

DRAFT: Comparing the speed of programs for sparse polynomial multiplication

Richard Fateman

July 2, 2002

Abstract

How should one program the multiplication of sparse polynomials? This is a simple question complicated by the details. We start with a single easily stated computation which we believe represents a useful midpoint benchmark for computer algebra systems (CAS).

Introduction: a specific task

Consider the expanded version of the polynomial $p^{20} = (1+x+y+z)^{20}$. This polynomial has 1771 distinct monomial terms. By one ordering, the first term is z^{20} and the 885th term is $1163962800x^9y^5z^4$. Its largest coefficient is 11732745024 which requires more than one 32-bit word to store, since it exceeds 2^{31} . The size of the polynomial p^n grows like this:

n	terms
1	4
2	10
4	35
8	165
16	969
20	1771
40	12341

That is, squaring p^{20} explicitly yields 12,341 terms.

A clever CAS will try to avoid computing such expansion since the factored form is so much more compact. Nevertheless, sometimes explicit multiplications are necessary in a sequence of operations including adds and cancellations.

We will try computing this product in various ways.

Some representations

Several sparse representations can be conveniently expressed in Common Lisp.

The computer algebra system Macsyma represents the relevant polynomial as a list of lists.

We have somewhat simplified the expression format but we use for an example the expansion of $(x + y + z + 1)^2$:

```
(z 2 1 1 (y 1 2 0 (x 1 2 0 2)) 0 (y 2 1 1 (x 1 2 0 2) 0 (x 2 1 1 2 0 1)))
```

which decodes to the *sparse recursive form*

$$z^2 * 1 + z^1 * (y^1 * 2 + y^0 * (x^1 * 2 + x^0 * 2)) + z^0 * (y^2 * 1 + y^1 * \dots)$$

The computer algebra system MockMMA [1] which we wrote a number of years ago uses a similar representation in that it is recursive,

```
#(z #(y #(x 1 2 1) #(x 2 2) 1) #(y #(x 2 2) 2) 1)
```

except that it uses *vectors* and (an inessential detail) the terms are given in reverse order. The Lisp notation `#(a b c)` is the print form for a vector of three items. Contrary to the popular notion of Lisp making everything out of linked lists, this vector is represented as three adjacent cells in memory: there are no pointers. In addition to not “wasting” space on links, MockMMA’s vector representation omits the exponents. This means that non-leading zero coefficients take space. This is a *dense recursive form* where the density is reconsidered at each recursive level.

To see how this trade-off works, for the expression $43x^9 + 5x + 1$ Macsyma would use a list `(x 9 43 1 5 0 1)` and MockMMA would use the vector `#(x 1 5 0 0 0 0 0 0 43)`. For the expression $43x^2 + 25x + 10$ Macsyma would use a list `(x 2 43 1 25 0 10)` and MockMMA would use the vector `#(x 10 25 43)`. Sometimes one form is comparatively more compact and sometimes the other.

Either the list or vector representation has to allow for the possibility of coefficients like 43 or 9 in the examples above to be *arbitrary-precision integers* or *bignums*. This potential complication results in considerable expense compared to programs written with the *a priori* restriction that all coefficients are short integers, integers modulo a small prime, or some fixed-size floats.

A third representation corresponds to a collection of monomials in a hash table. We can convert to this representation by traversing one of the

other representations and successively encoding each of the monomials by collecting exponents and the appropriate coefficient. In the same example shown above, $\{1*z^2*y^0*x^0, 2*z*y^2*x^0, 2*z^1*y^1*x^1, \dots\}$ would be put in the hash table. In this case it makes sense to “compress” the exponents by computing a single key: for the term $c * z^n*y^m*z^k$ the key could be $n + 100m + 10000k$. The coefficient c would be stored at that key location. This works as long as the polynomial degrees needed are not too high (in this case more than 100) and it is especially convenient if this combined key is a small integer (not a bignum). The exponent spacing of 100 can be altered to accomodate higher degrees, but the bound must be anticipated in some way. This idea of packing exponents subject to some pre-declared upper bound was used in ALTRAN, an early computer algebra system. It is not as general as we would like, but is not too burdensome to anticipate an upper bound on degrees, and a careful implementation (as in ALTRAN) can even detect overflows by padding the bit fields.

A final alternative worth considering is a vector: adjacent memory locations (like the MockMMA representation) but with explicit exponent and coefficient pairs. That is, no overhead for pointers. It will turn out that this representation *unordered* has the advantage of compactness over the hashtable, and for multiplication, this may not be bad at all. On the other hand, we may prefer to eventually sort this vector by exponent, since the cost may be modest and we would certainly benefit by this sorting if we were adding polynomials.

1 Algorithms

Fundamentally, any sparse algorithm is going to do $O(st)$ multiplications and $O(st - k)$ additions where the input polynomials have s and t terms and the result has k terms. An attempt to improve on this by using an FFT method, or Karatsuba multiplication will fail because these methods will generally fill in the terms. Our input polynomial $p^{20} = (1 + x + y + z)^{20}$, if represented as “dense to degree 20” in 3 variables would need to store 10,460,353,203 terms instead of the actual 1,771 terms. The way Karatsuba multiplication works best is if you have two dense polynomials of the same degree in some variable, say z , of degree $2n$. Consider splitting p as $A * z^n + B$ where A and B are of now of degree n or less in z and similarly split q into $C * z^n + D$. Normally we would multiply $p * q$ at cost $size(p) * size(q)$, to get $ACz^{2n} + (AD + BC)z^n + BD$. But the Karatsuba algorithm first computes $t_1 = (A + B) * (C + D)$, $t_2 = A * C$, $t_3 = B * D$, and

then $t_4 = t_1 - t_2 - t_3 = AD + BC$, to get $t_2 * z^{2n} + t_4 * z^n + t_3$. The Karatsuba algorithm is used recursively for the three multiplications, and for a dense polynomial would use $O(n^{1.5})$ rather than $O(n^2)$ multiplications. The algorithm's speed depends on the fact that $A + B$, being half the degree of p , is about half the size. Unfortunately if p is sparse, $A + B$ is likely to have about as many terms as p , and so we have reduced the original problem to three multiplications problems, one of which is about as difficult as the original one. In our benchmark problem the size of $A + B$ is 1486, much larger than $1771/2$. For our benchmark, Karatsuba multiplication takes more than 4,495,217 coefficient multiplies compared to the straightforward method which takes about $(1771)^2 = 3,136,441$ multiplies. Using the FFT for multiplying polynomials seems impractical for the sparse case, as well.

Given that algorithm complexity will apparently not provide a lever for faster results, the issue then becomes how to manage storage (especially in the face of cache memory) to optimize this loop: `for each term u in p do for each term v in q do addintosum(u*v)`.

The storage possibilities we investigate include storing the sum as a list as in the input form for Macsyma, using MockMMA's version of a vector, using some union/find structure, a tree, a hashtable, or a more elaborate structure such as cascading hashtables. We evaluate the most promising of these, namely hashing.

Another storage access issue to explore is blocking: the double-loop can be broken into blocks so that only part of each polynomial need be in fast memory at any given time.

Considerations of Coefficient Arithmetic

If we provide for polynomial operations over integers, rational numbers, finite fields or arbitrary rings (including recursively, polynomials in other variables), where do we check on the coefficient domain? We can try to make the code run faster by duplicating the framework of the routines and building separate particular versions of each program for each domain. These might include "arbitrary size integer coefficients", "finite field with small modulus (less than a word size)", "polynomials in exactly three variables with integer coefficients", and perhaps others for "finite field with large modulus (more than a word size)", "single-precision floats", "double-precision floats", "integers of no more than 256 bits" etc. Alternatively, we can do what is done in the Macsyma implementation (one of the programs compared below). Macsyma does not have separate polynomial routines over different

domains. Instead it repeatedly checks (yes, in the innermost loop doing coefficient arithmetic) to see if a global flag, in this case called `modulus` is set to a small prime, in which case modular arithmetic is performed. The default otherwise is to allow inputs that are arbitrary common Lisp numbers, which themselves could be single or double floats, small numbers, complex (etc.) So there is an overhead in this system of type dispatch at the single operation level.

The MockMMA and Hashing routines, also compared below, are programmed without consideration for finite-field arithmetic, eliminating the time taken for a “modulus check” but they do still allow for single/ double/ rational/ complex/ bignum arithmetic.

When Macsyma’s polynomial code was written (1967) memory was expensive and computers were slow. It was logical to write Macsyma with a certain mixture of concern for speed as well as code-size. Today there are several conflicting trends and design decision might be made in different ways.

Now it is plausible to write a separate module for finite-field arithmetic (maybe for several: GF(2), small (odd) primes, and big primes): memory is now so cheap. It is even plausible to do so automatically by reduplicating code by macro-expansion. On the other hand, in today’s object-oriented world it would be clean and logical to use inheritance, overloaded operators, or generic functions for coefficient operations.

Given the total Macsyma context, the time to check the modulus prior to each coefficient operation represents a small overhead, not worth much concern.

In the course of looking at this task we came across a variety of other computer algebra systems, many of them touting the use of an externally developed multiple-precision arithmetic package as a particular advantage.

We tried replacing the coefficient arithmetic with one of the fastest stand-alone arithmetic packages, GMP (Gnu Multiple Precision) version 4.1 for a Pentium-3 architecture. While we are not ready to abandon Lisp’s convenient built-in arithmetic and its overall consistent memory management, it is clear that for in-place repetitive arithmetic, a basic process that is needed for polynomial multiplication, GMP provides a facility that for our use can be faster even for 2-word long arithmetic. For longer integers, GMP’s advantage grows. In fact, the operation of “replace x_{i+j} with $x_{i+j} + a_i * b_j$ ” is not even supported by Lisp if the numerical result does not fit in a single word: the closest operation, which requires allocating new number object and pointing to it, is naturally slower.

Details of the benchmark, including some other programs

The benchmark task is to take the multivariate polynomial $q = p^{20}$ and multiply it by itself explicitly. In order to remove the possibility that a program knows that it is squaring q and thus can do less work, in some cases we may actually compute $q * (q + 1)$ and then forced the expansion.

We timed this computation in MockMMA (recursive vectors), Macsyma (recursive lists), Hashing (using a hash-table for accumulating the intermediate partial sums of the product, using the (built-in) lisp hash-table structure). The Lisp hash-table enlarges automatically if it gets too crowded.

We also tried two additional alternative bignum arithmetic systems. GMP is the Gnu Multiple-Precision package; we tried version 3.1 and 4.1.

Here are the times. (Storage allocation, or garbage collection time (GC) is indicated.)

Hashing/GMP4.1,ACL6.1	4.7 sec (no lisp GC)
MockMMA (in ACL 6.1)	5.8 sec (including 2.1 in GC)
Macsyma (in ACL 6.1)	6.9 sec (including 2.0 in GC)
Hashing/GMP3.1,ACL6.1	7.2 sec (no lisp GC)
Hashing (in ACL 6.1)	7.9 sec (including 2.4 in GC)

These times were measured on the same machine, a 933 Mhz Pentium 3 with 512 megabytes of RAM running Windows 2000; ACL 6.1 is Allegro Common Lisp version 6.1.

The times above are not perfectly reproducible since each depends to some extent on initial memory configurations. Arguably the times for systems written in Lisp could be compared with the GC times subtracted. It is not that the GC time does not count, just that non-GC CPU time tends to be more reproducible and GC time can generally be reduced substantially (or eliminated) by pre-allocating more memory at the start.

After seeing these results we thought it would be interesting to compare these times to other systems. In particular it would seem to be pointless to continue to compare various slow systems if there were much faster ones available (especially if they were similarly general in scope and available free)!

We tried several other Macsyma implementations (different underlying Lisps), and found that a few spent an inordinate amount of time in GC. We also tried some other systems (free and commercial), and are grateful¹

¹Well, actually, it kind of blunts the point of this paper as we had initially conceived it,

to readers of the newsgroup `sci.math.symbolic` for bringing some of these to our attention and in some cases providing copies to test on our machine. Also, we saw that substituting a non-functional state-changing style of arithmetic using the Gnu Multiple-Precision package could make some programs, even in Lisp, run faster². With these additional trials, our table looks like this:

which was to compare lists, vectors, and hashtables, with everything else held constant. It also bloats the paper with footnotes, comments, excuses and program fragments. But it also means we can make some other points beyond our initial goal.

²Also a distraction blunting our main point.

GP/Pari calc 2.0.17	2.3 sec
MockMMA ACL 6.1/GMP4.1	3.3 sec (including 0.4 in GC)
Fermat (see notes)	4.5 sec
Hashing/GMP4.1,ACL6.1	4.7 sec (no Lisp GC)
Reduce 3.7 (in CSL)	5.0 sec (including 1.3 in GC)
MockMMA (in ACL 6.1)	5.8 sec (including 2.1 in GC)
Singular 2.0.3	6.1 sec
Macsyma (in ACL 6.1)	6.9 sec (including 2.0 in GC)
Hashing/GMP3.1,ACL6.1	7.1 sec (including .04 in GC)
Hashing (in ACL 6.1)	7.9 sec (including 2.4 in GC)
Fermat (std)	8.6 sec (fastest of 2 runs.)
Form 3.0 sorted	21.6 sec (see notes)
MuPad 1.4.2 Pro	29.4 sec (poly form)
Maxima5.5 (in GCL)	32.0 sec (? GC not reported, second trial)
Macsyma (commercial)	37.5 sec (including 29.78 in GC)
Maxima5.5 (in GCL)	53.5 sec (? GC not reported, first trial)
Maple7	50.0 sec (? GC, it was slower on 2nd trial)
Mathematica 4.1	82.3 sec (? it was slower on 2nd trial)
Form 3.0	97.9 sec (see notes)
MuPad 1.4.2 Pro	118.6 sec (general form)

We also collected some data for times on other machines:

Reduce 3.7* (PSL)	6.2 sec (including 0.22 GC)
Maxima 5.9.0(CMUCL)*	7.9 sec
Giac 0.2.2* (C++ prg)	13.5 sec
MapleVR3*	16.7 sec
MapleVR4*	17.9 sec
MapleVR5*	18.4 sec
Maxima 5.9.0(GCL-2)*	18.8 sec
Giac 0.2.2*	19.5 sec
Axiom 2.3*	20.1 sec (including 6.56 GC)
Maxima 5.9.0(clisp)*	31.30 sec
Maxima5.6*	51.0 sec
Ginac 1.0.8*	57.3 sec
Maxima 5.9.0(GCL-1)*	112.2 sec
Mupad 2.0*	117.0 sec
Mupad 1.4.2*	119.2 sec
Meditor*	204.5 sec

Notes on this table

*(Giac 0.2.2 compiled with gcc 2.96 on 800Mhz P-III 256K cache. Data provided by Parisse Bernard. Numerous tests were run by Richard Kreckel on a P-III 1GHz 512kB cache Linux machine. Times for his machine seems to be about the same as mine for MuPad 1.4.2 (general), Macsyma 5.6, Mathematica, so probably the numbers for Ginac-1.0.8, Mupad 2.0, MapleVR3,4,5 are about comparable. Meditor took 204.5 sec on a 700Mhz processor. Meditor run by Raphael Jolly.

The collection of Maxima 5.9 timings are from James Amundson, running on a 700Mhz P-III linux machine. The GCL-1 time is, we believe, for the same system that took 53 seconds on our machine. The GCL-2 time is given a larger initial memory allocation. Clisp is a byte-coded Common Lisp that one expects to be slower compared to direct instruction execution because it interposes an extra layer of interpretation. CMUCL is a highly optimizing implementation of Common Lisp based on work at Carnegie Mellon University. Currently CMUCL runs only on variations of the Unix operating system (Linux, Sun Solaris or similar). SBCL is a descendant of CMUCL and may be available on other hosts. If it scales up comparably to the other measurements, it would be fastest of the Lisp systems. Axiom 2.3 time on Duron-600/Linux was is provided by Andrey Grozan. Reduce 3.7/PSL is from Andrey Grozin. PSL, Portable Standard Lisp, is a different version of Lisp, simpler than Common Lisp, and oriented toward support of Reduce. Paul Hsieh provided a copy of his program fscalc (version 0.55) to test: it produces the unsimplified answer in 0.8 seconds on my machine, but the answer is limited to double-precision floats; with this simplification, we did not consider the program comparable.)

Sparseness

As an additional test we tried to probe the sensitivity of the programs to additional sparseness. In this benchmark we used $(1 + x^2 + y^2 + z^2)$ which makes (only) MockMMA substantially less space efficient: it uses twice as much storage for the polynomial “backbones” during multiplication and in the final result. This alteration makes the program run 12 percent slower, uses 16 percent more cons cells and 4 percent more “other” bytes which constitute bignums and vectors. There are 156 megabytes of such other material used during the computation. The Macsyma and Hashing code size and times do not change.

Isolating bignum costs, and effect of memory cache

Of these execution times, how much of it is attributable to bignum arithmetic, and therefore irreducible in the context of improving data structures?

We reproduced the timing with the same programs but with the change that all the coefficients in p^{20} were replaced with the integer 1. Thus the same number of arithmetic operations was still being done, but none of it grew to bignum size in the course of the calculation. We provide this data where we have also added some space data for the ACL runs.

```
MockMMA (in ACL 6.1)  1.12 sec (including 0.12 in GC)
; 112,168 cons cells, 8,356,424 other bytes

Macsyma (in ACL 6.1)  2.39 sec (including 0.21 in GC)
; 3,174,312 cons cells, 96 other bytes

Hashing (in ACL 6.1), 1.7 sec (including 0.02 in GC)
; 12,343 cons cells, 916,680 other bytes

Reduce 3.7 (in CSL)    3.6 (including 0.8 in GC) [Thanks to AC Norman]
```

Hashing is less local

Based on these measurements it would seem that in Macsyma the bignum arithmetic costs about $6.9 - 2.4 = 4.5$ sec. For MockMMA, the time is $5.8 - 1.12 = 4.6$ sec. For Reduce, the time is $5.0 - 3.6 = 1.4$ sec. For Hashing, it seems the bignum arithmetic cost is about $7.9 - 1.7 = 6.2$ sec. This peculiar result says that the same arithmetic takes 1.7 seconds more when called from the Hashing program compared to being called from MockMMA. Since the same instructions are executed on the same data, this requires some further study!

The hashtable + bignum test runs slower because the input and output hashtables combined with the bignum routines, as well as the garbage collections prompted by the bignum arithmetic, in total cause more cache misses³. The data + keys of the two inputs is already several times the size of the L1 data cache. We must also figure that part of the output hashtable

³Pentium 3 specifications: L1 cache, half data and half instruction, is 16Kbytes or 4K fullwords, the L2 cache is 256Kbytes

must also be present in the cache, which grows to contain 12,341 bignums and 12,000 keys). If there is too much data accessed in random fashion (especially so for the output hashtable), successive accesses to are likely to provoke L1 cache misses. But they are mostly within the L2 (256Kbyte) cache. When we move from the fixnum-sized coefficients to bignum coefficients we force rather numerous garbage collections. Switching between the processes of “mutation” and “collection” means the L1 cache gets spoiled (twice) at each GC, But the more expensive problem may be the L2 cache miss rate which is raised by the many small “generation scavenging” GCs during this computation.

Cache performance also explains the relatively high speed of MockMMA, which uses vectors (recursively, of vectors, for each different variable). At the lowest level of the data structure we have a vector of coefficients which are bignums: as we pluck successive coefficients from the input to use as multipliers, the bignums and perhaps the lowest-level coefficient vector tends to persist in a cache.

Further tests on much larger problems than our benchmark here suggests that we are still well below the true “knee” of the cost curve on the L2 cache, which holds 256Kbytes. Even our 12,341-coefficient result fits in our L2 cache, if the L2 cache is not being flushed from activities of the garbage collection (or for that matter, activities of the operating system.) From other tests on this same machine we would expect to suffer an additional slowdown of perhaps a factor of 7 if our data grew routinely to be much larger than the L2 cache [3].

Having seen the importance of locality of data in our programs we should also make some efforts to assure that the design of garbage collection algorithms and policies are cognizant of the issues. Numerous small GCs may be a disadvantage compared to an occasional large one since it will certainly flush the L1 cache and possibly some of the L2 cache. On the other hand the GC algorithm should not be delayed so long that the L2 cache will be filled with its activities.

We have not seen any previous report of Lisp using real-life cache miss data (that is, not just simulations), so there may be some novelty to this part of the report. We have pursued such benchmarking in a different context separately[3], but briefly review the findings here. We used the PAPI system⁴ adapted to Lisp. We instrumented our system to count the total (instruction plus data) cache misses in the Pentium-3 L1 and L2 caches for various tests. All cache miss numbers should be multiplied by 10^7 . For

⁴icl.cs.utk.edu/projects/papi/

convenience we have also reproduced the associated run times.

HT=hash table version using bignums (the standard test)

MM= MockMMA for standard test

HT1 = hash table version with only 1's in the input

	L1 misses	L2 misses	runtime	cycles/L1 miss	cyc/L2miss
HT	4.15	2.79	7.9 sec	177	264
MM	2.62	1.45	5.8 sec	205	373
HT1	1.21	0.86	1.7 sec	131	184

A more detailed analysis would be required to assess the damage from the cache misses. In an attempt to rationalize the time difference between HT and HT1, assume that, *including some measure of related costs*, an L1 miss causes the waste of 12 cycles and an L2 miss similarly uses 120 cycles. Then we have lost $1.8 \cdot 10^9$ cycles or 1.92 seconds, accounting for most of the time difference between HT and MM. There are other memory related activities we have not accounted for, including TLB activity.

Blocking for cache performance

There is yet more that we can do, namely a block-multiply of polynomials, analogous to blocking of matrix operations.

A somewhat abstracted analysis of the cache performance of polynomial multiplication was done (for a class project) by Yozo Hida, in which he showed that performance can be improved by some 40-50% by optimizing block sizes for multiplication. That is, if the two input polynomials are broken up into smaller sets of monomial terms, the “batched” results may be computed faster. A speedup of up to 47% was measured in this paper[4]. We will discuss further results along these lines in a future paper.

Our expectation is that programs based on linked lists would be able to take advantage of blocking for smaller size computations: the cache memory would be about half-filled with pointers. Thus there is a considerable advantage to “denser” representations. The same argument goes for bignums: the ACL bignum package uses only about half the bits of the words (to make multiplication “more portable”). Other systems (in particular GMP) using a 32×32 bit multiply have more compact bignum representations.

Using ACL, we wrote and tested a larger program (with a 4-deep nested loop) that allowed us to experiment with different block sizes. We were

unable to find much improvement over straight multiplication, we think primarily because the result hash-table tended to have such poor memory locality.

Not all Lisp implementations are the same. In cross-Lisp comparisons we suspect that the bignum efficiency substantially affects the results. Certainly the Reduce CSL bignum arithmetic is more efficient than ACL, since its *overall* time is lower than the *arithmetic-only* time for ACL. A separate comparison of CMUCL and ACL (on a Sun Solaris machine) suggests they are approximately equivalent on the “fixnum only” hash version, but CMUCL is about 11 percent faster when bignums are used.

All the formulations do the same coefficient arithmetic, essentially 3,136,441 multiplications and almost that many adds.

Since it appears that the ACL arithmetic is not especially fast, we tried one well-regarded alternative, the Gnu Multiple Precision package, GMP, loaded into ACL. We tested versions 3.1 and 4.1, provided as libraries. An advantage of the 4.1 version, which is used both in our hashing program and in the MockMMA program, is a newly included combined add-multiply call.

The result is that by using GMP, our total hashing time is much lower (4.7 rather than 7.9 seconds) for the size of coefficients in this benchmark. The MockMMA benchmark dropped from 5.8 to 3.3 seconds, but there is a problem in that storage is used extravagantly. (That is, it is not collected properly, and grows rapidly.) A careful reprogramming of this recursive multiplication routine may help, but its functional orientation makes such a change very tricky. Using a variety of tricks to re-allocate GMP numbers by garbage collection or reference counts raises the time to about 9 seconds.

Using GMP more generally

From separate tests in ACL, we know that for raw arithmetic on much bigger numbers GMP is significantly faster: for squaring a number with 2^{17} bits GMP-3.1 is about 12 times faster, and GMP-4.1 is about 45 times faster.

For numbers with 2^6 bits (i.e. two-words long), ACL is twice as fast as GMP4.1 *if new space is allocated for each word*. GMP-4.1 is twice as fast as ACL if the same spot can be reused. At these rather small sizes the time for the arithmetic is dominated by the time to allocate space for a number (a necessary step if a bound on the number’s length is not known until run-time.)

But GMP, used appropriately, has another operation *not supported by Lisp or other typical functional programming systems* which is to modify a number *in place*. That is, GMP allows one to compute `a:=a*b` by mod-

ifying a 's value, thereby (usually) saving memory allocation time. Given this operation, or even better, $a:=a+b*c$, in GMP-4.1, the GMP system becomes a clear winner on this benchmark. The best that Lisp can provide, essentially $\text{newa}:=a+b*c$; $a:=\text{newa}$, creates new objects each time, only to discard them in the vast majority of cases. In this benchmark some 3,136,441 products and 3,124,100 sums are computed, but only 12,341 numbers survive. We have seen that collecting those discarded objects provokes garbage collections and cache misses.

Another cost (which we have not accounted for in the GMP times above) is determining that one can in fact return the GMP numbers to the storage in the GMP system, and actually returning them. The measured cost to return all remaining 12,341 coefficients in this answer is negligible (20 ms), *but this assumes we need do no computations to determine which GMP numbers are "garbage"*. In reality, using GMP as a substitute for all numbers in a Lisp system means we must separately collect garbage in this bignum space. We would miss the smooth transition managed by Lisp between single-precision and multiple-precision integers. Nevertheless, further work with GMP is possible and even has good prospects [2]. We have also programmed a half-page implementation of the Lisp GMP linkage where "finalizations" are attached to GMP object so that their internal storage can be returned when their "shell" is garbage-collected by Lisp. This keeps dead-storage from piling up, but itself adds storage overhead.

Can we change the benchmark to show GMP winning? Certainly: We increase the size of coefficients by starting with $10000000001(1+x+y+z)$. The coefficients in the answer are then about 10^{400} . Our tests show GMP-4.1 takes only 24 seconds (0.06 in GC), GMP-3.1 takes 149 seconds (0.06 in GC) while Lisp takes 167 seconds (21 in GC)⁵.

Running the same modified test takes Pari out of first place: it takes 77 seconds, whereas Singular takes 28 seconds. Each of these is slower than GMP-4.1 Lisp. Of the 24 seconds, at least 11.5 are within GMP, 9.4 seconds within `gmp_add_n` alone.

Using GMP did not speed up our MockMMA version of polynomial multiplication as much as we had hoped because the program produces, recursively, many sub-polynomials which are then added together and should be discarded. This creation of many numbers doesn't make good use of GMP-like facilities. We think a substantial revision of the MockMMA version in which the basic operations are "imperative" rather than functional, may

⁵Lisp takes 145 seconds if the coefficient multiplier is 10000000000. It appears some optimization in Lisp takes advantage of words of zeros!

be faster. That is, the polynomial operations would be based on in-place alterations.

Quibbles about the answers, and a few minor points

Arguably the Hashing version is returning less information in the answer: the monomials are not sorted, and thus such information as “what is the highest degree term in z ?” cannot be answered without some search. This kind of representation is also used in Maple, and we have argued previously that to be fair in comparisons, we should insist that Maple sort its final answers. A final sorting of the hash-table is a small fraction of a second.

Another quibble, perhaps more of a concern, is that the cost for Hashing will increase if the exponents cannot be encoded in a single word, or if we use a short vector of exponents as a key: the exponent arithmetic becomes more costly. The MockMMA version already allows such generality and is therefore another recommendation for it.

An anecdote about FORM

One of our initial tests suggested that FORM could compute the answer very fast: in about 0.9 seconds. But then a minor variation on that computation took 98 seconds. After an exchange of messages with Jos Vermaseren, we looked at the benchmark more carefully. Apparently FORM “expanded brackets” differently from the other programs, and instead of computing $p^{20} \cdot (1 + p^{20})$ it computed the equivalent $p^{20} + p^{40}$, a much quicker computation on any of the systems. It appeared that the 98 second time was more representative of the computation. Starting from this computation and following a suggestion by Jos Vermaseren to insert “sort” directives, we made FORM much faster, reducing the time from 98 to 22 seconds.

Conclusions

A computer algebra system (CAS) needs a multiplication program that is efficient on a typical mix of inputs. Not every CAS expression is a polynomial, but by suitable renaming of indeterminates many computations, we can reduce many bulky and time-consuming calculations to operations on polynomials. At the same time we are addressing multiplication, we presume the same underlying representation will support other operations, most particularly addition, subtraction and division with remainder. Our idea is to

allow arbitrary size integer coefficients, any number of variables, and their exponents could be any positive integers.

However, we have made some judicious simplifications. For example, do we really need to represent a polynomial whose degree exceeds the maximum 32-bit signed integer? While some computer algebra systems allow this, we might not actually find this useful and would resent any substantial inefficiency this might force upon us. Do we need a system with 80,000 different variables?

Given the context of our tests, the major conclusion of this paper is: *memory access speed critically affects computation time for sparse polynomial multiplication*. Hashtable organization which might ordinarily seem fruitful is not as good as expected, since at this size of computation it forces a less effective use of memory cache.

Our original intention at the outset was to see if we could write a compact program to do sparse polynomial multiplication faster than typical systems, but with comparable generality. Our hope was that the hash-table mechanism provided in Common Lisp would enable this. An exceedingly short first program computed the answer, but improving its efficiency made it somewhat larger in size.

It also became apparent that systems with many short garbage collections are at a disadvantage in a cache-memory system since the alternation between computation and collection forced cache misses. Tuning of a garbage collector to take larger chunks of time less frequently should provide better performance. In our tests, using bignums forced frequent GC interruptions. This had more of an effect on the hash program, which already had less locality of data and hence took more time to move into the cache.

Writing in Lisp, a functional language in which $x := x+a*b$ requires creating a new value and “pointing” x to it may be slower in certain applications like this than an alternative version of bignum arithmetic with “in-place” non-functional style operations. (That is, using unit operations that assign $x := x+a*b$ destroying the old value of x).

When we expanded the scope of this paper to see how other systems behaved, we found that Maple 7, Mathematica, Mupad and FORM are not nearly as fast as we had expected, but that GP/Pari, Reduce and Singular made good showings.

Thanks to Richard Kreckel, we learned that earlier versions of Maple were much faster on this benchmark.

Given a constant Lisp environment (Allegro Common Lisp), MockMMA’s mixture of dense arrays and recursion seems to work well for this problem.

Hashing (using “containers” to avoid re-hashing) is better than lists but slower than arrays. If all coefficient arithmetic is done in fixed-size chunks, the compactness of the hash program may be advantageous.

Using different Lisps and different initial allocations of working storage provides substantially different apparent performance. CSL, which is not a full Common Lisp, but a Lisp-based system building tool, provided an especially well-tuned base for Reduce. An important contributor to the speed of this system is its bignum arithmetic which was about 3 times faster. GCL and CLOE Lisp based Macsymas were apparently hobbled by inadequate storage allocations and/or less efficient garbage collection algorithms as initially configured.

Fermat, a relatively smaller interactive system tuned for performance is respectable but was not a top contender on this benchmark straight out of the box. Upon seeing his program’s rank on this test, R. Lewis, its author, looked at his algorithm and realized that the “efficient” algorithm he used, a Karatsuba-style polynomial multiplication, was actually slower than the naive algorithm. Forcing the use of the naive algorithm sped up the system from 8.6 to 4.5 seconds. We put the shorter speed on the chart with the notation “see notes”.

We anticipate that faster bignum arithmetic (for example, from GP/Pari) could be incorporated into any of the slower Lisp systems. In addition to GP/Pari’s arithmetic it may be worthwhile to continue to follow the revisions of GMP, as well as David Bailey’s MPFUN to see which off-the-shelf bignum library might make a net improvement on this benchmark. Such a library system, even if it is fast, still needs to be neatly merged into Lisp including techniques for managing storage. MPFUN packaged as a windows dynamic library can in principle be loaded into our standard Lisp environment, just as we have done with GMP.

A final conclusion: In May, 2002 there are at least 12 computer algebra programs that can be run on a Windows/Intel computer, each competent to do multivariate polynomial multiplication where there are over 10,000 terms in the answer.

Acknowledgments

This research was supported in part by NSF grant CCR-9901933 administered through the Electronics Research Laboratory, University of California, Berkeley. Thanks to the contributors of additional timings who responded to a post on the `sci.math.symbolic` newsgroup, and of course thanks to

the builders of the computer algebra systems tested above. Thanks also to Duane Rettig and John Foderaro for help on simplifying the interface between Allegro Common Lisp and GMP. Thanks to John Chen and Idler Fang for assistance with obtaining the GMP-4.1.dll. Thanks also to Ronen Gradwohl who has assisted in performing some of the experiments, and Derek Wollenstein for additional comments.

Program listings and notes.

Fermat program

```
&(m=0);
&(J=x);
&(J=y);
&(J=z);
w:=1+x+y+z;
r:=w^20;
r1:=r+1;
r*r1;
```

Singular program

```
timer=1;
ring r=0,(x,y,z),dp;
poly p=(1+x+y+z)^20;
poly q=p+1;
poly p*q;
```

GP/PARI Calculator version 2.0.17 (beta) program

```
#
r1=(1+x+y+z)^20;
r2=r1+1;
r1*r2;
```

Macsyma/Maxima program

```
showtime:all$
r:rat(1+x+y+z)^20$
```

```
r1:r+1$
r*r1$
```

Maple program

```
r:=expand((1+x+y+z)^20):
r1:r+1:
st:=time();
expand(r*r1):
time()-st;
```

Hashing program, loaded into macsyma after running above (see Lisp code below)

```
h:hex(r)$
h1:hex(r1)$
hashtimes(h,h1)$
```

Mathematica program

```
r=Expand[(1+x+y+z)^20];
r1=r+1;
Timing[Expand[r*r1];]
```

MockMMA program

```
r=Rat[(1+x+y+z)^20+1]; (* force expansion of power *)
r1=r-r1;
Timing[Rat[r*r1+1];]
```

MuPad 1.4.2 / general form

```
q:=expand(1+x+y+z)^20):
q:=q+1:
time(expand(q*q1));
```

Mupad 1.4.2 using poly form, suggested by Christopher Creutzig,
faster but specialized for polynomials.

```
a:=expand((1+x+y+z)^20):
q:=poly(a):
q1:=poly(a+1):
time(q*q1);
```

```

GINAC program
#include <ginac/ginac.h>
using namespace GiNaC;
int main()
{ symbol x("x"), y("y"), z("z");
  ex r = pow(1+x+y+z, 20);
  ex r1 = r+1;
  ex m = expand(r1*r);
  return 0; }

```

```

Reduce program
on time;
a:=(1+x+y+z)^20$
length(a*(a+1));

```

```

GIAC C++ program
//You can replace giac and gen by any library/class representing bignum
//(gen is a generic type that can represent other data types than just
//large integers)

```

```

#include <giac/giac.h>
using namespace std;
using namespace giac;

```

```

typedef map<int,gen> pol;

```

```

void fastmult(const pol & p1,const pol & p2,pol & product){
  product.clear();
  pol::const_iterator it1=p1.begin(),it1end=p1.end();
  pol::const_iterator it2,it2end=p2.end();
  pol::iterator prod_it;
  int sum1,sum;
  gen coeff1;
  for (;it1!=it1end;++it1){
    sum1=it1->first;
    coeff1=it1->second;
    for (it2=p2.begin();it2!=it2end;++it2){
      sum=sum1+it2->first;
      prod_it=product.find(sum);
      if (prod_it==product.end())

```

```

product[sum]=coeff1*it2->second;
    else
prod_it->second = prod_it->second +coeff1*it2->second;
    }}}

void fastpower(const pol & p,int k,pol & res){
    if (k==0){
        res.clear();
        res[0]=1;
        return;
    }
    res=p;
    pol tmp;
    for (int i=1;i<k;++i){
        tmp=res;
        fastmult(tmp,p,res);
    }
}

ostream & operator << (ostream & of,const pol & p){
    pol::const_iterator it=p.begin(),itend=p.end();
    for (;it!=itend;++it){
        of << it->second << "[" << it->first << "]" ";
    }
    return of;
}

int main(){
    int k;
    cout << "Enter power";
    cin >> k;
    pol xyz;
    xyz[0]=1; // cst coeff
    xyz[1]=1; // x coeff
    xyz[1<<8]=1; // y coeff
    xyz[1<<16]=1; // z coeff
    cout << xyz << endl;
    // make xyz20
    pol xyz20,xyz40;
    fastpower(xyz,k,xyz20);
    cout << xyz20.size() << endl;
}

```

```

    clock_t start=clock();
    fastmult(xyz20,xyz20,xyz40);
    // cout << xyz40 << endl;
    clock_t end=clock();
    cout << end-start << endl;
    cout << "Number of monomials" << xyz40.size() << endl;
}

```

Form programs (3 of them)

*This one takes 0.98 sec

Symbols x,y,z;

Local q2=(1+x+y+z)²⁰*(1+(1+x+y+z)²⁰);

.end

*This one takes 97.9 sec

Symbols x,y,z;

Local q2=(1+x+y+z)²⁰*(2+x+y+z)²⁰;

.end

*This one takes 21.6 seconds

Symbols x,y,z;

Local p = (1+x+y+z)²⁰;

.sort

Local q = p+1;

.sort

Local r = p*q;

.end

Meditor

```

// The polynomials are represented in distributed form, ordered using
// a beta-tree with monomials as keys and coefficients as values. The
// monomials are arrays of integer exponents. The coefficients are java
// BigIntegers.

```

```

import spoke.math.*;

```

```

import spoke.math.analyse.*;

```

```

public class Essai {

```

```

    public static void main(String args[]) {

```

```

Arithmetique a=Polynome.valueOf(Expression.valueOf
      ("(1+x+y+z)20"),new
Variable[] {new Constante("x"),new Constante("y"),new
Constante("z")},Lexicographique.comparerer,0);
System.out.println(a);
long t=System.currentTimeMillis();
Arithmetique b=a.multiply(a);
System.out.println(System.currentTimeMillis()-t);
}}

```

Hashing in Lisp

```

;; Hashing polynomial multiplication
;; Code equivalent to that timed for Allegro Common Lisp is given here.
;; Any Common Lisp, including CSL should be able to compile this
(defun hexinit() ; create a hash expression for (1+x+y+z)
  (let ((h(make-hash-table :test #'eq)))
    (setf (gethash 1 h)(cons 1 1)) ;; 1
    (setf (gethash 100 h) (cons 1 100)) ;; x
    (setf (gethash 10000 h) (cons 1 10000));; y
    (setf (gethash 1000000 h) (cons 1 1000000)) ;;z
    h )) ;; return the hash table

(defun hi20() (let* ((h (hexinit)) (res h))
  (dotimes (i 19 res) (setf res (hashtimes res h))))

(defun hashtimes(r1 r2)
  (let((res (make-hash-table :test #'eq)))
    (declare (optimize (speed 3)(safety 0)(debug 0)))
    (maphash #'(lambda(key val)
      (declare (fixnum key))
      (maphash #'(lambda(key2 val2)
        (declare (fixnum key2))
        (let* ((newk (+ key key2))
          (newcell (gethash newk res))
          (newval (* (car val)(car val2))))
          (declare (fixnum newk)(integer newval))
          (if newcell
            (incf (car newcell) newval)
            (setf (gethash newk res)(cons newval newk))))))
        r2)) r1)

```

```

(maphash #'(lambda(key val)
            (if (eql 0 (car val))(remhash key res)))res) ;optional, remove 0
res))

;; note: a simpler version which stored coefficients only,
;; not (coefficient . exponent) did not run as fast
;; because it computed hash values about 2X as often.
;;
(defun oldhashtimes(r1 r2) ;not used in timings. Simpler though.
  (let((res (make-hash-table :test #'eq)))
    (maphash #'(lambda(key val)
                (maphash #'(lambda(key2 val2)
                            (incf (gethash (+ key key2) res 0)
                                      (* val val2)))
              r2)) r1)
    (maphash #'(lambda(key val)
                (if (eql 0 val)(remhash key res)))res) ;optional, remove 0 terms
    res))

;; set up test
(setf test20 (hi20))
;; this is the program to time
(time (hashtimes test20 test20))

```

Additional notes on Reduce and CSL

Here is a discussion of REDUCE's internal representation (Thanks to Arthur Norman).

An expression is always held as a rational function as a Lisp cons (num . den). Each of these is a poly, where a poly is either

nil	zero
Lisp integer	number
(symbol ...)	used eg for complex, bigfloats, coeff domains
((var . power) . coeff) . reductum	

where var is a prefix form (often a symbol), power is a Lisp integer (so could be a bignum - all exponent arithmetic supports that), coeff is another poly (either a domain elements such as a number, or a poly in less important vars) and the reductum is another poly to add on that is of lower degree

than this term. The prefix form used at top level has a (!*sq . . t) that wraps around the canonical structure described above.

Thus REDUCE and Macsyma's forms are quite similar.

Where REDUCE represents $300x^{10} + 200 * x^5 + 100$ as essentially

```
((x . 10) . 300) ((x . 5) . 200) . 100) ; 6 cons cells
```

Macsyma uses

```
(x 10 300 5 200 0 100) ;7 cons cells
```

But for polys with more terms Macsyma will start to win since it uses just 2 more cons cells per term while REDUCE uses 3 (because REDUCE puts the variable in each time). For low degree polys REDUCE wins by avoiding the “. nil” on the end of Macsyma's list and by avoiding the “0 .” for the constant term, but for a poly with 5 terms or more it looks as if Macsyma would win.

CSL, Codemist System Language, is designed to be a platform for building and delivering products (Reduce, and in a variant form, Axiom) rather than a fully free-standing Lisp. One technique of CSL is compilation of most of the Lisp into a bytecode stream. It then interprets that. This gives a performance hit to some extent, but saves considerable memory which can help speed on machines where memory caches really matter. For better overall speed, systems are benchmarked and the procedures which contribute most of the cost are (statically) compiled from Lisp into C and linked into the binary code. As expected, the programs selected by the profiling tool usually include the polynomial arithmetic.

Thus adding a hashcoding multiply to CSL using the bytecode would not be competitive with the usual compiled polynomial multiply: in fact it is about 2.5 times slower on this benchmark. Compiling the hashcode representation into C and linking it in to the CSL system provides about equal speed.

Here is the R-Lisp version of the hashing code given above.

```
Lisp;  
on backtrace, echo;  
on comp;  
  
symbolic procedure hexinit();  
  begin  
    scalar h;  
    h := mkhash(10, 0, 2.0);
```

```

    puthash(1, h, 1);
    puthash(100, h, 1);
    puthash(10000, h, 1);
    puthash(1000000, h, 1);
    return h
end;

fluid '(ddd);
ddd := 20;

symbolic procedure hi20();
begin
    scalar h, res;
    h := hexinit();
    res := h;
    for i := 2:ddd do res := hashtimes(res, h);
    return res
end;

fluid '(r2 key val res);

symbolic procedure inner(key2, val2);
begin
    scalar n;
    n := key #+ key2;    % uses fixnum-only plus here
    puthash(n, res, gethash(n, res, 0) + val*val2)
end;

symbolic procedure outer(key, val);
    maphash(function inner, r2);

symbolic procedure tidy(key, val);
    if zerop val then remhash(res, key);

symbolic procedure hashtimes(r1, r2);
begin
    scalar res;
    res := mkhash(100, 0, 3.0); % initial size 100, EQ test, grows by
                                % factor of 3 when needed
    maphash(function outer, r1);

```

```

    maphash(function tidy, res);
    return res
end;

fluid '(n);

symbolic procedure incn(key, val);
  n := n #+ 1;

symbolic procedure counthash v;
  begin
    scalar n;
    n := 0;
    maphash(function incn, v);
    return n
  end;

counthash hexinit();
test20 := hi20()$
counthash test20;
on time;
counthash hashtimes(test20, test20);
algebraic;
w := (1+x+y+z)^20$
length (w*w);
end;

```

References

- [1] R. Fateman. “A Lisp-language Mathematica-to-Lisp Translator,” *SIGSAM Bulletin* 24 no. 2 p 19-21 (April, 1990). Also reprinted in *Computer Algebra Nederland Nieuwsbrief* 6, October, 1990.
- [2] R. Fateman. “Importing Pre-packaged Software into Lisp: Experience with Arbitrary-Precision Floating-Point Numbers,” Poster session, ISSAC 2000 International Symposium on Symbolic and Algebraic Computation, St. Andrews, Scotland, UK, August 2000., <http://www.cs.berkeley.edu/~fateman/papers/mpflis.pdf>

- [3] R. Fateman. “Memory Cache and Lisp: Faster list processing via automatically rearranging memory,” Draft, May 2002.
- [4] Yozo Hida. “Data Structures and Cache Behavior of Sparse Polynomial Multiplication,” Class project CS282, UCB May, 2002.