

```

gap> g:= SymmetricGroup( 4 );
Sym( [ 1 .. 4 ] )
gap> tbl:= CharacterTable( g );; HasIrr( tbl );
i5 : betti(t,Weights=>{1,0})
false
      0 1 2 3 4 gap> tblmod2:= CharacterTable( tbl, 2 );
o5 = total: 1 4 13 14 4 BrauerTable( Sym( [ 1 .. 4 ] ), 2 )
      0: 1 . . . .
      1: . 2 2 4 2 gap> tblmod2 = CharacterTable( tbl, 2 );
      2: . 2 5 6 . true
      3: . . 4 . 2
      4: . . . 4 . gap> tblmod2 = BrauerTable( tbl, 2 );
      5: . . 2 . . true
gap> tblmod2 = BrauerTable( tbl, 2 );
o5 : BrauerTable( Sym( [ 1 .. 4 ] ), 2 )
i6 : betti(t,Weights=>{0,1})
      0 1 2 3 4 gap> libtbl:= CharacterTable( "M" );
o6 = total: 1 4 13 14 4 CharacterTable( "M" )
      0: 1 . . . . gap> CharacterTableRegular( libtbl, 2 );
      1: . 2 2 . 2 BrauerTable( "M", 2 );
      2: . 2 5 6 . BrauerTable( libtbl, 2 );
      3: . . 4 . 2 gap> BrauerTable( libtbl, 2 );
      4: . . . 4 . fail
      5: . . 2 . .
gap> CharacterTable( "Symmetric", 4 );
o6 : BettiTally CharacterTable( "Sym(4)" )
i7 : t1 = betti(t,Weights=>{1,1})
gap> ComputedBrauerTables( tbl );
      0 1 2 3 4 [ , BrauerTable( Sym( [ 1 .. 4 ] ), 2 ) ]
o7 = total: 1 4 13 14 4
      0: 1 . . . .
      1: . . . . .
      2: . . . . .
      3: . 2 . . .
      4: . . . . .
      5: . 2 . . .
      6: . . 1 . .
      7: . . 8 6 .
      8: . . 4 8 4
ring r1 = 32003,(x,y,z),ds;
int a,b,c,t=11,5,3,0;
poly f = x^a+y^b+z^(3*c)+x^(c+2)*y^(c-1)+x^(c-2)*y^c*(y^2+t*x)^2;
option(noprot);
timer=1;
ring r2 = 32003,(x,y,z),dp;
poly f=imap(r1,f);
ideal j=jacob(f);
vdim(std(j));
==> 536
vdim(std(j+f));
==> 195
timer=0; // reset timer
o7 : BettiTally
i8 : peek t1
o8 = BettiTally{(0, {0, 0}, 0) => 1 }
      (1, {2, 2}, 4) => 2
      (1, {3, 3}, 6) => 2
      (2, {3, 7}, 10) => 2
      (2, {4, 4}, 8) => 1
      (2, {4, 5}, 9) => 4
      (2, {5, 4}, 9) => 4
      (2, {7, 3}, 10) => 2
      (3, {4, 7}, 11) => 4
      (3, {5, 5}, 10) => 6
      (3, {7, 3}, 10) => 2
      (4, {5, 7}, 12) => 2
      (4, {7, 5}, 12) => 2

```

Journal of Software for Algebra and Geometry

Threaded Gröbner bases: a Macaulay2 package

SONJA PETROVIĆ AND SHAHRZAD ZELENBERG

Threaded Gröbner bases: a Macaulay2 package

SONJA PETROVIĆ AND SHAHRZAD ZELENBERG

ABSTRACT: The complexity of Gröbner computations has inspired many improvements to Buchberger’s algorithm over the years. Looking for further insights into the algorithm’s performance, we offer a threaded implementation of the classical Buchberger algorithm in Macaulay2. The output of the main function of the package includes information about *lineages* of nonzero remainders that are added to the basis during the computation. This information can be used for further algorithm improvements and optimization.

1. INTRODUCTION. The importance in computational algebra of Gröbner bases and therefore of Buchberger’s algorithm, as well as its many variants, is indisputable. Yet it is still a challenge to apply brute force algorithms to larger problems primarily due to considerations in computer science. That is, the current computing paradigm favors clusters of CPUs, or nodes, rather than one massive CPU. As a result, there is a need to distribute this algorithm that is automated for the user (in that it does not require a user to know how it should be distributed).

Past work in this area has focused on synchronized methods as detailed in [Mityunin and Pankrat’ev 2005]. One method spreads a key step in Buchberger’s algorithm — the reduction of S-pairs by division — across nodes; another sends tasks to individual nodes while a central, coordinating node waits for all threads to complete. Each of these still requires some central node and synchronization, which leads to bottlenecks in the computation. A truly distributed algorithm would be decentralized and asynchronous. Zelenberg [2018] discusses an asynchronous, decentralized distributed version of Buchberger’s algorithm done generically with the potential of very good speedups. Zelenberg implemented a threaded version in [Python] to explore this further, and as a result some important discoveries were made. It should be noted that multithreaded algorithms are not necessarily distributed across distinct nodes; rather, threads are sharing computation and passing information back and forth.

Of the discoveries made, the most important is this: in order for a distributed process to be both generically usable and automated for the user, an effective algorithm will need to account for features of the polynomials (relative to the starting basis) when deciding what tasks to assign to what node. This is because transferring information between nodes is a very slow process and so needs to be minimized.

Given the need to analyze aspects of these polynomials, [Macaulay2] offers some clear advantages over *Python*. Moreover, the *Macaulay2* engine is written in C/C++, a language well suited for writing distributed algorithms. Threads within *Macaulay2* work differently than within *Python* and, as such,

MSC2020: 13P10.

Keywords: Gröbner basis, distributed computation, lineages, S-pairs.

ThreadedGB version 1.1

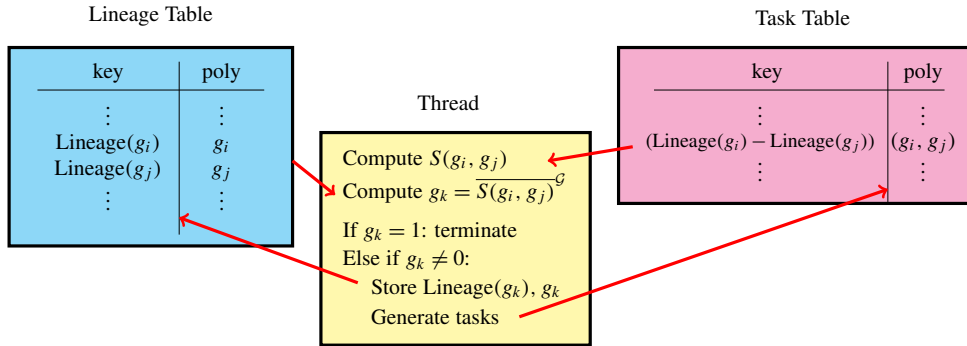


Figure 1. The lineage table is initialized with the starting basis, and the task table is initialized with a task for each pair of starting generators. Then, all threads perform the same task: they pull a pair of polynomials from the tasks hash table along with an associated lineage key. They compute the S-polynomial and reduce it with respect to the current basis. If the remainder, r , is nonzero, it is stored in the lineage table along with its lineage key, and, for each g in the current basis, a new task indexed by the pair (g, r) is added to the task table. If the remainder 1 is found, the process of creating tasks stops. The process is repeated using n parallel threads, with n specified by the user, until the task table is empty.

some design changes were necessary from the implementation discussed in [Zelenberg 2018]. Queues are replaced altogether with hash tables — an improvement since threads access the most up-to-date version of the generating set at the time of reduction. Even so, this cannot eliminate redundancies as threads may compute the same result (virtually) simultaneously.

One of the goals of our package *ThreadedGB* is to allow a user to analyze what we refer to as *lineages* of polynomials in a Gröbner basis.

Definition 1.1. Let G be a Gröbner basis of $I = (f_0, \dots, f_k)$. A *lineage* of a polynomial in G is a natural number, or an ordered pair of lineages, tracing its history in the given Gröbner basis computation. It is defined recursively as follows:

- For the starting generating set, $\text{Lineage}(f_i) = i$,
- For any subsequently created S-polynomial $S(f, g)$, the lineage of its remainder r on division is the pair $\text{Lineage}(r) = (\text{Lineage}(f), \text{Lineage}(g))$.

To illustrate, suppose $I = (x^2 - y, x^3 - z) \subset \mathbb{Q}[x, y, z]$ with graded reverse lexicographic order. Then $\text{Lineage}(x^2 - y) = 0$ and $\text{Lineage}(x^3 - z) = 1$. Two additional elements are added to create a (nonminimal) Gröbner basis: $xy + z$ and $y^2 - xz$, with lineages $(0, 1)$ and $((0, 1), 0)$, respectively. According to $\text{Lineage}(y^2 - xz)$, this element is constructed from $S(xy + z, x^2 - y)$. Lineages are expressions of the starting generating set and thus dependent on the choice and order of its elements. More importantly, a lineage is not necessarily unique, as the same polynomial can be constructed multiple ways. The lineage tables produced by *ThreadedGB* (see Figure 1) do not provide all possible lineages — only a particular choice based on the order in which the generators are provided by the user.

2. EFFECT OF ORDERING OF POLYNOMIALS ON LINEAGES: A SIMPLE EXAMPLE. Consider the polynomial ring $\mathbb{Q}[x_1, x_3, x_0, x_4, x_2]$ with lexicographic order and the ideal of the rational normal curve in \mathbb{P}^4 . The six generators have lineages $0, \dots, 5$, and Buchberger's algorithm adds three new elements to the Gröbner basis before final reduction. This can be seen by turning on the `gbTrace` option in *Macaulay2*, which tells us three new polynomials are added to the basis. The function `tgb` lets us know exactly which ones and their lineages. Specifically, a run of `tgb` reveals these are $x_0x_4 - x_2^2$, $-x_3x_0x_4 + x_3x_2^2$, $-x_0x_4x_2 + x_2^3$ with lineages $(2, 3)$, $(1, 4)$, $(1, 2)$, respectively.

```
i1 : needsPackage "ThreadedGB";
i2 : QQ[x_1,x_3,x_0,x_4,x_2,MonomialOrder=>Lex];
i3 : rnc = minors(2, matrix{{x_0..x_3},{x_1..x_4}});
o3 : Ideal of QQ[x_1, x_3, x_0, x_4, x_2]
i4 : allowableThreads = 4;
i5: g = tgb(rnc)
o5 = LineageTable{(1, 2) => - x x x + x3
                    0 4 2    2
(1, 4) => - x x x + x x2
                    3 0 4    3 2
(2, 3) => x x - x2
                    0 4    2
0 => - x2 + x x
1 => - x1x2 + x3x0
2 => x x - x2
1 3    2
3 => - x x + x x
1 3    0 4
4 => x x - x x
1 4    3 2
5 => - x2 + x x
3    4 2
o5 : LineageTable
```

Running the command `reduce g` will produce a reduced Gröbner basis; in particular, the lineage table entries with keys $(1, 2)$, $(1, 4)$ and 2 will be replaced by `null`. This allows the user to see which nonzero polynomials produced during the computation turn out not to be needed. Of course, to continue computing with the given basis, one wishes to have it in standard *Macaulay2* format, which is a matrix.

```
i6 : matrix reduce g
o6 = | x_1^2-x_0x_2 x_1x_2-x_3x_0 x_1x_3-x_2^2
      x_1x_4-x_3x_2 x_3^2-x_4x_2 x_0x_4-x_2^2 |
o6 : Matrix (QQ[x_1, x_3, x_0, x_4, x_2]) <--- (QQ[x_1, x_3, x_0, x_4, x_2])6
```

One can use the package to study, for example, how reordering the input basis affects the algorithm. In Gröbner computations, *Macaulay2* creates and processes S-polynomials in lexicographic order of pairs (first and second, then first and third, and so on). Let $S = \mathbb{Q}[a, b, c, d]$ and $I = (abc - 1, abc, a + bd - c)$; clearly $I = S$. But the order of generators listed affects the complexity of the particular run; namely, listing the quadratic first makes the algorithm perform more steps. The method `tgb` can be verbose and

can tell us what is going on behind the scenes for each lineage.

```
i7 : QQ[a, b, c, d];
i8 : I = ideal (a*b*c, a*b*c - 1, a+b*d-c);
i9 : tgb(I,Verbose=>true)
Scheduling a task for lineage (0,1)
Scheduling a task for lineage (0,2)
Scheduling a task for lineage (1,2)
Adding the following remainder to GB: 1 from lineage (0,1)
Found a unit in the Groebner basis; reducing now.
o9 = LineageTable{(0, 1) => 1}
      0 => null
      1 => null
      2 => null
o9 : LineageTable
```

Compare this to the following run of the threaded Buchberger’s algorithm under a different input order.

```
i10 : I = ideal (a+b*d-c, a*b*c-1, a*b*c);
o10 : Ideal of QQ[a, b, c, d]
i11 : tgb(I)
o11 = LineageTable{(0, 1) => null}
      (0, 2) => null
      (1, 2) => 1
      0 => null
      1 => null
      2 => null
o11 : LineageTable
```

Three new elements are added to the basis, namely (0,1), (0,2), (1,2), if the quadratic generator is listed first, but if it is listed last, then only the polynomial with lineage (0,1) is added—because it already equals 1—and the algorithm stops.

3. NUTS AND BOLTS. Given a list of polynomials L or an ideal I and an integer n , the main method `tgb` uses `Tasks` in *Macaulay2* to compute a Gröbner basis of I or (L) using n threads. It returns an object of type `LineageTable`, which is an instance of `HashTable`, whose values are a Gröbner basis of I or (L) . The keys are polynomial lineages.

The starting basis L (meaning, the input list L or $L=I_*$) populates the entries of a lineage table G with keys from 0 to one less than the number of elements of L . The method creates all possible S-polynomials of L and schedules their reduction with respect to G as tasks. Throughout the computation, every nonzero remainder added to the basis is added to G , with its lineage, as defined above, being the key. Each such remainder also triggers the creation of S-polynomials using it and every element in G and scheduling the reduction thereof as additional tasks. The process is done when there are no remaining tasks.

There is a way to track the tasks being created by turning on the option `Verbose`, or provide the reduced or a minimal Gröbner basis using the functions `reduce` or `minimalize`, respectively. The users who expect just a Gröbner basis in usual *Macaulay2* format, without the lineages, can call `matrix` on the `LineageTable`.

4. IMPROVEMENTS AND SPEED-UPS. As with any *Macaulay2* package, improvements are easy to make via GitHub. Our package’s GitHub repository will be made public shortly, so other users can implement any extensions or add improvements to this threaded implementation of Buchberger’s algorithm.

These may include known speed-ups as optional ways to run the algorithm; for example, if one wishes to study lineages produced by the F4 algorithm [Faugère 1999], then one can build that option into this threaded computation.

The current goal is to explore algorithm performance and complexity and how input basis features affect these; the lineages are designed specifically to aid in this goal. Of course, speed-ups should come “naturally” from a threaded implementation. However, with *Macaulay2*’s current implementation of threads, speed-ups aren’t observed with interpreted code, hence to achieve effective speed-ups in practice, we plan to implement `tgb` in the engine, using C/C++.

SUPPLEMENT. The online supplement contains version 1.1 of *ThreadedGB*.

REFERENCES.

- [Faugère 1999] J.-C. Faugère, “A new efficient algorithm for computing Gröbner bases (F_4)”, pp. 61–88 in *Effective methods in algebraic geometry* ((Saint-Malo, 1998)), vol. 139, 1999. MR Zbl
- [Macaulay2] D. R. Grayson and M. E. Stillman, “Macaulay2, a software system for research in algebraic geometry”, available at <https://math.uiuc.edu/Macaulay2/>.
- [Mityunin and Pankrat’ev 2005] V. A. Mityunin and E. V. Pankrat’ev, “Parallel algorithms for the construction of Gröbner bases”, *Sovrem. Mat. Prilozh.* 30 (2005), 46–64. Translated in *J. of Math. Statistics*, **142**(4): 2248–2266, 2007. MR
- [Python] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA.
- [Zelenberg 2018] S. J. Zelenberg, *Distributed computational systems*, Ph.D. thesis, 2018, available at etda.libraries.psu.edu/catalog/15329sxj937.

RECEIVED: 19 Jul 2020

REVISED: 6 Aug 2021

ACCEPTED: 8 Oct 2021

SONJA PETROVIĆ:

sonja.petrovic@iit.edu

Department of Applied Mathematics, Illinois Institute of Technology, Chicago, IL, United States

SHAHRZAD ZELENBERG:

szelenberg@mx.lakeforest.edu

Lake Forest College, Department of Mathematics and Computer Science, Lake Forest, IL, United States

<i>Phylogenetic trees</i>	1
Hector Baños, Nathaniel Bushek, Ruth Davidson, Elizabeth Gross, Pamela E. Harris, Robert Krone, Colby Long, Allen Stewart and Robert Walker	
<i>Software for doing computations in graded Lie algebras</i>	9
Clas Löfwall and Samuel Lundqvist	
<i>The relative canonical resolution: Macaulay2-package, experiments and conjectures</i>	15
Christian Bopp and Michael Hoff	
<i>The FrobeniusThresholds package for Macaulay2</i>	25
Daniel J. Hernández, Karl Schwede, Pedro Teixeira and Emily E. Witt	
<i>Computing theta functions with Julia</i>	41
Daniele Agostini and Lynn Chua	
<i>Decomposable sparse polynomial systems</i>	53
Taylor Brysiewicz, Jose Israel Rodriguez, Frank Sottile and Thomas Yahl	
<i>A package for computations with sparse resultants</i>	61
Giovanni Staglianò	
<i>ExteriorModules: a package for computing monomial modules over an exterior algebra</i>	71
Luca Amata and Marilena Crupi	
<i>The Schur–Veronese package in Macaulay2</i>	83
Juliette Bruce, Daniel Erman, Steve Goldstein and Jay Yang	
<i>admcycles - a Sage package for calculations in the tautological ring of the moduli space of stable curves</i>	89
Vincent Delecroix, Johannes Schmitt and Jason van Zelm	
<i>Coding theory package for Macaulay2</i>	113
Taylor Ball, Eduardo Camps, Henry Chimal-Dzul, Delio Jaramillo-Velez, Hiram López, Nathan Nichols, Matthew Perkins, Ivan Soprunov, German Vera-Martínez and Gwyn Whieldon	
<i>Threaded Gröbner bases: a Macaulay2 package</i>	123
Sonja Petrović and Shahrzad Zelenberg	
<i>Standard pairs of monomial ideals over nonnormal affine semigroups in SageMath</i>	129
Byeongsu Yu	
<i>Computations with rational maps between multi-projective varieties</i>	143
Giovanni Staglianò	