) = \deg(\beta) \\ -1, \text{otherwise.} \end{cases}$$

Then we can define $\alpha > \beta$ for the above mentioned simple monomial orderings by:

lp:	$\mathrm{Lex}(\alpha,\beta) = 1$
ls:	$\mathrm{Lex}(\alpha,\beta) = -1$
Dp:	$\mathrm{Deg}(\alpha,\beta) = 1$, or $\mathrm{Deg}(\alpha,\beta) = 0$ and $\mathrm{Lex}(\alpha,\beta) = 1$
Ds:	$\mathrm{Deg}(\alpha,\beta) = -1$, or $\mathrm{Deg}(\alpha,\beta) = 0$ and $\mathrm{Lex}(\alpha,\beta) = 1$
dp:	$\mathrm{Deg}(\alpha,\beta) = 1$, or $\mathrm{Deg}(\alpha,\beta) = 0$ and $\mathrm{RevLex}(\alpha,\beta) = 1$
ds:	$\mathrm{Deg}(\alpha,\beta) = -1$, or $\mathrm{Deg}(\alpha,\beta) = 0$ and $\mathrm{RevLex}(\alpha,\beta) = 1$

We furthermore call a monomial ordering $>$ a *degree based monomial ordering* if $\forall \ \alpha, \beta : \deg(\alpha) > \deg(\beta) \Rightarrow \alpha > \beta$ (e.g., Dp and dp and their weighted relatives are degree based orderings).

Due to the nature of the GB algorithm, monomial operations are by far the most frequently used primitive operations. For example, monomial comparisons are performed much more often than, and monomial additions at least as often as, arithmetic operations over the coefficient field. The number of divisibility tests depends very much on the given input ideal, but is usually very large, as well (see also Table 1).

Nevertheless, whether or not monomial operations dominate the running time of a GB computation depends on the coefficient field of the underlying polynomial ring: monomial operations are certainly run–time dominating for finite fields with a small[1] characteristic (e.g., integers modulo a small prime number), since an arithmetic operation over these fields can usually be realized much faster than a monomial operation. However, for fields of characteristic 0 (like the rational numbers), GB computations are usually dominated by the arithmetic operations over these fields, since the time needed for these operations is proportional to the size of the coefficients which tend to grow rapidly during a GB computation.

Therefore, improvements in the efficiency of monomial operations will have less of an impact on GB computations over fields of characteristic 0.

2.2 monomial representations

As a first question, one might wonder which kind of polynomial (and, consequently, monomial) representation is best suited for GB computations. Although there are several alternatives to choose from (e.g., distributive/recursive, sparse/dense), it is generally agreed upon that efficiency considerations lead to only one real choice: a dense distributive representation. That is, a polynomial is stored as a sequence of terms where each term consists of a coefficient and an array of exponents which encodes the respective monomial. With such a representation, one has not only very efficient access to the leading monomial of a polynomial (which is the most frequently needed part of a polynomial during GB computations), but also to the single (exponent) values of a monomial.

Now, what type should the exponent values have? Efficiency considerations lead again to only one realistic choice, namely to fixed–length integers whose size is smaller or equal to the size of a machine word. While assuring the most generality, operations on and representations of arbitrary–length exponent values usually incur an intolerable slow–down of GB computations. Of course, a fixed–length exponent size restricts the range of representable exponent

[1] say, smaller than the maximal representable machine integer, i.e. smaller than `INT_MAX`

values. However, exponent bound restrictions are usually not very critical for GB computations: on the one hand, the (worst case) complexity of GB computations grows exponentially with the degree of the input polynomials, i.e., large exponents usually make a problem practically uncomputable. On the other hand, checks on bounds of the exponent values can be realized by degree bound checks in the outer loops of the GB computation (e.g., during the S–pair selection) which makes exponent value checks in subsequent monomial operations unnecessary.

Furthermore, the degree of a monomial is so often needed during GB computations, that, as experience has shown, it is advantageous to add an additional degree field to the monomial data structure.

As an illustration, and for later reference, we show below SINGULAR's internal `Term_t` data structure:

```
struct  Term_t
{
  Term_t*      next;
  void*        coef;
  long         order;
  Exponent_t   exp[1];
};
```

Following the arguments outlined above, a SINGULAR polynomial is represented as a linked list of terms, where each term consists of a coefficient (implemented as a hidden type: could be a pointer or a `long`) and a monomial. A monomial is represented by its exponent vector, together with its degree field (`order`). The data type of the exponent values (`Exponent_t`) can be set at configuration time with the restriction that it must be an integer type whose size is a multiple of the word size of the used machine[2]. The size of the `Term_t` structure is dynamically set at run–time, depending on the number of variables in the current polynomial ring.

Based on a monomial representation like SINGULAR's, the basic monomial operations are "traditionally" implemented by straightforward realizations of their definitions. Only monomial additions need also to update the `order` field of the result monomial, which is simply accomplished by adding the degree values of the operands (since the values of the `order`–field are additive).

3 Vectorized monomial operations

The main idea behind what we call *vectorized monomial operations* is the following: provided that the size of the machine word is a multiple (say, m) of the size of one exponent, we perform monomial operations on machine words, instead of directly on exponents. By doing so, we process a vector of m exponents with word operations only, thereby reducing the length of the inner loops of the monomial operations and avoiding non–aligned accesses of exponents.

The PoSSo library [14] already uses parts of this idea to speed up monomial additions, assignments, and tests for equalities, but has not (yet?) carried these ideas over to the much more time-critical monomial comparisons and divisibility tests.

The details of this technique are based on the following lemma whose proof is straightforward and, therefore, omitted:

[2] By default, we declare `Exponent_t` to be of type `short` which seems to be a good alternative between flexibility and efficiency.

Lemma 1: Let $s_e, m \in \mathbf{N}$, and $\alpha = (\alpha_0, \ldots \alpha_{m-1}), \beta = (\beta_0, \ldots \beta_{m-1}) \in M_m$ with $\alpha_i, \beta_i < 2^{s_e - 1}$. Furthermore, let $s_w = s_e\, m$ and $a, b, \widehat{a}, \widehat{b}, \gamma_0, \ldots \gamma_{m-1}, \delta_0 \ldots \delta_{m-1} \in \mathbf{N}$ with $\gamma_i, \delta_i < 2^{s_e}$ such that

$$\begin{aligned}
a &= \alpha_0 + \alpha_1 2^{s_e} + \ldots + \alpha_{m-1} 2^{(m-1)s_e}, \\
b &= \beta_0 + \beta_1 2^{s_e} + \ldots + \beta_{m-1} 2^{(m-1)s_e} \\
\widehat{a} &= \alpha_{m-1} + \alpha_{m-2} 2^{s_e} + \ldots + \alpha_0 2^{(m-1)s_e}, \\
\widehat{b} &= \beta_{m-1} + \beta_{m-2} 2^{s_e} + \ldots + \beta_0 2^{(m-1)s_e}, \\
b + a \pmod{2^{s_w}} &= \gamma_0 + \gamma_1 2^{s_e} + \ldots + \gamma_{m-1} 2^{(m-1)s_e}, \\
b - a \pmod{2^{s_w}} &= \delta_0 + \delta_1 2^{s_e} + \ldots + \delta_{m-1} 2^{(m-1)s_e}.
\end{aligned}$$

Then the following holds:

$$\begin{aligned}
b - a \pmod{2^{s_w}} < 2^{s_w - 1} \quad &\text{iff} \quad \mathrm{Lex}(\alpha, \beta) = 1 && (1) \\
\widehat{b} - \widehat{a} \pmod{2^{s_w}} < 2^{s_w - 1} \quad &\text{iff} \quad \mathrm{RevLex}(\alpha, \beta) = -1 && (2) \\
(\gamma_0, \ldots, \gamma_{m-1}) \quad &= \quad \alpha + \beta && (3) \\
\forall\, 0 \le i < m : \delta_i < 2^{s_e - 1} \quad &\text{iff} \quad \alpha \,|\, \beta && (4)
\end{aligned}$$

Lemma 1 shows how the monomial operations on the right–hand sides of (1) – (4) can be "vectorized", such that, on a 2's complement machine with word–size s_w, the checks (resp. operations) on the left hand sides of (1) – (4) can be performed in at most three single machine instructions (see the source code examples below for further details).

However, before Lemma 1 can be applied in an actual implementation, some technical difficulties have to be overcome:

Firstly, stricter bounds on the values of the single exponents have to be assured (i.e., the exponent values need to be less than $2^{s_e - 1}$ instead of 2^{s_e}).

Secondly, the condition $s_w = s_e\, m$ implies that the total length of the exponent vector has to be a multiple of the word–size which requires that $(-n) \pmod m$ "unnecessary" exponents (whose value is set to and always kept at 0) might have to be added to the exponent vector .

Thirdly, and most importantly, the order and arrangement of the exponents within the exponent vector has to be adjusted, depending on the monomial ordering and on the endianess of the used machine. On big–endian machines, the order of the exponents has to be reversed for reverse lexicographical orderings whereas on little–endian machines, the order of the exponents has to be reversed for lexicographical orderings. In practice, this fact can be hidden behind appropriate macro (or inline) definitions for accessing single exponents. In our implementation, we used a global variable called `pVarOffSet` and implemented the exponent access macro as follows:

```
#define pGetExp(p, i)    \
        p->exp[(pVarOffSet ? pVarOffSet - i : i)]
```

Provided that n is the number of variables of the current ring then we set the value of `pVarOffSet` as follows:

type of ordering	type of machine big–endian	little–endian
lexicographical	0	$n-1$
reverse lexicographical	$n-1$	0

Some source code fragments can probably explain it best: Figure 1 shows (somewhat simplified) versions of our implementation of the vectorized monomial operations. Some explanatory remarks are in order:

```
inline long MonComp(Term_t* a, Term_t* b)
{
  // return = 0, if a = b
  //        > 0, if a > b
  //        < 0, if a < b

  // first, check order
  long d = a->order - b->order;
  if (d) return d*pOrdSgn;

  // second, check Lex, resp. RevLex
#ifdef WORDS_BIGENDIAN
   for (long i = 0; i<n_w; i++)
#else
   for (long i = n_w -1; i; i--)
#endif
   {
     d=((long*)a->exp)[i]-((long*)b->exp)[i];
     if (d) return d*pLexSgn;
   }
   return 0;
}

inline void MonAdd(Term_t* c, Term_t* a, Term_t* b)
{
   // Set c = a + b
   for (long i=0; i<n_w; i++)
     ((long*)c->exp)[i] =
       ((long*)a->exp)[i] + ((long*)b->exp)[i];

  // update order field
  c->order = a->order + b->order;
}

inline bool MonDivBy(Term_t* a, Term_t* b)
{
  // return true,  if a divides b
  //        false, otherwise
#if SIZEOF_LONG == 4
  #if SIZEOF_EXPONENT == 1
    #define DIV_MASK 0x80808080
  #else // SIZEOF_EXPONENT == 2
    #define DIV_MASK 0x80008000
  #endif
#else // now assume SIZEOF_LONG == 8
  #if SIZEOF_EXPONENT == 1
    #define DIV_MASK 0x8080808080808080
  #elif  SIZEOF_EXPONENT == 2
    #define DIV_MASK 0x8000800080008000
  #else // SIZEOF_EXPONENT == 4
    #define DIV_MASK 0x8000000080000000
  #endif
#endif

  for (long i=0; i<n_w; i++)
  {
    long d = ((long*) b->exp)[i]
             - ((long*) a->exp)[i];
    if (d & DIV_MASK) return false;
  }
  return true;
}
```

Figure 1: Vectorized monomial operations in Singular

`n_w` is a global variable denoting the length of the exponent vectors in machine words (i.e., if s_w is the size of a machine word, s_e the size of an exponent, n the number of variables, then $n_w = \lceil n\, s_e/s_w \rceil$).

`pLexSgn` and `pOrdSgn` (used in `MonComp`) are global variables which are used to appropriately manipulate the return value of the comparison routine and whose values are set as follows:

	lp	ls	Dp	Ds	dp	ds
pOrdSgn	1	-1	1	-1	1	-1
pLexSgn	1	-1	1	1	-1	-1

Notice that `MonAdd` works on three monomials and it is most often used as a "hidden" initializer (or, assignment), since monomial additions are the "natural source" of most new monomials.

Our actual implementation contains various, tediously to describe, but more or less obvious, optimizations (like loop unrolling, use of pointer arithmetic, replacement of multiplications by bit operations, etc). We apply, furthermore, the idea of "vectorized operations" to monomial assignments and equality tests, too. However, we shall not describe here the details of these routines, since their implementation is more or less obvious and they have less of an impact on the running time than the basic monomial operations.

So much for the theory, now let us look at some actual timings: Table 1 shows various timings illustrating the effects of the vectorized monomial operations described in this section. In the first column, we list the used examples — details about those can be found in the Appendix. All GB computations were done using the degree reverse lexicographical ordering (dp) over the coefficient field $\mathbf{Z}/32003$. We measured the following times (all in seconds):

$t_{1.0}$ Running time of Singular version 1.0 which uses the straightforward implementation of monomial operations.

$t_{n,s}$, $t_{n,c}$ Running time of Singular *without* vectorized monomial operations and exponents of type `short` (for $t_{n,s}$) and `char` (for $t_{n,c}$)[3].

t_s, t_c Running time of Singular *with* vectorized monomial operations and exponents of type `short` (for t_s) and `char` (for t_c).

%mon,%comp,%div,%add Percentage of the running time (of $t_{n,s}$ runs) spent in basic monomial operations, monomial comparisons, divisibility tests, and additions, respectively (i.e., %mon = %comp + %div + %add)

Before evaluating these numbers, we should like to mention that running the same tests on different machines and/or with different simple monomial orderings leads to a similar picture (see also Appendix B for more timings).

Now, what do these numbers tell us?

Firstly, they support our assertion that for GB computations over finite fields of small characteristic, the running time is largely dominated by basic monomial operations (see column %comp). However, in which of the (three) basic monomial operations the most time is spent varies very

[3] Using our previously introduced terminology, we have $s_e = 2$ for $t_{n,s}$ (resp. 4 for $t_{n,c}$); and $m = s_w/s_e = 2$ (resp. 4 for $t_{n,c}$) on a 32–bit machine (like the HP C160, or Pentium Pro) and $m = 4$ for $t_{n,s}$ (resp. 8 for $t_{n,c}$) on a 64–bit machine (like the DEC Alpha).

Example	$t_{n,s}$	%mon	%comp	%div	%add	$\frac{t_{n,s}}{t_s}$	$\frac{t_{n,c}}{t_c}$	$\frac{t_s}{t_c}$	$\frac{t_{1.0}}{t_{n,s}}$	$\frac{t_{1.0}}{t_c}$
ecyclic 6	3.9	79.4	68.9	1.6	8.8	1.7	2.0	1.3	1.3	2.8
homog cyclic 7	125.8	72.5	33.6	11.8	27.1	1.5	1.9	1.3	1.3	2.6
homog cyclic 6	0.9	69.5	30.5	13.3	25.8	1.6	1.9	1.3	1.3	2.6
katsura 7	8.8	71.5	41.5	3.7	26.3	1.6	1.9	1.2	1.3	2.6
rcyclic 15	53.2	76.2	29.6	26.0	20.6	1.6	2.0	1.4	1.2	2.5
rcyclic 14	22.9	76.0	28.1	28.5	19.4	1.5	1.9	1.3	1.2	2.5
rcyclic 16	106.4	78.2	26.5	30.7	21.0	1.6	2.0	1.4	1.2	2.5
rcyclic 19	740.0	80.3	29.5	29.9	20.9	1.5	1.9	1.4	1.2	2.5
rcyclic 18	396.5	80.0	29.4	29.2	21.4	1.5	1.9	1.4	1.2	2.4
rcyclic 13	10.2	74.8	27.6	26.6	20.7	1.5	1.8	1.3	1.2	2.4
rcyclic 12	4.3	74.7	28.4	29.8	16.5	1.6	1.9	1.2	1.2	2.4
rcyclic 11	1.6	67.4	22.9	24.7	19.8	1.6	1.8	1.2	1.2	2.4
bjork 8	4.0	67.8	32.0	13.2	22.5	1.5	1.8	1.2	1.3	2.3
rcyclic 17	210.9	78.2	26.5	30.3	21.4	1.5	1.8	1.3	1.2	2.3
gerhard 1	1.7	68.7	48.1	4.5	16.0	1.3	1.5	1.2	1.5	2.3
katsura 8	94.3	73.1	45.1	3.5	24.5	1.5	1.7	1.2	1.3	2.3
homog 2mat3	126.0	81.9	77.5	1.9	2.5	1.3	1.5	1.4	1.2	2.2
cyclic 7	190.0	66.9	31.0	11.2	24.7	1.4	1.8	1.3	1.2	2.2
rcyclic 10	0.5	53.7	16.2	20.0	17.5	1.6	1.7	1.1	1.2	2.2
cyclic 6	0.7	61.6	26.8	16.3	18.6	1.4	1.7	1.2	1.2	2.1
homog gonnet	125.4	76.8	43.0	31.5	2.2	1.5	1.7	1.2	1.1	2.1
gerhard 2	13.8	72.3	54.8	3.2	14.3	1.3	1.4	1.1	1.4	2.1
homog alex 2	12.9	71.7	58.2	4.8	8.8	1.3	1.4	1.1	1.4	2.0
schwarz 11	131.6	65.4	30.6	18.4	16.5	1.3	1.6	1.3	1.2	2.0
2mat3	123.1	80.2	75.7	1.9	2.6	1.2	1.4	1.4	1.2	2.0
schwarz 10	21.9	62.5	28.2	18.4	15.9	1.3	1.5	1.2	1.2	2.0
cohn2	9.7	64.4	29.7	4.1	30.6	1.4	1.5	1.1	1.2	1.9
gerhard 3	24.4	67.1	45.5	13.1	8.6	1.3	1.3	1.1	1.3	1.8
symmetric 6	60.6	68.0	25.6	14.5	28.0	1.4	1.4	1.1	1.2	1.8
schwarz 9	3.3	55.5	23.0	20.0	12.5	1.3	1.4	1.1	1.2	1.8
alex 2	7.1	63.6	52.1	3.3	8.2	1.3	1.4	1.1	1.2	1.8
schwarz 8	0.6	50.6	24.7	14.6	11.2	1.3	1.5	1.1	1.1	1.7
ecyclic 7	1076.4	86.4	84.2	0.3	2.0	1.2	1.4	1.2	1.1	1.6
alex 3	2.0	65.9	53.5	7.7	4.7	1.3	1.3	1.1	1.2	1.6
homog alex 3	2.0	66.2	47.2	14.8	4.1	1.2	1.2	1.1	1.3	1.6
gonnet	1.2	43.8	6.2	34.2	3.4	1.1	1.3	1.2	1.1	1.5
averages										
Pentium Pro	103.3	69.8	38.4	15.6	15.8	1.4	1.6	1.2	1.2	2.2
HP C160	125.4	69.0	41.7	17.9	9.5	1.2	1.5	1.3	1.8	2.9
DEC Alpha	221.4	69.8	37.7	16.1	16.0	1.7	1.7	1.2	1.1	2.3

Table 1: Detailed timings for vectorized monomial operations: SINGULAR compiled with gcc version 2.7.2 for a Pentium Pro 200 (32 bit, little–endian) running Linux and average of timings for a HP C160 (32 bit, big–endian) running HP–UX 10.20 and a DEC Alpha (64 bit, little endian) running Linux.

much from example to example (compare, e.g., line "gonnet" with line "ecyclic 6").

Secondly, and most importantly: the impact of the vectorized monomial operations is quite substantial (see columns $t_{n,s}/t_s$, $t_{n,c}/t_c$ which show the speedup gained by vectorized operations). As expected, the larger the ratio $m = s_w/s_e$ (i.e, the number of exponents packed in one machine word), the more speedup is usually gained (compare column $t_{n,s}/t_s$ and $t_{n,c}/t_c$). However, notice that we cannot conclude a direct correlation between the percentage of time spent in monomial operations and the efficiency gains of vectorized operations. This is due to the fact that the number of inner loop iterations (and, hence, reduction of inner loop iterations) for comparisons and divisibility tests is not constant, but depends on the input monomials.

Thirdly: As we would expect, the more exponents are encoded into one machine word, the faster the GB computation is accomplished (see the t_s/t_c column). This has two main reasons: first, more exponents are handled by one machine operation; and second, less memory is used, and therefore, the memory performance increased (e.g., number of cache misses is reduced). However, we also need to keep in mind that the more exponent are encoded into one word, the smaller are the upper bounds on the value of a single exponents. This is especially crucial for computations with `char` exponents, since these require that an exponent may not be larger than 127. Ideally, we would like to "dynamically" switch from one exponent size to the next larger one, whenever it becomes necessary. With SINGULAR, we cannot do this yet, at the moment, but we intend to implement this feature in one of the next major upgrades.

Fourthly: We would like to point out that the only difference between the SINGULAR versions resulting in $t_{1.0}$ and $t_{n,s}$ is that monomial comparisons for the different orderings are not reduced to one inlined routine but instead are realized by function pointers (and, therefore, each monomial comparisons requires a function call). Hence, column $t_{1.0}/t_{n,2}$ shows that already the in–place realization of monomial comparisons results in a considerable efficiency gain which by far compensates the additional overhead incurred by the additional indirection for accessing one (single) exponent value (i.e., the overhead incurred by the pGetExp macro shown above).

Last, but not least, the combination of all these factors into our newest and fastest version of SINGULAR, leads to a rather significant improvement (see column $t_{1.0}/t_c$) over the "old" SINGULAR version 1.0.

4 Rank-based monomial representation and comparisons

A different monomial representation is what we call the *rank-based representation*. This monomial representations was, to the best of our knowledge, first developed for and used in the CAS Macaulay [4]. The main idea behind this monomial representation stems from the following property of a degree based ordering $>$ on M_n:

$$\forall m \in M_n : \#\{m' \in M_n : 1 \leq m' \leq m\} < \infty \qquad (5)$$

Therefore, it is possible to bijectively enumerate the monomials in the order given by the monomial ordering and to represent an entire monomial by just a single integer, which we call the *rank* of a monomial.

The bijective maps between the monomials and the integers depend on the used monomial ordering: for the degree reverse lexicographic ordering it is given by the following lemma:

Lemma 2: Let $\alpha \in M_n$ and $r_{dp} : M_n \to \mathbf{N}$ given by

$$r_{dp}(\alpha) := \sum_{i=1}^{n} \binom{\sum_{j=1}^{i} \alpha_j + i - 1}{i}. \qquad (6)$$

Then the following holds:

$$\forall \alpha, \beta \in M_n : r_{dp}(\alpha) > r_{dp}(\beta) \Leftrightarrow \alpha >_{dp} \beta$$

$$\forall \alpha, \beta \in M_n : r_{dp}(\alpha) = r_{dp}(\beta) \Leftrightarrow \alpha = \beta.$$

A similar bijective map can be given for the ordering Dp (degree lexicographic). For negative orderings like Ds or ds the same idea can be exploited but it is necessary to change the signs since the zero monomial is, with respect to these orderings, the largest, not the smallest monomial.

Using the rank-based representation for monomials, monomial comparisons can very easily and efficiently be realized by simply comparing their corresponding ranks (we also call such monomial comparisons simply *rank-based comparisons*). Furthermore, a pure rank-based polynomial representation is very compact, i.e. requires an almost minimal amount of memory. Unfortunately, we have to pay for these advantages with the following difficulties:

(i) Monomial additions and divisibility tests can not be easily accomplished, since the bijective enumeration maps are neither additive nor do they allow any direct conclusions about monomial divisibility.

(ii) An unrestricted realization of the enumeration maps would require that arbitrary precision integers are used which would in turn result in considerable performance losses. On the other hand, restricting the range of the enumeration maps to machine integers, imposes a limit on the largest representable monomial (which we also call the *maximal integer monomial*). Clearly, this limit depends on the word–length of the machine and on the number of variables of the polynomial ring. For example, using a 32–bit integer for the realization of (6) restricts the degree of the maximal integer monomials as shown in the following table:

variables	3	4	5	7	10	15	20	25
max. deg	2340	471	186	66	31	17	12	9

(iii) Non–degree orderings (like the lexicographical or negative lexicographical ordering) do not admit a bijective enumeration map which is compatible with the monomial ordering, since they do not enjoy (5).

Problem (i) is overcome in the Macaulay system by keeping two representations of the the leading monomial of a polynomial (the integer representation and the exponent vector representation) and by computing the inverse of the enumeration map for the non–leading monomials of a polynomial.

Problem (ii) is overcome by restricting the input polynomials to being homogeneous and by Macaulay's "infamous" degree bound messages and enforcement mechanisms.

Problem (iii) is simply ignored by Macaulay — no computations with non–degree orderings are possible.

In SINGULAR, we choose a slightly different way to overcome problems (i) – (iii): first, we always keep the exponent vector representation of a monomial. And, second, if a monomial is smaller than the maximal integer monomial, we store its rank in the negative range of the `order` field (i.e., `a->order + INT_MAX` yields the rank of the monomial `a`). If, on the other hand, a monomial is larger than the maximal integer monomial, then its `order` field contains (as before) the (positive) degree of the monomial.

Based on such a data representation, monomial comparisons can simply be accomplished as follows:

```
inline long MonComp(Term_t* a, Term_t* b)
{
  if (a->order != b->order)
  {
    if (a->order > b->order) return 1;
    return -1;
  }

  // a->order==b->order: are we already done?
  if (a->order < 0) return 0;  // monom's are equal

  /* now compare using the exponent values,
     assuming the degrees are equal ... */
}
```

Notice that we do not need to distinguish between the two different interpretations of the `order` field: if we compare a monomial which is smaller than the maximal integer monomial (`order` value is negative) with one that is larger (`order` value is positive), then the different signs of the `order` fields result in an immediate and correct return of the monomial comparison. Therefore, only a sign–check (for the case of equal `order` fields) needs to be added to the previously given monomial comparison routines to have them take advantage of ranks.

Since we always keep the exponent vector representation alongside with the ranks, monomial divisibility tests

exponent type monomial operations	char vectorized				short vectorized				short not vectorized			
Example	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$
2mat3	1.3	7.6	69.0	1.1	1.7	6.5	74.4	1.4	1.9	5.0	76.3	1.6
homog gonnet	1.3	6.7	34.5	1.0	1.8	4.7	38.0	1.1	2.7	4.8	43.9	1.3
homog alex 2	1.3	9.7	51.1	0.9	1.4	8.5	51.2	1.0	1.8	7.8	57.4	1.2
alex 2	1.3	7.9	52.1	0.9	1.1	7.4	49.0	1.0	1.6	7.5	56.8	1.2
alex 3	1.3	5.2	47.8	1.0	1.0	2.9	49.1	0.9	1.4	2.8	53.7	1.2
gerhard 1	1.4	15.6	41.9	0.8	1.2	12.8	41.5	0.9	2.0	13.4	48.1	1.2
homog 2mat3	1.2	7.2	70.6	1.0	1.3	4.6	75.4	1.0	1.6	4.1	78.2	1.2
katsura 8	1.4	26.3	39.6	0.8	1.5	25.0	38.5	0.9	2.2	21.6	44.6	1.2
gerhard 2	1.1	10.4	47.4	0.9	1.3	12.6	48.4	0.9	1.7	10.6	54.3	1.1
katsura 7	1.2	28.3	32.7	0.7	1.4	23.0	30.2	0.8	2.1	22.5	38.2	1.1
gerhard 3	1.3	8.0	40.0	0.9	1.5	7.6	43.4	0.9	1.7	7.6	45.4	1.1
homog alex 3	1.2	5.0	39.1	1.0	1.2	5.8	42.0	1.0	1.4	5.8	43.1	1.1
cohn2	1.4	19.6	22.2	0.8	1.3	19.8	23.8	0.9	2.1	19.6	29.9	1.1
bjork 8	1.2	23.6	21.5	0.7	1.6	23.2	21.1	0.8	2.3	21.8	30.1	1.1
cyclic 7	1.3	25.0	22.0	0.8	1.6	22.3	23.8	0.8	2.6	19.6	31.4	1.1
schwarz 11	1.2	20.1	24.4	0.8	1.6	15.9	25.4	0.9	2.2	14.7	30.1	1.1
symmetric 6	1.3	22.1	18.3	0.8	1.5	20.0	19.1	0.9	2.4	18.0	26.9	1.1
schwarz 10	1.2	19.8	20.1	0.8	1.5	15.6	21.1	0.9	2.7	15.5	28.5	1.1
homog cyclic 7	1.4	28.6	24.8	0.7	1.6	24.3	24.2	0.8	2.8	22.0	33.7	1.1
rcyclic 13	1.5	31.1	21.9	0.7	1.7	24.5	19.8	0.9	2.9	19.2	26.1	1.0
rcyclic 14	1.4	28.5	19.1	0.7	2.0	20.3	19.5	0.9	3.5	20.6	27.7	1.0
schwarz 9	1.2	21.2	15.7	0.8	1.4	15.3	18.7	0.9	2.0	16.2	22.4	1.0
rcyclic 12	2.0	25.1	24.2	0.7	1.3	20.1	16.5	0.9	3.1	20.3	24.0	1.0
rcyclic 11	1.4	26.2	21.1	0.7	1.3	23.1	19.1	0.8	3.7	18.7	30.1	1.0
rcyclic 15	1.3	26.2	20.9	0.7	1.5	19.3	20.4	0.9	2.2	17.2	30.3	1.0
homog cyclic 6	1.1	27.3	15.2	0.6	3.1	19.0	32.9	0.8	3.3	19.3	28.0	1.0
gonnet	1.4	7.9	5.6	0.9	1.3	9.3	4.0	0.9	2.3	6.6	9.2	1.0
ecyclic 7*	1.0	3.2	81.4	1.0	1.0	1.7	82.3	1.0	1.0	1.3	84.5	0.9
rcyclic 16*	1.1	17.6	19.4	0.8	1.0	11.8	18.4	0.9	1.0	9.6	26.6	0.9
rcyclic 17*	0.9	15.9	20.5	0.8	0.9	10.8	19.1	0.9	0.9	8.2	27.1	0.9
rcyclic 18*	0.9	15.9	18.2	0.8	1.0	10.0	19.6	0.9	1.0	7.8	29.5	0.9
rcyclic 19*	1.0	16.0	20.4	0.8	1.0	10.1	19.9	0.9	1.1	7.8	29.7	0.9
ecyclic 6*	0.9	18.0	55.3	0.8	1.0	10.9	51.4	0.8	1.2	8.6	66.3	0.8
averages												
Pentium Pro	1.3	17.1	33.6	0.8	1.4	13.8	33.8	0.9	2.1	12.6	40.3	1.1
HP C160	1.4	9.1	36.6	0.8	1.6	8.4	39.0	0.9	1.7	9.1	41.6	1.1
DEC Alpha	1.8	11.8	33.1	0.8	1.5	11.7	31.5	0.8	2.4	10.2	35.5	1.0

Table 2: Detailed timings for rank-based monomial operations – same settings as in Table 1.

and additions can be accomplished as usual, except that we need to recompute and set the `order` field at the end of the monomial addition routine. For monomials smaller than the maximal integer monomial, this basically amounts to computations of the respective ranks. We accomplish the latter by a simple implementation of (6) using a pre–computed table of binomial coefficients.

As with vectorized monomial additions, overflow or degree checks can largely be moved to the outer loops of the GB computations so that we can almost always avoid the costly checks on whether or not the sum of two monomials is smaller than the maximal integer monomial.

As a summary, we can conclude that with a relatively small effort, it is possible to extend a "traditional" exponent vector–based monomial representation so that it can fully take advantage of ranks while, at the same time, the limitations of the Macaulay system (degree bounds, no support for non–degree orderings, repeated inverse rank computations) can be avoided.

But now again, let us examine some timings (Table 2) to see how much all that buys us in practice: we again used the examples of the Appendix and performed all GB computations using the degree reverse lexicographical ordering (dp) over the coefficient field $\mathbf{Z}/32003$. We obtained timings for three basic SINGULAR configurations: `char` and `short` exponents with vectorized monomial operations, and `short` exponents with "traditional" (i.e., not vectorized) monomial operations. For each of these "conventional" configurations, we also obtained a SINGULAR version which uses rank-based comparisons and measured the following:

$t_<, t_{m,<}$**:** the time spent in monomial comparisons, only ($t_<$–conventional comparisons, $t_{m,<}$–rank-based comparisons).

%R, %comp: the percentage of time spent in rank computations (%R) and in conventional monomial comparisons (%comp).

t, t_r**:** the overall running time for the GB computation (t–conventional configuration, t_m–rank-based configuration). For the conventional configuration, these timings are the same as those used in the previous timing, Table 1.

And again, let us see what we can conclude from these experiments:

Firstly, rank-based comparisons are realized substantially faster than traditional and vectorized monomial comparisons (see the $t_</t_{r,<}$ columns), except in such examples, where a majority of the monomials is larger than the maximal integer monomial (these examples are marked with a $*$).

Secondly, the time spent for rank computations during monomial additions is significant, and cannot be neglected (see the %R columns).

Thirdly, whether or not the time saved by rank-based comparisons pays off against the additional time spent for rank computations, depends on the relative efficiency of the "conventional" monomial comparisons and on the characteristics of the example — where the percentage of time spent in monomial comparisons seems to be the most important factor (compare the %comp and t/t_r columns).

Finally, evaluating the t/t_r columns, we can conclude good news and bad news: the good news is that rank-based comparisons, by and large, lead to an overall speedup when used in combination with traditional (i.e., not vectorized) monomial operations. That is, our results make a strong case for such a usage of the rank-based representation. The bad news is, that rank-based comparisons by and large do *not* lead to an overall speedup when used in combination with vectorized monomial operations. Our results show that the more efficient vectorized monomial operations result in smaller speedups of the rank-based comparisons, i.e., speedups which are generally too small to outweigh the additional cost of computing ranks.

5 Summary

In this paper, we examined monomial representations and operations for GB computations. In section 2 we argued that (i) the canonical polynomial representation is a linked list of terms where each term consists of a coefficient, a degree field and an array of fixed-length exponent values, and, that (ii) the canonical implementation of monomial operations is simply the straight-forward realization of their definitions.

In section 3 we introduced the idea of vectorized monomial operations: Instead of working on one exponent at a time, we work with as many exponents at a time, as fit into one machine word (say, m). Vectorized monomial operations can be applied to monomial additions, divisibility tests, assignments, and comparisons w.r.t. simple monomial orderings; and they have the following advantages:

1. The maximal number of iterations in the inner loops of the monomial operations is cut by approximately a factor of m.
2. Non–aligned access of single exponents are avoided.
3. Monomial comparisons for the different monomial orderings are reduced to *one* routine which does not need to explicitly distinguish the different simple monomial orderings.

We have illustrated with our timings that these advantages directly translate in significant overall speedups of GB computations (which are in the range of 1.5 to 3 in our implementation).

The concept of vectorized monomial comparisons could be extended to non-simple orderings (like elimination, block, or even matrix orderings) based on the following idea: Let A be a matrix describing an arbitrary monomial ordering. Besides representing a monomial by a vector α of exponents we keep an additional vector $A\,\alpha$ and accomplish divisibility tests as before, comparisons by vectorized lexicographical comparisons of the $A\,\alpha$ vectors and additions by vectorized additions of both, α and $A\,\alpha$. However, the practicality of this idea remains to be investigated.

In section 4 we examined a rank-based representations of monomials. For degree-based orderings, we can uniquely associate an order-preserving rank with each monomial, and use this rank to reduce a monomial comparisons to just one integer comparison. Unfortunately, our timings indicate that this technique does not generally lead to efficiency gains over vectorized monomial operations, since the time saved in monomial comparisons is usually not enough to make up for the additional time spent in rank computations.

In other words, we highly recommend the usage of vectorized monomial operations in GB computations, whereas we can not recommend rank-based monomial representations and comparisons.

Although parts of our conclusions are based on the timings obtained from a SINGULAR implementation of these techniques, they are system-independent in their nature and should therefor apply in similar ways to other GB implementations.

With our results we have shown that it can be rewarding to systematically examine implementational aspects of GB computations. We hope to continue along these lines and, especially, to get more programmers of the various GB systems to join in such discussions and investigations.

6 Acknowledgments

We should like to thank all members of the SINGULAR developer team for their support of our work and for giving us some time off from our day-to-day developer tasks. We furthermore are grateful to F. O. Schreyer and his colleagues for placing their DEC Alpha machines at our disposal.

References

[1] ATTARDI, G., AND FLAGELLA, T. Memory management in the PoSSo solver. *J. Symbolic Computation 21*, 3 (March 1996), 293–312.

[2] BAYER, D., AND MUMFORD, D. What can be computed in algebraic geometry? Cambridge University Press, Cambridge, 1993, pp. 1–48.

[3] BAYER, D., AND STILLMAN, M. A theorem on refining division orders by the revers lexicographical order. *Duke J. Math. 55* (1987), 321–328.

[4] BAYER, D., AND STILLMAN, M. *Macaulay Classic*: A computer algebra system for algebraic geometry, 1993. Available via anonymous ftp from `ftp://math.harvard.edu/Macaulay`.

[5] BJÖRCK, G., AND FRÖBERG, G. A faster way to count the solution of inhomogeneous systems of algebraic equations, with applications to cyclic n-roots. *J. Symbolic Computation 12*, 3 (September 1991), 329–336.

[6] BOEGE, W., GEBAUER, R., AND KREDEL, H. Some Examples for Solving Systems of Algebraic Equations

by Calculating Groebner Bases. *J. Symbolic Computation 2*, 1 (March 1986), 83–98.

[7] BUCHBERGER, B. Groebner bases: an algorithmic method in polynomial ideal theory. D. Reidel Publishing Company, 1985, pp. 184–232.

[8] CABOARA, M., DE DOMINICIS, G., AND ROBBIANO, L. Mutltigraded Hilbert Functions and Buchberger Algorithm. In *Proc. of the International Symposium on Symbolic and Algebraic Computation (ISSAC'96)* (Zurich, Switzerland, July 1996), ACM Press, pp. 72–78.

[9] COX, D., LITTLE, J., AND O'SHEA, D. *Ideals, Varieties, and Algorithms*, 2nd ed. Springer-Verlag, 1997.

[10] GIOVINI, A., MORA, T., NIESI, G., ROBBIANO, L., AND TRAVERSO, C. One sugar cube, please or Selection strategies in Buchberger algorithms. In *Proceedings of the 1991 International Symposium on Symbolic and Algebraic Computations, ISSAC'91* (1991), S. Watt, Ed., ACM press, pp. 49–54.

[11] GREUEL, G.-M., PFISTER, G., AND SCHÖNEMANN, H. Singular Reference Manual. In *Reports On Computer Algebra*, no. 12. Centre for Computer Algebra, University of Kaiserslautern, May 1997. `http://www.mathematik.uni-kl.de/~zca/Singular`

[12] KAHN, K. private communications, 1993.

[13] KATSURA, W., FUKUDA, S., INAWASHIRO, N., FUJIKI, N. M., AND GEBAUER, R. *Cell Biophysics*, vol. 11. 1987, pp. 309–319.

[14] *PoSSo*: Polynomial System Solving, 1995. `http://posso.dm.unipi.it/`.

[15] ROBBIANO, L. Term Orderings on the Polynomial Ring. In *Proceedings of EUROCAL 85, Lecture Notes in Computer Science* **204** (1985), pp. 513–517.

Appendix A: Benchmark examples

I this appendix, we show some more details about the sets of polynomial we used for our GB computation test runs.

Example	#vars	#polys	homog	Deg_s	$\mathrm{Deg}_{\max}$	Ref
ecyclic 7	43	7	no	7	27	
ecyclic 6	31	6	no	6	17	
rcyclic $10 \le i \le 19$	i	$i-1$	no	$i-1$	$\lfloor i/2 \rfloor * 2+4$	
homog 2mat3	19	8	yes	4	13	[14]
2mat3	18	8	no	4	13	
homog gonnet	18	19	yes	2	11	[6]
gonnet	17	19	no	2	11	[6]
schwarz 11	11	11	no	2	13	
schwarz 10	10	10	no	2	12	
katsura 8	9	9	no	2	10	[13]
katsura 7	8	8	no	2	9	[13]
bjork 8	8	9	no	8	18	[5]
homog cyclic 7	8	7	yes	7	20	[5]
cyclic 7	7	7	no	7	27	[5]
homog cyclic 6	7	6	yes	6	17	[5]
cyclic 6	6	6	no	6	17	[5]
homog alex 3	6	4	yes	14	51	
alex 3	5	4	no	14	51	
gerhard 1	5	3	yes	10	32	
symmetric 6	5	5	yes	6	23	[12]
homog alex 2	5	3	yes	12	40	
cohn2	4	4	no	6	20	[14]
alex 2	4	3	no	12	33	
gerhard 2	4	3	yes	9	44	
gerhard 3	4	3	yes	23	81	

Table 3: Summary of properties of benchmark examples

Table 3 lists a summary of their properties: column *#vars* shows the number of occurring variables, column *#polys* the number of elements (polynomials), column *homog* gives the homogeneity, and Deg_s shows the maximal degree of the input sets. $\mathrm{Deg}_{\max}$ gives the maximal degree occurring during the GB computation w.r.t. the degree reverse lexicographical ordering. The last column gives references to, our sources of these examples[4]. Those without a reference are from the collection of examples of the SINGULAR team.

Finally, in the rest of this appendix, we completely list the all the used examples.

cyclic n: $[x_1, \dots x_n]$, n generators p_k:

$$1 \le k < n : p_k = \sum_{j=1}^{n} \prod_{i=j}^{j+k-1} x_i, \quad p_n = \prod_{i=1}^{n} x_i - 1$$

For example, cyclic 4:

$$\begin{pmatrix} x_1 + x_2 + x_3 + x_4 \\ x_1x_2 + x_2x_3 + x_3x_4 + x_4x_1 \\ x_1x_2x_3 + x_2x_3x_4 + x_3x_4x_1 + x_4x_1x_2 \\ x_1x_2x_3x_4 - 1 \end{pmatrix}$$

homog cyclic n: $[x_1, \dots x_{n+1}]$, n generators p_k:

$$1 \le k < n : p_k = \sum_{j=1}^{n} \prod_{i=j}^{j+k-1} x_i, \quad p_n = \prod_{i=1}^{n} x_i - x_{n+1}^n$$

[4] We tried to follow the paths of these examples to their origins as much as possible.

For example, homog cyclic 4:

$$\begin{pmatrix} x_1 + x_2 + x_3 + x_4 \\ x_1x_2 + x_2x_3 + x_3x_4 + x_4x_1 \\ x_1x_2x_3 + x_2x_3x_4 + x_3x_4x_1 + x_4x_1x_2 \\ x_1x_2x_3x_4 - x_5^4 \end{pmatrix}$$

rcyclic n: $[x_1, \dots x_n]$, n generators p_k:

$$1 \le k < n : p_k = \sum_{j=1}^{3} \prod_{i=j}^{j+k-1} x_i, \quad p_n = \prod_{i=1}^{n} x_i - 1$$

For example, rcyclic 4:

$$\begin{pmatrix} x_1 + x_2 + x_3 \\ x_1x_2 + x_2x_3 + x_3x_4 \\ x_1x_2x_3 + x_2x_3x_4 + x_3x_4x_1 \\ x_1x_2x_3x_4 - 1 \end{pmatrix}$$

ecyclic n: $[x_1, \dots x_{n(n-1)+1}]$, n generators p_k:

$$1 \le k < n : p_k = \sum_{j=1}^{n} \prod_{i=j}^{j+k-1} x_{(j-1)n+i}, \quad p_n = \prod_{i=1}^{n} x_{(i-1)n+1}$$

For example, ecyclic 4:

$$\begin{pmatrix} x_1 + x_2 + x_3 + x_4 \\ x_1x_5 + x_2x_6 + x_3x_7 + x_4x_8 \\ x_1x_5x_9 + x_2x_6x_{10} + x_3x_7x_{11} + x_4x_8x_{12} \\ x_1x_5x_9x_{13} - 1 \end{pmatrix}$$

homog 2mat3: $[a,b,c,d,e,f,g,h,i,A,B,C,D,E,F,G,H,I,t]$

$$\left(\begin{array}{c} bD+cG-dB-gC,\\ aB-bA+bE+cH-eB-hC,\\ aC+bF-cA+cI-fB-iC,\\ -aD+dA-dE+eD+fG-gF,\\ -bD+dB+fH-hF,\\ -cD+dC+eF-fE+fI-iF,\\ -aG-dH+gA-gI+hD+iG,\\ -bGt^2-cGt^2-eHt^2-fHt^2-ghCI+gBt^2+hEt^2+hFt^2+iHt^2 \end{array}\right)$$

2mat3: $[a,b,c,d,e,f,g,h,i,A,B,C,D,E,F,G,H,I]$

$$\left(\begin{array}{c} bD+cG-dB-gC,\\ aB-bA+bE+cH-eB-hC,\\ aC+bF-cA+cI-fB-iC,\\ -aD+dA-dE+eD+fG-gF,\\ -bD+dB+fH-hF,\\ -cD+dC+eF-fE+fI-iF,\\ -aG-dH+gA-gI+hD+iG,\\ -bG-cG-eH-fH-ghCI+gB+hE+hF+iH \end{array}\right)$$

homog gonnet: $[x,y,z,t,u,v,w,a,b,c,d,e,f,g,h,i,j,k]$

$$\left(\begin{array}{c} ag,\\ wk+ah+bg+df,\\ ci,\\ wf,\\ tk+ai+cg+cj+ck+di+ei,\\ yk+uk+vk+2wk+ag+aj+ak+bi+cf+ch+dg+dj+dk+eg+ej+ek,\\ xk+zk+wg+wj+wk+af+ah+bg+bj+bk+df+dh+ef+eh-k^2,\\ wf+wh+bf+bh,\\ 2wf+wh+bf,\\ xk+2wg+wj+wk+2af+ah+bg+df+ef,\\ yk+wi+2ag+aj+ak+cf+dg+eg,\\ ai+cg,\\ wg+af,\\ wh+bf+2bh,\\ zk+wk+ah+bg+bj+2bk+df+2dh+eh,\\ uk+ak+bi+ch+dg+dj+2dk+ek,\\ ck+di,\\ wh+bf,\\ ak+dg \end{array}\right)$$

gonnet: $[x,y,z,t,u,v,w,a,b,c,d,e,f,g,h,i,j]$

$$\left(\begin{array}{c} ag,\\ w+ah+bg+df,\\ ci,\\ wf,\\ t+ai+cg+cj+c+di+ei,\\ y+u+v+2w+ag+aj+a+bi+cf+ch+dg+dj+d+eg+ej+e,\\ x+z+wg+wj+w+af+ah+bg+bj+b+df+dh+ef+eh-1,\\ wf+wh+bf+bh,\\ 2wf+wh+bf,\\ x+2wg+wj+w+2af+ah+bg+df+ef,\\ y+wi+2ag+aj+a+cf+dg+eg,\\ ai+cg,\\ wg+af,\\ wh+bf+2bh,\\ z+w+ah+bg+bj+2b+df+2dh+eh,\\ u+a+bi+ch+dg+dj+2d+e,\\ c+di,\\ wh+bf,\\ a+dg \end{array}\right)$$

schwarz 10: $[x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10}]$

$$\begin{pmatrix} x_1^2 + x_1 + 2x_2x_{10} + 2x_3x_9 + 2x_4x_8 + 2x_5x_7 + x_6^2, \\ 2x_1x_2 + x_2 + 2x_3x_{10} + 2x_4x_9 + 2x_5x_8 + 2x_6x_7, \\ 2x_1x_3 + x_2^2 + x_3 + 2x_4x_{10} + 2x_5x_9 + 2x_6x_8 + x_7^2, \\ 2x_1x_4 + 2x_2x_3 + x_4 + 2x_5x_{10} + 2x_6x_9 + 2x_7x_8, \\ 2x_1x_5 + 2x_2x_4 + x_3^2 + x_5 + 2x_6x_{10} + 2x_7x_9 + x_8^2, \\ 2x_1x_6 + 2x_2x_5 + 2x_3x_4 + x_6 + 2x_7x_{10} + 2x_8x_9, \\ 2x_1x_7 + 2x_2x_6 + 2x_3x_5 + x_4^2 + x_7 + 2x_8x_{10} + x_9^2, \\ 2x_1x_8 + 2x_2x_7 + 2x_3x_6 + 2x_4x_5 + x_8 + 2x_9x_{10}, \\ 2x_1x_9 + 2x_2x_8 + 2x_3x_7 + 2x_4x_6 + x_5^2 + x_9 + x_{10}^2, \\ 2x_1x_{10} + 2x_2x_9 + 2x_3x_8 + 2x_4x_7 + 2x_5x_6 + x_{10} \end{pmatrix}$$

schwarz 11: $[x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10}, x_{11}]$

$$\begin{pmatrix} x_1^2 + x_1 + 2x_2x_{11} + 2x_3x_{10} + 2x_4x_9 + 2x_5x_8 + 2x_6x_7, \\ 2x_1x_2 + x_2 + 2x_3x_{11} + 2x_4x_{10} + 2x_5x_9 + 2x_6x_8 + x_7^2, \\ 2x_1x_3 + x_2^2 + x_3 + 2x_4x_{11} + 2x_5x_{10} + 2x_6x_9 + 2x_7x_8, \\ 2x_1x_4 + 2x_2x_3 + x_4 + 2x_5x_{11} + 2x_6x_{10} + 2x_7x_9 + x_8^2, \\ 2x_1x_5 + 2x_2x_4 + x_3^2 + x_5 + 2x_6x_{11} + 2x_7x_{10} + 2x_8x_9, \\ 2x_1x_6 + 2x_2x_5 + 2x_3x_4 + x_6 + 2x_7x_{11} + 2x_8x_{10} + x_9^2, \\ 2x_1x_7 + 2x_2x_6 + 2x_3x_5 + x_4^2 + x_7 + 2x_8x_{11} + 2x_9x_{10}, \\ 2x_1x_8 + 2x_2x_7 + 2x_3x_6 + 2x_4x_5 + x_8 + 2x_9x_{11} + x_{10}^2, \\ 2x_1x_9 + 2x_2x_8 + 2x_3x_7 + 2x_4x_6 + x_5^2 + x_9 + 2x_{10}x_{11}, \\ 2x_1x_{10} + 2x_2x_9 + 2x_3x_8 + 2x_4x_7 + 2x_5x_6 + x_{10} + x_{11}^2, \\ 2x_1x_{11} + 2x_2x_{10} + 2x_3x_9 + 2x_4x_8 + 2x_5x_7 + x_6^2 + x_{11} \end{pmatrix}$$

katsura 8: $[u_0, u_1, u_2, u_3, u_4, u_5, u_6, u_7, u_8]$

$$\begin{pmatrix} u_0 + 2u_1 + 2u_2 + 2u_3 + 2u_4 + 2u_5 + 2u_6 + 2u_7 + 2u_8 - 1, \\ u_0^2 - u_0 + 2u_1^2 + 2u_2^2 + 2u_3^2 + 2u_4^2 + 2u_5^2 + 2u_6^2 + 2u_7^2 + 2u_8^2, \\ 2u_0u_1 + 2u_1u_2 - u_1 + 2u_2u_3 + 2u_3u_4 + 2u_4u_5 + 2u_5u_6 + 2u_6u_7 + 2u_7u_8, \\ 2u_0u_2 + u_1^2 + 2u_1u_3 + 2u_2u_4 - u_2 + 2u_3u_5 + 2u_4u_6 + 2u_5u_7 + 2u_6u_8, \\ 2u_0u_3 + 2u_1u_2 + 2u_1u_4 + 2u_2u_5 + 2u_3u_6 - u_3 + 2u_4u_7 + 2u_5u_8, \\ 2u_0u_4 + 2u_1u_3 + 2u_1u_5 + u_2^2 + 2u_2u_6 + 2u_3u_7 + 2u_4u_8 - u_4, \\ 2u_0u_5 + 2u_1u_4 + 2u_1u_6 + 2u_2u_3 + 2u_2u_7 + 2u_3u_8 - u_5, \\ 2u_0u_6 + 2u_1u_5 + 2u_1u_7 + 2u_2u_4 + 2u_2u_8 + u_3^2 - u_6, \\ 2u_0u_7 + 2u_1u_6 + 2u_1u_8 + 2u_2u_5 + 2u_3u_4 - u_7 \end{pmatrix}$$

katsura 7: $[u_0, u_1, u_2, u_3, u_4, u_5, u_6, u_7]$

$$\begin{pmatrix} u_0 + 2u_1 + 2u_2 + 2u_3 + 2u_4 + 2u_5 + 2u_6 + 2u_7 - 1, \\ u_0^2 - u_0 + 2u_1^2 + 2u_2^2 + 2u_3^2 + 2u_4^2 + 2u_5^2 + 2u_6^2 + 2u_7^2, \\ 2u_0u_1 + 2u_1u_2 - u_1 + 2u_2u_3 + 2u_3u_4 + 2u_4u_5 + 2u_5u_6 + 2u_6u_7, \\ 2u_0u_2 + u_1^2 + 2u_1u_3 + 2u_2u_4 - u_2 + 2u_3u_5 + 2u_4u_6 + 2u_5u_7, \\ 2u_0u_3 + 2u_1u_2 + 2u_1u_4 + 2u_2u_5 + 2u_3u_6 - u_3 + 2u_4u_7, \\ 2u_0u_4 + 2u_1u_3 + 2u_1u_5 + u_2^2 + 2u_2u_6 + 2u_3u_7 - u_4, \\ 2u_0u_5 + 2u_1u_4 + 2u_1u_6 + 2u_2u_3 + 2u_2u_7 - u_5, \\ 2u_0u_6 + 2u_1u_5 + 2u_1u_7 + 2u_2u_4 + u_3^2 - u_6 \end{pmatrix}$$

bjork 8: $[a, b, c, d, e, f, g, h]$

$$\begin{pmatrix} a + b + c + d + e + f + g + h, \\ ab + ah + bc + cd + de + ef + fg + gh, \\ abc + abh + agh + bcd + cde + def + fgh, \\ abcd + abch + abgh + afgh + bcde + cdef + defg + efgh, \\ abcde + abcdh + abcgh + abfgh + aefgh + bcdef + cdefg + defgh, \\ abcdef + abcdeh + abcdgh + abcfgh + abefgh + adefgh + bcdefg + cdefgh, \\ abcdefg + abcdefh + abcdegh + abcdfgh + abcefgh + abdefgh + acdefgh + bcdefgh, \\ a * (b + f) * (b + g) * (c + f) * (c + g) * (c + h) * (d + g) * (d + h) - 1 \end{pmatrix}$$

homog alex 3: $[t,x,y,z,w]$

$$\begin{pmatrix} t^5x^2yw^6+5t^3x^8y^2z+5tw^{13}+2x^2yw^{11}, \\ 4t^6xw^2+9t^3y^2zw^3+6x^2yw^6+7xy^6z^2, \\ 4t^6zw^2+2t^3x^3y^2w+8tx^5y^3+3xw^8 \end{pmatrix}$$

alex 3: $[t,x,y,z,w]$

$$\begin{pmatrix} t^5x^2y+5t^3x^8y^2z+5t+2x^2y, \\ 4t^6x+9t^3y^2z+6x^2y+7xy^6z^2, \\ 4t^6z+2t^3x^3y^2+8tx^5y^3+3x \end{pmatrix}$$

gerhard 1: $[w,t,x,y,z]$

$$\begin{pmatrix} 5w^4t^3x^2z+2t^2x^5y^3, \\ 7w^2y+2wtz+4x^2y+xy^2, \\ 3w^3tz+3w^2yz^2+2yz^4 \end{pmatrix}$$

symmetric 6: $[a,b,c,d,e]$

$$\begin{pmatrix} a^6-b^6, \\ b^6-c^6, \\ c^6-d^6, \\ d^6-e^6, \\ a^5b+ae^5+b^5c+c^5d+d^5e \end{pmatrix}$$

homog alex 2: $[t,x,y,z,w]$

$$\begin{pmatrix} t^5x^2yw^3+2t^3x^5y^2z+2x^2yw^8, \\ 3t^6x^2yz^2+9t^3y^2zw^5+2xy^3z^2w^5, \\ 2t^5zw^6+t^3x^5y^2z^2+tx^2y^3w^6 \end{pmatrix}$$

cohn2: $[t,x,y,z]$

$$\begin{pmatrix} -x^3y^2+2x^2y^2z-144x^2y^2-x^2yz^2-207x^2yz-3456x^2y+288xy^2z \\ \qquad -5184xy^2+78xyz^2-9504xyz+7192xy+xz^3-432xz^2-1798xz-9705x, \\ -t^2x^3z+32t^2x^3+5t^2x^2z^2-72t^2x^2z-87t^2xz^2+2592t^2xz-t^2z^3+864t^2z^2+5394t^2z \\ \qquad -6tx^3z-432tx^2z-414txz^2+11267txz+3456tz^2+11149tz-8x^3z-1728x^2z+3596xz-9705z, \\ t^3x^2y-2t^3xy^2-24t^3xy+128t^3x+t^3y^3+24t^3y^2+8t^2x^2y-12t^2xy^2-160t^2xy+2304t^2x+4t^2y^3 \\ \qquad +96t^2y^2+20tx^2y-20txy^2+96txy+13824tx+16x^2y+1152xy-4355x, \\ t^3y^3-t^3y^2z+72t^3y^2+71t^3yz+1728t^3y-464t^3z+13824t^3+4t^2y^3-2t^2y^2z+288t^2y^2+360t^2yz \\ \qquad +6912t^2y+6t^2z^2-4320t^2z-8710t^2+432tyz+8tz^2-13824tz+z^2-13824z \end{pmatrix}$$

alex 2: $[t,x,y,z]$

$$\begin{pmatrix} t^5x^2y+2t^3x^5y^2z+2x^2y, \\ 3t^6x^2yz^2+9t^3y^2z+2xy^3z^2, \\ 2t^5z+t^3x^5y^2z^2+tx^2y^3 \end{pmatrix}$$

gerhard 2: $[w,x,y,z]$

$$\begin{pmatrix} 3w^6x^2y+3w^6x^2z+4w^6xy^2+10w^6xyz+4w^6xz^2+w^6y^3 \\ \qquad +5w^6y^2z+5w^6yz^2+w^6z^3+3w^5x^2yz+4w^5xy^2z+4w^5xyz^2+w^5y^3z+2w^5y^2z^2+w^5yz^3+10x^9, \\ w^6x^3+4w^6x^2y+5w^6x^2z+3w^6xy^2+10w^6xyz+5w^6xz^2 \\ \qquad +3w^6y^2z+4w^6yz^2+w^6z^3+w^5x^3z+4w^5x^2yz+2w^5x^2z^2+3w^5xy^2z+4w^5xyz^2+w^5xz^3+10y^9, \\ w^6x^3+5w^6x^2y+4w^6x^2z+5w^6xy^2+10w^6xyz+3w^6xz^2 \\ \qquad +w^6y^3+4w^6y^2z+3w^6yz^2+w^5x^3y+2w^5x^2y^2+4w^5x^2yz+w^5xy^3+4w^5xy^2z+3w^5xyz^2+10z^9 \end{pmatrix}$$

gerhard 3: $[w,x,y,z]$

$$\begin{pmatrix} 6w^{12}x^5y^6+8w^{11}x^7y^5+10w^{11}x^4y^8+5w^{11}x^4y^5z^3+4w^{10}x^3y^{10}+24x^{23}, \\ 6w^{11}x^6y^5+5w^{10}x^8y^4+16w^{10}x^5y^7+5w^{10}x^5y^4z^3+10w^9x^4y^9+23y^{22}, \\ 3w^5x^5y^5z^2+18z^{17} \end{pmatrix}$$

Appendix B: Timings on other machines

Example	$t_{n,s}$	%mon	%comp	%div	%add	$\frac{t_{n,s}}{t_s}$	$\frac{t_{n,c}}{t_c}$	$\frac{t_s}{t_c}$	$\frac{t_{1.0}}{t_{n,s}}$	$\frac{t_{1.0}}{t_c}$
homog alex 2	10.6	74.0	62.5	5.9	5.6	1.5	1.3	0.9	3.3	4.6
ecyclic 6	3.9	72.4	66.9	2.0	3.5	1.5	1.8	1.2	1.3	2.3
katsura 7	6.9	69.3	47.4	4.2	17.7	1.4	1.8	1.3	2.0	3.6
katsura 8	89.4	71.1	53.3	3.5	14.2	1.4	1.6	1.2	1.9	3.3
alex 3	1.7	68.0	57.7	8.6	1.7	1.4	1.3	1.0	3.2	4.3
gerhard 1	1.4	68.5	55.1	6.0	7.4	1.4	1.5	1.1	2.9	4.2
cohn2	9.3	66.1	37.2	5.0	23.9	1.4	1.6	1.2	1.7	2.8
2mat3	161.0	79.7	76.9	1.5	1.4	1.3	1.4	1.4	1.7	3.1
alex 2	5.9	69.2	59.9	3.6	5.7	1.3	1.6	1.4	3.1	5.6
homog gonnet	154.0	79.8	49.5	29.1	1.2	1.3	1.5	1.7	1.6	3.5
homog 2mat3	158.7	81.3	78.2	1.6	1.4	1.3	1.5	1.5	1.8	3.5
rcyclic 14	20.4	73.7	27.2	36.5	10.0	1.3	1.5	1.4	1.5	2.7
ecyclic 7	1260.3	85.4	83.8	0.3	1.2	1.3	1.4	1.2	1.3	1.9
homog alex 3	1.8	66.5	51.8	12.3	2.3	1.3	1.1	0.9	2.9	3.3
bjork 8	3.9	57.6	33.0	11.6	13.0	1.3	1.6	1.5	1.6	2.9
rcyclic 12	3.3	69.6	28.2	30.4	11.0	1.2	1.7	1.3	1.7	2.7
homog cyclic 7	147.8	74.9	36.1	22.0	16.7	1.2	1.5	1.3	1.5	2.5
rcyclic 13	8.0	70.1	30.9	30.2	9.1	1.2	1.4	1.1	1.7	2.3
homog cyclic 6	0.7	63.5	28.8	22.0	12.7	1.2	1.7	1.4	1.7	2.9
gerhard 2	11.0	71.4	58.8	3.0	9.6	1.2	1.9	1.6	3.4	6.6
schwarz 10	24.9	56.0	29.9	17.0	9.2	1.2	1.3	1.3	1.4	2.2
rcyclic 16	112.0	77.4	29.1	36.8	11.5	1.2	1.6	1.5	1.4	2.6
schwarz 9	3.1	58.4	31.4	21.1	5.9	1.2	1.3	1.1	1.5	1.9
cyclic 7	226.8	59.8	30.2	13.2	16.4	1.2	1.4	1.3	1.5	2.3
rcyclic 17	233.6	78.0	29.7	36.7	11.6	1.2	1.4	1.3	1.5	2.3
rcyclic 11	1.3	64.4	25.3	28.4	10.7	1.2	1.5	1.3	1.6	2.4
rcyclic 15	50.0	74.1	27.6	35.8	10.8	1.2	1.6	1.6	1.5	2.9
schwarz 8	0.6	50.0	30.2	12.3	7.6	1.2	1.4	1.2	1.4	2.0
rcyclic 19	916.6	79.8	29.4	38.9	11.6	1.1	1.5	1.5	1.4	2.4
gerhard 3	28.6	61.9	43.8	12.3	5.8	1.1	1.4	1.2	2.4	3.4
rcyclic 18	477.2	79.6	28.6	39.4	11.6	1.1	1.4	1.5	1.4	2.4
symmetric 6	74.0	70.8	28.2	24.0	18.6	1.1	1.3	1.2	1.5	2.0
gonnet	1.1	56.7	12.7	40.8	3.2	1.1	1.1	1.1	1.1	1.4
schwarz 11	178.2	60.1	34.4	16.8	9.0	1.1	1.3	1.4	1.4	2.1
averages										
HP C160	125.4	69.0	41.7	17.9	9.5	1.2	1.5	1.3	1.8	2.9
Pentium Pro	103.3	69.8	38.4	15.6	15.8	1.4	1.6	1.2	1.2	2.2
DEC Alpha	221.4	69.8	37.7	16.1	16.0	1.7	1.7	1.2	1.1	2.3

Table 4: Detailed timings for vectorized monomial operations for a HP C160 (32 bit, big–endian) running HP-UX 10.20

Example	$t_{n,s}$	%mon	%comp	%div	%add	$\frac{t_{n,s}}{t_s}$	$\frac{t_{n,c}}{t_c}$	$\frac{t_s}{t_c}$	$\frac{t_{1.0}}{t_{n,s}}$	$\frac{t_{1.0}}{t_c}$
rcyclic 12	10.1	74.8	28.2	28.1	18.5	2.1	2.2	1.3	1.1	3.0
ecyclic 6	9.4	69.1	60.7	1.6	6.8	2.0	2.2	1.3	1.2	3.2
rcyclic 11	3.8	68.4	26.1	24.8	17.4	2.0	2.0	1.2	1.1	2.7
rcyclic 15	124.6	80.9	30.3	31.0	19.6	2.0	2.4	1.4	1.2	3.2
homog cyclic 6	2.0	69.9	30.2	15.6	24.1	2.0	2.1	1.2	1.1	2.7
rcyclic 19	1838.3	83.8	30.6	31.8	21.4	2.0	2.0	1.2	1.1	2.6
rcyclic 14	55.0	79.7	29.5	31.1	19.2	1.9	2.3	1.4	1.2	3.1
rcyclic 13	24.9	77.9	30.3	28.1	19.5	1.9	2.3	1.4	1.1	3.1
rcyclic 18	986.4	84.5	30.7	34.6	19.2	1.9	1.9	1.2	1.1	2.6
rcyclic 16	256.6	82.3	30.5	32.1	19.8	1.9	2.0	1.2	1.2	2.7
katsura 7	21.1	71.1	41.4	3.9	25.9	1.9	1.6	1.0	1.1	2.1
homog gonnet	310.1	81.1	44.9	33.9	2.3	1.9	1.9	1.1	1.2	2.4
bjork 8	10.1	69.0	33.0	12.6	23.4	1.9	1.8	1.1	1.0	2.1
homog cyclic 7	300.2	77.9	37.4	13.2	27.4	1.8	1.8	1.1	1.2	2.4
rcyclic 17	527.1	83.2	30.4	33.3	19.5	1.8	1.9	1.2	1.1	2.5
cyclic 7	458.9	68.6	32.5	10.3	25.7	1.7	1.7	1.2	1.1	2.1
cyclic 6	1.7	61.4	25.3	13.5	22.6	1.7	1.7	1.2	1.0	2.0
ecyclic 7	1805.0	60.7	55.9	0.6	4.2	1.7	1.8	1.1	1.2	2.4
homog 2mat3	256.8	78.3	74.1	1.8	2.5	1.7	1.8	1.2	1.4	2.6
gerhard 1	4.0	71.2	50.5	5.4	15.2	1.6	1.6	1.2	1.3	2.4
katsura 8	229.9	73.3	46.6	3.0	23.7	1.6	1.7	1.2	1.1	2.2
gonnet	2.8	51.7	8.2	40.4	3.1	1.6	1.6	1.1	1.0	1.7
cohn2	24.6	65.9	30.3	3.2	32.4	1.6	1.5	1.2	1.0	1.9
schwarz 8	1.3	50.0	19.6	14.9	15.4	1.6	1.5	1.0	1.0	1.7
schwarz 11	317.1	65.4	32.4	16.5	16.4	1.6	1.6	1.1	1.1	1.9
schwarz 10	51.2	62.9	29.4	17.3	16.2	1.5	1.6	1.1	1.1	1.9
symmetric 6	147.2	72.4	29.1	13.8	29.5	1.5	1.5	1.2	1.1	2.1
homog alex 2	32.1	74.8	60.6	4.4	9.7	1.5	1.5	1.2	1.3	2.3
schwarz 9	7.5	57.9	23.6	17.8	16.5	1.5	1.6	1.1	1.1	1.9
gerhard 2	33.2	72.7	56.2	2.1	14.3	1.5	1.4	1.2	1.3	2.3
2mat3	248.0	70.5	66.4	1.6	2.4	1.4	1.5	1.2	1.2	2.0
alex 2	17.4	65.1	52.6	2.9	9.6	1.4	1.2	1.1	1.1	1.7
homog alex 3	5.2	65.5	48.7	12.5	4.3	1.3	1.3	1.1	1.2	1.8
gerhard 3	61.4	66.6	46.7	9.7	10.2	1.3	1.3	1.2	1.2	1.9
alex 3	4.4	60.8	50.1	6.3	4.5	1.2	1.2	1.1	1.2	1.7
averages										
DEC Alpha	221.4	69.8	37.7	16.1	16.0	1.7	1.7	1.2	1.1	2.3
Pentium Pro	103.3	69.8	38.4	15.6	15.8	1.4	1.6	1.2	1.2	2.2
HP C160	125.4	69.0	41.7	17.9	9.5	1.2	1.5	1.3	1.8	2.9

Table 5: Detailed timings for vectorized monomial operations for DEC Alpha (64 bit, little–endian) running Linux.

exponent type monomial operations	char vectorized				short vectorized				short not vectorized			
Example	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$
2mat3	1.3	2.6	75.5	1.0	1.2	2.1	75.3	1.1	1.2	1.8	78.1	1.4
homog alex 2	1.4	6.8	56.1	1.0	1.7	6.4	58.9	0.9	1.9	7.9	61.4	1.3
alex 3	1.5	3.3	48.2	1.0	1.4	1.4	55.9	0.9	1.5	3.1	58.0	1.3
alex 2	1.6	5.7	56.0	0.8	1.5	5.6	60.2	1.0	1.4	5.9	57.7	1.2
katsura 8	1.6	13.4	48.6	0.9	1.6	12.5	50.5	1.0	2.0	13.7	54.4	1.2
homog alex 3	1.9	2.6	49.9	1.1	2.1	1.8	50.9	1.0	1.9	4.2	46.6	1.2
homog gonnet	1.7	2.5	42.6	0.9	1.4	1.6	45.2	1.1	1.6	1.6	48.4	1.2
gerhard 1	1.7	7.7	55.5	0.9	1.5	9.5	51.9	0.9	2.3	12.9	56.6	1.2
katsura 7	1.6	17.6	39.6	0.8	1.9	16.2	49.1	0.9	1.8	15.9	48.2	1.2
gerhard 2	1.4	8.2	55.0	0.7	1.7	9.7	61.4	1.0	1.6	9.3	57.5	1.1
homog 2mat3	1.3	2.7	76.8	0.9	1.0	1.7	76.6	1.0	1.2	1.9	78.0	1.1
cyclic 7	1.3	10.7	25.0	0.9	1.5	10.6	25.4	0.9	1.8	12.1	30.1	1.1
gerhard 3	1.4	4.8	36.2	0.9	1.7	5.4	39.4	1.0	1.6	5.5	40.7	1.1
schwarz 11	1.2	7.8	30.4	0.9	1.3	6.4	30.5	1.0	1.5	7.2	34.0	1.1
bjork 8	1.1	10.6	23.2	0.8	1.7	10.5	33.7	0.9	2.2	12.5	32.8	1.0
schwarz 10	1.5	8.8	26.7	0.9	1.3	7.1	24.7	0.9	1.6	9.3	30.7	1.0
homog cyclic 7	1.4	11.9	27.7	0.8	1.8	11.2	29.7	0.9	2.0	12.5	34.3	1.0
cohn2	1.2	12.6	28.9	0.7	1.6	13.7	31.8	0.8	1.9	15.8	35.1	1.0
rcyclic 13	2.0	14.1	24.0	0.8	1.8	10.9	25.7	0.9	2.9	12.0	29.5	1.0
rcyclic 12	1.6	11.3	18.8	0.8	1.9	9.7	22.6	0.8	2.4	11.4	28.4	1.0
rcyclic 14	2.0	12.0	23.4	0.8	1.6	10.8	20.9	0.8	2.5	10.3	28.7	1.0
rcyclic 15	1.4	12.1	23.1	0.8	1.5	9.3	23.3	0.9	1.5	9.8	28.3	1.0
schwarz 9	1.7	11.7	23.5	0.9	1.6	13.5	23.4	0.9	2.5	12.0	28.4	1.0
symmetric 6	1.3	10.9	21.8	0.8	1.4	11.3	24.9	0.9	1.8	13.6	28.9	1.0
homog cyclic 6	1.0	9.6	24.5	0.7	2.3	16.9	29.1	0.8	2.6	17.5	25.4	1.0
schwarz 8	1.4	10.9	15.7	0.8	1.9	9.5	23.3	0.9	2.9	5.9	21.7	1.0
rcyclic 11	2.0	13.0	26.9	0.8	2.2	7.2	22.3	0.8	2.6	14.9	29.0	1.0
gonnet	1.0	5.0	3.5	0.9	2.2	4.0	9.6	0.9	1.0	4.1	6.4	1.0
cyclic 6	0.8	13.8	16.5	0.8	2.5	15.5	24.1	0.8	1.7	17.0	32.4	1.0
ecyclic 7	1.0	2.5	85.3	0.9	1.0	1.8	83.9	0.9	0.8	1.2	84.8	0.9
rcyclic 19	1.1	9.6	23.7	0.8	1.0	7.3	24.2	0.8	0.8	6.5	29.5	0.9
ecyclic 6	1.1	16.7	61.7	0.7	1.1	11.2	62.2	0.8	0.6	9.2	67.8	0.8
rcyclic 18	0.9	9.0	22.7	0.9	0.9	7.8	23.6	0.8	0.9	7.0	28.3	0.8
rcyclic 16	1.1	11.1	23.2	0.7	1.2	9.4	24.1	0.8	0.8	8.0	29.0	0.8
rcyclic 17	1.0	9.4	22.1	0.8	0.9	8.1	24.9	0.8	1.1	8.1	30.1	0.8
averages												
HP C160	1.4	9.1	36.6	0.8	1.6	8.4	39.0	0.9	1.7	9.1	41.6	1.1
Pentium Pro	1.3	17.1	33.6	0.8	1.4	13.8	33.8	0.9	2.1	12.6	40.3	1.1
DEC Alpha	1.8	11.8	33.1	0.8	1.5	11.7	31.5	0.8	2.4	10.2	35.5	1.0

Table 6: Detailed timings for rank-based monomial operations on an HP C160 (32 bit, big–endian) running HP-UX 10.20.

exponent type monomial operations	char vectorized				short vectorized				short not vectorized			
Example	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$	$\frac{t_<}{t_{r,<}}$	%R	%comp	$\frac{t}{t_r}$
homog alex 2	2.4	9.7	60.7	1.0	2.0	9.3	56.4	1.0	2.6	8.8	59.4	1.4
homog alex 1	2.2	5.7	55.9	1.3	1.6	5.5	44.0	1.1	2.2	5.6	48.6	1.3
homog gonnet	1.3	5.0	41.1	1.1	1.7	5.2	45.6	1.1	2.9	4.4	44.7	1.3
alex 1	1.7	2.4	65.2	1.3	1.5	3.2	62.8	1.3	1.4	2.7	59.9	1.3
homog 2mat3	1.7	2.4	65.2	1.3	1.5	3.2	62.8	1.3	1.4	2.7	59.9	1.3
gerhard 2	2.1	10.8	57.8	0.9	1.7	10.1	53.0	1.0	2.1	10.2	53.6	1.3
alex 2	1.6	8.6	56.6	1.1	1.4	7.3	54.6	1.2	1.6	6.9	56.1	1.3
gerhard 1	2.4	13.4	51.3	0.8	1.8	12.8	45.1	0.9	2.3	11.8	48.1	1.3
gerhard 3	1.9	7.4	44.1	1.0	1.5	7.2	40.8	1.0	1.9	7.0	44.3	1.2
homog alex 3	1.9	4.6	50.2	1.0	1.6	5.0	46.2	1.1	2.0	5.2	49.1	1.2
alex 3	1.4	5.4	53.5	1.0	1.3	5.3	51.8	1.0	1.7	5.5	55.5	1.2
katsura 8	1.7	19.8	47.1	0.8	1.7	17.9	44.8	0.9	2.9	17.4	46.9	1.2
katsura 7	1.7	20.1	39.4	0.8	1.7	19.0	36.8	0.8	2.8	17.2	41.0	1.1
cyclic 7	2.0	17.1	27.6	0.8	1.5	17.1	26.5	0.8	2.7	15.0	32.8	1.1
bjork 8	1.8	17.8	26.0	0.8	1.6	16.7	25.8	0.8	3.0	16.4	32.1	1.1
homog cyclic 7	1.7	17.1	29.7	0.7	1.8	18.2	29.2	0.7	3.0	16.2	35.0	1.1
cohn2	2.4	19.7	30.4	0.8	1.4	15.8	24.5	0.9	2.1	14.5	29.9	1.1
schwarz 11	1.7	14.7	31.0	0.8	1.6	16.0	30.5	0.9	2.4	13.3	32.9	1.0
rcyclic 14	1.7	20.7	24.8	0.6	2.0	21.9	24.9	0.7	3.8	15.7	28.7	1.0
rcyclic 13	1.8	20.4	20.9	0.7	2.0	21.1	23.6	0.7	3.9	16.1	28.9	1.0
symmetric 6	2.0	16.4	26.1	0.7	1.5	16.1	21.6	0.8	2.4	14.3	27.1	1.0
cyclic 6	2.9	19.3	23.5	0.7	1.3	16.7	16.5	0.7	2.6	14.5	25.3	1.0
schwarz 10	1.7	14.3	25.9	0.8	1.6	17.1	25.3	0.9	2.4	13.4	29.7	1.0
schwarz 8	1.8	13.3	13.4	0.7	1.8	12.8	15.6	0.7	2.7	12.5	19.3	1.0
symmetric 5	2.1	13.8	18.5	0.8	1.3	11.4	11.8	0.8	2.2	12.2	17.5	1.0
homog cyclic 6	2.7	18.2	26.7	0.6	1.7	19.2	20.5	0.6	3.2	19.1	26.8	1.0
rcyclic 11	1.9	17.2	19.4	0.6	1.9	20.0	19.2	0.6	3.6	15.3	25.5	1.0
rcyclic 15	1.5	15.6	26.0	0.7	1.6	16.1	25.2	0.8	2.3	12.2	29.9	1.0
rcyclic 12	1.9	19.9	22.1	0.7	1.9	20.1	22.2	0.7	3.7	15.5	27.7	1.0
ecyclic 7	1.9	19.9	22.1	0.7	1.9	20.1	22.2	0.7	3.7	15.5	27.7	1.0
schwarz 9	1.6	15.2	19.7	0.8	1.6	14.0	18.9	0.8	2.8	13.7	23.8	0.9
gonnet	1.4	9.1	4.4	0.8	1.3	9.8	5.0	0.8	3.1	6.8	8.3	0.9
rcyclic 10	2.8	16.2	13.8	0.6	2.0	18.8	14.6	0.7	4.2	15.3	20.2	0.9
ecyclic 6	0.9	1.0	57.7	0.6	0.9	1.1	57.3	0.7	1.0	0.8	67.3	0.9
rcyclic 18	0.9	1.0	23.4	0.7	1.0	1.1	28.6	0.7	1.0	0.6	31.8	0.8
rcyclic 19	1.1	1.0	25.8	0.7	1.0	1.1	24.8	0.7	0.9	1.0	28.8	0.8
rcyclic 17	0.9	1.5	21.5	0.7	1.0	1.6	25.6	0.7	1.0	1.0	29.5	0.8
rcyclic 16	0.9	3.3	21.3	0.7	1.0	3.3	24.7	0.7	1.1	2.2	29.9	0.8
averages												
DEC Alpha	1.8	11.8	33.1	0.8	1.5	11.7	31.5	0.8	2.4	10.2	35.5	1.0
Pentium Pro	1.3	17.1	33.6	0.8	1.4	13.8	33.8	0.9	2.1	12.6	40.3	1.1
HP C160	1.4	9.1	36.6	0.8	1.6	8.4	39.0	0.9	1.7	9.1	41.6	1.1

Table 7: Detailed timings for rank-based monomial operations on a DEC Alpha (64 bit, little–endian) running Linux.

List of papers published in the Reports on Computer Algebra series

[1] H. Grassmann, G.-M. Greuel, B. Martin, W. Neumann, G. Pfister, W. Pohl, H. Schönemann, and T. Siebert. Standard bases, syzygies and their implementation in singular. 1996.

[2] H. Schönemann. Algorithms in singular. 1996.

[3] R. Stobbe. FACTORY: a C++ class library for multivariate polynomial arithmetic. 1996.

[4] O. Bachmann and H. Schönemann. A Manual for the MPP Dictionary and MPP Library. 1996.

[5] O. Bachmann, S. Gray, and H. Schönemann. MPP: A Framework for Distributed Polynomial Computations. 1996.

[6] G.M. Greuel and G. Pfister. Advances and improvements in the theory of standard bases and szyzygies. 1996.

[7] G.M. Greuel. Description of SINGULAR: A computer algebra system for singularity theory, algebraic geometry, and commutative algebra. 1996.

[8] T. Siebert. On strategies and implementations for computations of free resolutions. September 1996.

[9] B. Reinert. Introducing reduction to polycyclic group rings – a comparison of methods. October 1996.

[10] O. Bachmann, S. Gray, and H. Schönemann. A proposal for syntactic data integration for math protocols. January 1997.

[11] O. Bachmann. Effective simplification of cr expressions. January 1997.

[12] O. Bachmann, S. Gray, and H. Schönemann. MP Prototype Specification. January 1997.

[13] O. Bachmann. MPT – a library for parsing and Manipulating MP Trees. January 1997.

[14] K. Madlener and B. Reinert. Relating rewriting techniques on monoids and rings: Congruences on monoids and ideals in monoid rings. September 1997.

[15] B. Martin and T. Siebert. Splitting Algorithm for vector bundles. September 1997.

[16] K. Madlener and B. Reinert. String Rewriting and Gröbner Bases – A General Approach to Monoid and Group Rings. October 1997.

[17] Thomas Siebert. An algorithm for constructing isomorphisms of modules. January 1998.

[18] O. Bachmann and H. Schönemann. Monomial Representations for Gröbner Bases Computations. January 1998.

[19] B. Reinert, K. Madlener, and T. Mora. A note on nielsen reduction and coset enumeration. February 1998.