
DEPARTMENT OF MATHEMATICS
TECHNICAL REPORT

ON PARALLELIZING THE CLIFFORD
ALGEBRA PRODUCT FOR CLIFFORD

RAFAL ABLAMOWICZ
BERTFRIED FAUSER

April 2012

No. 2012-2



TENNESSEE TECHNOLOGICAL UNIVERSITY
Cookeville, TN 38505

On Parallelizing the Clifford Algebra Product for CLIFFORD

Rafał Abłamowicz and Bertfried Fauser

Abstract We present, as a proof of concept, a way to parallelize the Clifford product in $\mathcal{C}_{p,q}$ for a diagonalized quadratic form as a new procedure `cmulWpar` in the CLIFFORD package for Maple®. The procedure uses a new `Threads` module available under Maple 15 (and later) and a new CLIFFORD procedure `cmulW` which computes the Clifford product of any two Grassmann monomials in $\mathcal{C}_{p,q}$ with a help of Walsh functions. We benchmark `cmulWpar` and compare it to two other procedures `cmulNUM` and `cmulRS` from CLIFFORD. We comment on how to improve `cmulWpar` by taking advantage of multi-core processors and multithreading available in modern processors.

1 Introduction

In [5] we have described how to use (graded) tensor products and periodicity isomorphisms of real Clifford algebras to accomplish computations in Clifford algebras over vector spaces of dimensions higher than 8. We accomplish these computations with Maple packages CLIFFORD and Bigebra which have been described thoroughly in [1–3, 10]. These packages have proven to be indispensable when deriving mathematical results presented in, for example, [6–9, 11]. Often CLIFFORD and Bigebra have been used to prepare examples in support of a mathematical theory in place of hand computations, e.g., [14], and especially when computing in higher dimensions.

Rafał Abłamowicz
Department of Mathematics, Box 5054, Tennessee Technological University, Cookeville, TN
38505, U.S.A. e-mail: rablamowicz@math.tntech.edu

Bertfried Fauser
The University of Birmingham, School of Computer Science, Edgbaston-Birmingham, B15 2TT,
England. e-mail: b.fausser@cs.bham.ac.uk

Recent applications in engineering use real Clifford (geometric) algebras like $\mathcal{C}_{8,2}$ when modeling geometric transformations in robotics. [13] Thus, there is a need for efficient and fast symbolic computations which not only take advantage of the mathematical theory, for example by using the periodicity theorems, but take also full advantage of recent multi core hardware and software models supporting parallel computing.

In this note, we present an experimental procedure `cmulWpar` from `CLIFFORD` which utilizes the `threading` module available in Maple 15 and later. Maple supports a coarse grained tasks based model for parallel computing, which abstracts the need to actually deal with threads, locks and other low level constructs. The procedure `cmulWpar`, for now, computes, the Clifford product of two arbitrary symbolic elements of type `clipolynom` in the real Clifford algebra $\mathcal{C}_{p,q}$ for a diagonalized quadratic form.¹

In `CLIFFORD` the user can chose between two algorithms² to compute the Clifford product, or supply his/her own routine (not necessarily computing the Clifford product). The two main procedures are `cmulNUM` and `cmulRS` which compute the Clifford product of any two basis monomials of type `clibasmon` in Clifford algebras $\mathcal{C}(B)$ of any bilinear form B . The former is based on Chevalley's recursive definition of the product and performs usually better for bilinear forms with numeric entries, especially if many of them are zero. The latter is based on the Hopf algebraic Rota-Stein cliffordization process and computes faster on fully symbolic bilinear forms. These routines are highly optimized for speed as they use internal features of Maple like hashing of already computed results (using the `remember` option). Although we have succeeded in parallelizing them after making all procedures internal to them thread-safe, in this note we concentrate on computations in real Clifford algebras $\mathcal{C}_{p,q}$ of a non-degenerate quadratic form and the simpler `cmulWpar` procedure.

A third experimental procedure available to `CLIFFORD` is `cmulW`. It belongs (for now) to `Walshpackage` developed by the authors. `cmulW` uses binary coding of basis elements and Walsh functions, see for example [15], to compute the Clifford product of any two basis monomials in $\mathcal{C}_{p,q}$ for a quadratic form of signature (p, q) in an orthogonal basis. We like to recall that `CLICAL`, a stand-alone semi-symbolic "calculator" for $\mathcal{C}_{p,q}$ designed by Lounesto *et al.* [12, 16] already in 1987, was based on binary coding and Walsh functions for internal data handling and storage.

In Section 2, we display and briefly discuss the code of `cmulW` and `cmulWpar` which internally uses `cmulW` for a product of any two basis monomials. We describe a mechanism in `CLIFFORD` which permits the user to select which of the procedures `cmulNUM`, `cmulRS`, `cmulW`, or even a user provided procedure, is used internally by the active, non parallel, top-level procedure `cmul` furnishing the Clifford product in $\mathcal{C}(B)$ (the first two) or $\mathcal{C}_{p,q}$ (the third one). Then, we benchmark the

¹ In worksheets posted at [4], we show how to extend this multi-threading to quantum Clifford algebras $\mathcal{C}(B)$ of any arbitrary bilinear form B .

² The newest version of `CLIFFORD` [3] contains now three algorithms: the original `cmulNUM` and `cmulRS` working in $\mathcal{C}(B)$ for any B , and the newer `cmulWalsh3` which is based on the procedure `cmulW` discussed in this paper.

procedures, namely, the parallel `cmulWpar` against the sequential `cmulW`, `cmul` with `cmulRS`, and `cmul` with `cmulNUM` for some test computations of the most general Clifford polynomials in $\mathcal{C}_{p,q}^{\ell}$ for $p+q \leq 9$. Commented code of all Maple worksheets showing these computations including parallelized `cmulNUM` and `cmulRS` is available at [4].

2 Code of `cmulW` and `cmulWpar`

2.1 The Clifford product based on Walsh functions

First, we present the code of `cmulW` which we use later in the parallel procedure `cmulWpar`. The latter procedure relies on several other procedures, which we do display here for the sake of completeness, and which handle things like producing the Clifford product on basis monomials (`Walsh`) and the data conversion (`convert(<bas>, <data-type1>)`) from CLIFFORD's internal data structures for basis monomials and their representations as binary tuple used by the `oplus` and `Walsh` procedures. As `cmulRS` and `cmulNUM` do not have to perform these conversions, there is a slight loss of speed here due to the data conversion. `twist` provides the proper sign factor due to the grading which is easily computed from the binary (Gray code) representation of the Clifford monomials.

Listing 1 Clifford product on basis monomials e_I, e_J using Walsh functions in $\mathcal{C}_{p,q}^{\ell}$

```

cmulW:=proc(eI::clibasmon,eJ::clibasmon,
             B1::{matrix,list(nonnegint)})
local a,b,ab,monab,Bsig,flag,i,dim_V_loc,ploc,qloc,
        _BSIGNATUREloc;
# -- this procedure depends on external variables
global dim_V,_BSIGNATURE,p,q;
if type(B1,list) then
    ploc,qloc:=op(B1);
    dim_V_loc:=ploc+qloc;
    _BSIGNATUREloc:=[ploc,qloc];
else
    ploc,qloc:=p,q;    ###<<<-- this reads global p and q
    dim_V_loc:=dim_V; ###<<<-- this reads global dim_V
    _BSIGNATUREloc:=[ploc,qloc];
    if not _BSIGNATURE=[ploc,qloc] then _BSIGNATURE:=[p,q] end if;
end if;
# -- data structure conversion: string to binary
a,b:=convert(eI,clibasmon_to_binarytuple,dim_V_loc),
          convert(eJ,clibasmon_to_binarytuple,dim_V_loc);
# -- mod 2 binary addition
ab:=oplus(a,b);
# -- data structure conversion: binary to string
monab:=convert(ab,binarytuple_to_clibasmon);
return
    twist(a,b,_BSIGNATUREloc)*Walsh(a,hinversegGrayCode(b))*monab;

```

```
end proc;
```

2.2 Maple's threading mechanism for coarse grained parallel computing

The following example is taken from Maple's help page `?Threads:-Task:-Start`.³ It explains how to split a computation into pieces when the computation is 'large' enough to profit from a parallel execution, and then execute the parallel tasks and use a continuation function to produce the result. The example computes $\sum_{i=1}^{10^7} i$.

Listing 2 Task threading example

```
continuation := proc ( a, b ) # add two results
    return a + b;
end proc;
task := proc ( i, j )
    # distributes the computation into tasks
    local k;
    if ( j-i < 1000 ) then
        # if the range is small, just compute
        return add( k, k=i..j );
    else
        # split computation into two parts
        k := floor( (j-i)/2 )+i;
        # produce two child tasks, by calling task recursively
        Threads:-Task:-Continue( continuation,
            Task=[ task, i, k ], Task=[ task, k+1, j ] );
    end if;
end proc;
# compute sum 1..10^7 parallel and using add
Threads:-Task:-Start( task, 1, 10^7 ) = add( i, i=1..10^7 );
```

The parallelism is coarse-grained, the user does not have to deal with threads, and, for a large part, with locks. However, the involved routines have to be programmed in a thread-safe fashion.⁴ Since we want to demonstrate how to parallelize the Clifford product `cmulW` in $\mathcal{C}_{p,q}$, in the next section 2.3, we will discuss only `cmulWpar`. The code of `cmulNUMpar` and `cmulRSpar` will be available in the worksheets [4] accompanying this paper.

³ See [17].

⁴ The highly optimized routines `cmulRS` and `cmulNUM` turned up initially *not* to be thread-safe due to relying on an internal function which was not defined in a thread-safe manner. After fixing this problem, both have been now successfully parallelized.

2.3 The parallel procedure *cmulWpar* for the Clifford product

We discuss briefly the code of `cmulWpar`, the parallelized version of the Clifford multiplication based on the Walsh functions core multiplication of the Clifford monomials used in `cmulW` shown in section 2.1. This code needs at least Maple 15, and it has been tested in Maple 15 and Maple 16.

The idea is to implement the Clifford multiplication along the lines of the example given above in the listing 2. As we parallelize a procedure with two arguments, we need to deal with each argument separately which slightly complicates the procedure.

Listing 3 Parallelized version `cmulWpar` of the Clifford product using `cmulW`

```

cmulWpar:=proc (x, y)
  local i, cf, term, co1, lst1, co2, lst2, lst, task, addUp; global p, q;
  # -- process x : turn Clifford polynomial x into a
  # list of type :: list [coeff, monom]
  if type(x, '+' ) then # x is a sum
    lst:=op(x); lst1:=[]:
    for i in lst do # split each term into [coeff, monom]
      if type(i, clibasmon) then
        lst1:=op(lst1), [1, i];
      else
        term, cf:=selectremove(type, i, clibasmon);
        lst1:=op(lst1), [cf, term];
      end if;
    end do;
  elif type(x, '*' ) then # x is a term
    term, cf:=selectremove(type, x, clibasmon);
    lst1:=op([cf, term]);
  else # x is a monom
    lst1:=op([1, x]);
  end if;
  # -- process y : turn Clifford polynomial y into a
  # list of type :: list [coeff, monom]
  if type(y, '+' ) then # y is a sum
    lst:=op(y);
    lst2:=[]:
    for i in lst do # split each term into [coeff, monom]
      if type(i, clibasmon) then
        lst2:=op(lst2), [1, i];
      else
        term, cf:=selectremove(type, i, clibasmon);
        lst2:=op(lst2), [cf, term];
      end if;
    end do;
  elif type(y, '*' ) then # y is a term
    term, cf:=selectremove(type, y, clibasmon);
    lst2:=op([cf, term]);
  else
    lst2:=op([1, y]); # y is a monom
  end if;
  #=====

```

```

# -- set up multitasking
# -- continue function, add up results of task processes
addUp:=proc(a,b) a+b end proc:
# -- task definition
task:=proc(lst1,lst2)
  local i,j, packsize,lst,lstpair;
  packsize:=4;          # size for sequential processing
  # -- if x and y are small, just compute
  if max(nops(lst1),nops(lst2))<=packsize then
    add(add(lst1[i][1]*lst2[j][1]*cmulW(lst1[i][2],
      lst2[j][2],[p,q]),i=1..nops(lst1)),j=1..nops(lst2));
  # -- split the larger list for parallel processing
  elif nops(lst1) < nops(lst2) then # process lst2
    # -- split lst2 (y)
    lstpair:=lst2[1..packsize],lst2[packsize+1..-1];
    # -- produce two new tasks for the split list lst2
    Threads:-Task:-Continue(addUp,
      Task=[ task, lst1, lstpair[1]],
      Task=[ task, lst1, lstpair[2]] );
  else # process lst1
    # -- split lst1 (x)
    lstpair:=lst1[1..packsize],lst1[packsize+1..-1];
    # -- produce two new tasks for the split list lst1
    Threads:-Task:-Continue(addUp,
      Task=[ task, lstpair[1], lst2],
      Task=[ task, lstpair[2], lst2] );
  end if;
end proc:
# -- start computation and collect results
Threads:-Task:-Start( task, lst1,lst2);
end proc:

```

The procedure `cmulWpar` starts off by processing the inputs x, y , which may be Clifford polynomials. First, it splits x into a list `lst1` of lists of type `::List [coeff, monom]`, where `coeff` as a base ring element, and `monom` is a Clifford basis monomial e_I . This splitting is done using the type `clibasmon` from `CLIFFORD`. Similarly, y is split into the list `lst2`. This conversion could be made external by defining a new (external) procedure, however, as `CLIFFORD` deals internally differently with (multi)linearity we keep it inline here, saving also two function calls. The signature (p, q) of the current quadratic form is passed on to `cmulW` through two Maple variables p and q which are declared “global” to `cmulWpar`.

The parallel processing starts with the definition of `addUp`, which adds the results later provided by two tasks. The main routine is `task`, which operates on pairs of type `:: List [coeff, monom]`. We pass here Cartesian products of the two lists in effect. Maple provides in the `combinat` package a way to pass an iterator, which saves memory. However, regarding thread safety we refrained from using this device yet. The parameter `packsize` sets a threshold from which size onwards parallel processing is applied. If both lists are small compared to `packsize`, `task` just computes the result directly, as in the summation example in listing 2. Otherwise, one of the lists is ‘large’ and we split the larger list recursively to produce two new tasks. To do so we use the `Threads:-Task:-Continue(...)` function. This proceeds unless

both lists are ‘small’ and are actually computed in their respective threads. Finally, the `Threads:-Task:-Start(...)` routine initializes the threading mechanism and starts producing the `task` in separate threads and also collects the results.

The number of tasks produced is also the number of threads Maple produces. On a 4-core cpu one would like to have 4 threads only, all taking equal time to compute. For that reason we should compute the parameter `packsize` dynamically. The Maple procedures like `Add`, `Seq`, etc., do this. At the moment we use a static `packsize` and have to compromise between an optimal number of threads and losing parallelism at all. Experiments show that the input and dimension of \mathcal{C} have a large impact on a good choice for `packsize`, a reasonable setting is about 16. However, let us consider two Clifford polynomials x, y with 1,000 terms each. Then the above procedure with `packsize=16` will produce roughly 3,906 tasks and hence as many threads. This clearly contradicts the idea of a coarse grained parallelism and calls for a dynamical setting of `packsize`.

As long as all involved procedures are thread-safe, that is, they can be used without any further (negative) side effects at a threat of miscomputing, parallelizing Maple procedures is formally straightforward. However, to achieve efficiency one needs some understanding of what is going on internally.

3 Benchmarking results for `cmulWpar` versus sequential multiplications `cmulW`, `cmulRS` and `cmulNum`

Benchmarking Maple procedures is not an easy task, and it becomes even more complicated if threads and parallel computing are in use. Firstly, Maple has a garbage collection and it is largely out of control when it does that. Using ‘`cputime`’ as a measure, one needs to take into account that Maple adds up the cpu times on different cores. So, running a procedure on 2 cores for 3 seconds each will be reported as 6 seconds of the cpu time usage. Having the administrative overhead, parallel computations will take *more* cputime than single threaded computations.

The second possibility is to use ‘`realtime`’ which measures the clock time when executing a process. If the two processes above run in parallel, then this should take 3 seconds, but Maple has to share the processor with the operating system and possibly other applications which currently run. Hence, benchmarking has to be done on a clean idle system to get reproducible and comparable results, and this is what we have ensured to be the case. A useful tool for such benchmarking is the `CodeTools:-Usage(...)` procedure of Maple. However, the user should be warned that `profile` and `CodeTools` packages of Maple are *not* yet thread safe and especially `profile` shows at times even negative run times.

We have tested the above given parallelized Clifford product on two machines. The first one is a dual core Windows XP (SP3) machine with Intel (R) Core (TM) 2 Duo CPU 2.19 GHz and 2.9 GB RAM. The second machine is a core i7-2640QM at 2.8-3.5 GHz and 8 GB RAM running ubuntu 11.10 Linux and it has a physical dual core with 4 hyper threading virtual cores seen by linux. The two machines give, up

to a scaling factor of about 2, the same results, so we show only one set of data here. See the appendix for Maple's perception of how many cores are available.

Having a dual or quad core available, we can expect a theoretical speedup of at most a factor of 2 or 4, which in practice suffers from administrative overhead of the threading software. The below given speedups cannot then be attributed to the parallelizing alone, but must also be partially caused by the different ways the procedures compute and by the involved data structures. It looks like as if thinking about the threading model leads one to more efficient code per se. A sort of a measure to check if parallelization does really occur is given by

- a) checking the system load during computation to see if all cores are running under full load,
- b) computing the quotient `cputime/realtime`, which somehow measures the 'effective number of cores in use'.

The ratio `cputime/realtime` varies a lot over the input to `cmulWpar` and it seems to vary from 0.9 to 1.8 on the second machine given above. Maple allows one to set the number of cores (called `cpu's` in Maple) in use. In this way one can, with care (see the appendix) benchmark the parallelized threaded code on a single core versus several cores. As we do not have quad and octo-core machines available, we cannot demonstrate such results which would really show how the threading mechanism scales on the number of cores.

In Table 1 we summarize some of our benchmarking results.⁵ We have computed Clifford products of two most general Clifford polynomials X and Y in the Clifford algebras $\mathcal{C}\ell(p, 0)$ for $p \leq 9$. The table shows CPU times in seconds as they were returned by the Maple procedure `CodeTools:-Usage(...)` where t_1 is the CPU time taken to compute such product by the parallel procedure `cmulWpar`, whereas t_2 , t_3 , and t_4 are the cpu times needed by the ordinary, non-parallel Clifford product procedure `cmul` using internally `cmulW`, `cmulRS`, and `cmulNUM`, respectively. The computations in dimensions 8 and 9 with `cmulW`, `cmulRS`, or `cmulNUM` were not completed due to running out of the RAM memory. For example, in dimension 9 they were stopped after 7.8GB RAM had been consumed and the machine started to swap.

4 Conclusions

We have shown how to use Maple's task model, a coarse-grained parallel computing framework, to parallelize the Clifford product for the `CLIFFORD` package. As long as such computations are side-effects free and thread safe, this is easily achieved by using the `Threads` package from Maple.

We hit on some problems when trying to parallelize the core multiplication procedures `cmulRS` and `cmulNUM`, which are highly optimized mathematically in their

⁵ These times were obtained on Windows XP (SP3) machine with Intel (R) Core (TM) 2 Duo CPU 2.19 GHz and 2.9 GB RAM.

Table 1 Benchmarking CPU times: t_1 of `cmulWpar`; t_2 of `cmulW`, t_3 of `cmulRS` and t_4 of `cmulNUM`

dimV	t_1 [sec]	t_2 [sec]	t_2/t_1	t_3 [sec]	t_3/t_1	t_4 [sec]	t_4/t_1
2	0	0.063	∞	0.047	∞	0.047	∞
3	0.046	0.172	3.7	0.250	5.44	0.203	4.41
4	0.250	0.766	3.1	0.828	3.31	0.984	3.93
5	1.094	3.000	2.74	4.594	4.20	5.532	5.05
6	4.687	14.68	3.13	36.09	7.7	43.61	9.3
7	22.000	136.0	6.18	506.1	22.9	448.6	22.5
8	112.47	NA	NA	NA	NA	NA	NA
9	647.20	NA	NA	NA	NA	NA	NA

algorithms, and, on the software side, by extensive hashing of precomputed results (Maple remember tables). However, after removing some internal procedures these are now thread safe too, and have now parallel versions `cmulNUMpar` and `cmulRSpar`.

In this paper we have discussed the very fast procedure `cmulW` for computing the Clifford product in orthogonal bases and arbitrary signature as it was already thread-safe and allowed immediate parallelization. Amazingly the parallelized version is *much faster* than the theoretical limit allows, so that this speedup is not solely due to our parallelizing the computation. It seems that dealing with the threading package of Maple forced us to produce more efficient coding especially of the multilinear features. We have provided detailed benchmark results showing the speedups and also discussed how to separate the speedup by the given different coding and that coming from actual parallel computing. We find a speedup on large Clifford polynomials due to parallelizing of up to 1.8 on a dual core machine, which is what one can expect and shows that parallel computing is now feasible in a symbolic computer algebra. This is possibly good news for engineers and roboticists who do computations in higher dimensional Clifford algebras like $\mathcal{C}_{8,2}$ for geometric computations.

Thus, we must be cautious when examining the speedup ratios t_4/t_1 and t_3/t_1 shown in Table 1. We repeat to caution the reader that, for example, the speedup factor of around 22 in dimension 7 cannot be attributed exclusively to the parallelization process alone. This speedup is a combination of factors, due to, for example, making the overhead in `cmulWPar` dealing with the bilinearity much smaller and faster than the resources- and time-demanding procedure `clibilinear` from CLIFFORD. The latter includes among other things time-consuming type checking of the input on many levels of recursion, especially when `cmulNUM` is used. At the moment we do actually profit from multi-threading seen by computing the number of effective cores $\text{cputime/realtime} \simeq 0.9\text{--}1.5$. However, we saw that a bit of reorganization of the data structures and the recursive way to do the products (saving memory) gives us an even more substantial speedup. It is questionable if multi-threading is only valuable when one knows that one's code is at its theoretical limit with respect to space and time complexity, and CLIFFORD is not yet at that limit.

We suspect that CLIFFORD could be faster at least by an overall factor of more than 20-30, based on this current experience, by a generic rewrite using better data structures and avoiding all the repetitious parsing and type checking where it can be avoided, and using the recursive way to split (multi)linearity, etc. Optimizing CLIFFORD and its related packages like Bigebra, Cliplus, Octonion, etc. [3] is a priority whose urgency has been emphasized by this exercise in parallelizing the Clifford product.

The results discussed here are accompanied by Maple worksheets posted on [4]. These well-documented worksheets contain further results and alternatives like using the inherently parallel procedures `Add`, `Seq`, `Map` of Maple or producing threads directly. There we further discuss the efficient usage of Maple's `Threads` package. We are working to make all of CLIFFORD thread safe after we have succeeded parallelizing the more complex and complicated `cmulRS` and `cmulNUM` routines. While `cmulRS` is based on a provable optimal algorithm, the above discussion still sheds some light on efficiency of the implementations due to different data structures or recursive computing models (saving memory usage). In that respect, this is a very open area of research.

Acknowledgements Bertfried Fauser wants to thank Darin Ohashi from Maplesoft for his kind help with and email discussions about Maple's threading mechanism. Both authors thank referees for their comments as they have helped us extend this work and improve its clarity.

Appendix

In this appendix we make some remarks about programming practice using the different versions of parallelizing code in Maple, using the threading module of Maple (from version 15 onwards), and the inherently parallel routines such as `Add`, `Mul`, `Map`, and `Seq`.

4.1 Thread safety

The `Threads` package was introduced in Maple 15 and it was improved for Maple 16. Still, large parts of Maple are not yet 'thread safe', that is, the code cannot be run in parallel as it may cause side effects which can interfere with other threads. A common source of such problems are global variables and name space conflicts. For example, Maple's parser will not complain when a running variable in `add(f(i), i=1..10)`; or `seq(i^2, i=1..10)`; is not declared local: it will simply miscompute. The following procedure `dummy` will miscompute in a threading environment unless the local variables `i` and `j` are explicitly declared local to `dummy`.

Listing 4 Local running variables **have** to be declared local

```
dummy:=proc(x::List[expression], N::Integer)
```

```

# local i, j;    # <== needed!
# assignment of j produces a warning if not declared local
j:=x[1];
add(x[i], i=1..N);
end proc:

```

Note, that Maple's parser does **not** complain about an undeclared local variable, so this issue slips through unnoticed, as it will for the (here unused) variable `j` if not declared local.

The next issue is more subtle. If in a procedure `dummy` one has a helper procedure `fun`, declared local to it, and if this function is defined with a remember table, either by option or by assigning certain values to it, then these assignments are seen globally, hence are visible to all threads! In effect, any thread can access values set other threads, and this will ultimately lead to errors.

Listing 5 Avoid procedures local to a procedure

```

# NOT thread safe procedure
dummy:=proc(x::List[expression], N::Integer)
  local fun, i;
  # -- local helper procedure
  fun:=proc(y::any)
    # -- either remember here
    option remember;
    return 3*y^2+4; # just some computation
  end proc:
  # -- or use special cases (implements a remember table)
  fun(x[1]):=x[1];
  fun(x[N]):=x[N];
  add(fun(x[i]), i=1..N);
end proc:

```

CLIFFORD used such a construction to implement permutations in reordering wedge products of Grassmann basis monomials, and this rendered `cmulNUM` and `cmulRS` not thread safe at first. The reordering had to be done sequentially or it needed to be done differently.

A further issue with threading comes from the fact that the programmer must check *every* procedure the given package uses whether it is 'thread safe'. Each Maple procedure has a help page and there it is marked if this particular procedure is thread safe, and for which version of Maple onwards. If no such statement is given, one must assume that the procedure is not thread safe. A suggestion of the referee to use the `combinat:-cartprod` construction to iterate over Cartesian products of sets cannot be realized, as this function is *not* (yet) declared by Maple thread safe.

Unfortunately for benchmarking issues, neither `profile` nor the `CodeTools` package is thread safe yet either. While `profile` seems to be broken as it reports once in a while negative running times, the function `CodeTools:-Usage` seems to be reasonably stable. For that reason we used this utility to do our benchmarks, but ran also checks on the results via external timing.

4.2 Overhead versus gain in parallel code

Given the example from listing 2 one observes that just adding integers in parallel is *slower* than doing so sequentially. A similar result is obtained if one uses the parallel code snippet `Add(i, i=1..10^7);`. This shows clearly that the overhead introduced by Maple to produce threads is *too large* to provide any gain in speed. However, the situation changes when one computes a more complicated sum, e.g., $\sum_{i=1}^{10^7} i^{2/3}$ evaluated as float. The code snippet looks like `Add(evalf(i^(2/3)), i=1..10^7);`. Similar effects are encountered when one uses threads directly. To benefit from parallelizing code in Maple, one has to make sure that the work done in a single task is *as large as possible*, and that reflects the idea of coarse grained parallelism implemented by Maple.

Another suggestion of the referee to use `Add`, `Mul`, `Seq`, or `Map` *inside* a thread is also not advisable. On a processor with n cores (cpus), once n threads already have been produced by other devices, parallelizing will only create superficial threads which cannot be processed in parallel since all cpus are already busy running assigned to them threads.

Finally, there is a difference in the `Threads` packages for Maple 15 and 16 how the number of cpus is set. The procedure `kernelopts(numcpus);` reports the number of ‘cpus’ Maple sees and uses for threading. Maple 15 sets the number of cpus to the number of virtual cores of a physical cpu (for example, 4 for a core i7-2640QM with 2 physical and 4 virtual hyper-threading cores), while Maple 16 uses the number of physical cores, that is 2 here. However, on modern core i7 processors the virtual cores allow a substantial speed up due to interlocking processes when queues and/or pipelines are filled etc. It is therefore advisable to set the number of cpus at the *beginning* of the worksheet to the number of virtual cores by `kernelopts(numcpus=4);` for 4 virtual cores. Note that when the threading mechanism is used, it is initialized using this number, and the number of cpus *cannot* be reset later again (despite what `kernelopts(numcpus);` later reports).

4.3 Other parallelization mechanisms of Maple

As we have already mentioned, the creation of *tasks* is not the only way you can use Maple’s threading package. We have investigated if different approaches give significantly different results.

The first group of seemingly simple to use procedures includes `Add`, `Mul`, `Seq`, and `Map`, which parallelize the corresponding sequential (lower case) procedures. The advantage is that Maple does some load balancing in computing how many threads are created and that it is very simple to use these procedures. As the previous section shows, one is nevertheless left with benchmarking these procedures as a too naive usage may result in slower code. Given the Clifford product we were able to get roughly a similar speedup as with the method described above, which does not

yet use a dynamical setting of `packsize`. But using `Add` to sum up the terms in a Clifford product is very memory intensive, and this favors other solutions.

The second method is to use the task method, using the `Threads:-Task` package providing the procedures `Start`, `Return` to leave a task, and `Continue`. As we have chosen this model above, there is not much to add here.

A third way to create threads is to directly use the `Threads:-Create` command. Maple provides locks, mutexes and a synchronization using `Threads:-Wait` to deal with these threads directly. For example, to compute the sum of integers $\sum_{i=1}^{10^7}$ in 2 threads one can use this code:

```
restart;
# -- define two _functions_ performing the work
p1:=proc(x) local i; add(i,i=1 ..5*10^6) end proc;
p2:=proc(x) local i; add(i,i=5*10^6+1..10^7 ) end proc;
# -- create 2 threads executing the work
id1:=Threads:-Create(p1(),out1);
id2:=Threads:-Create(p2(),out2);
# -- wait for the two threads id1, id2 to be finished
Threads:-Wait(id1,id2);
# produce the result
out1+out2;
```

The above code produces an output `id1=1, id2=2` –showing the id’s of these threads–, and the numerical sum `50000005000000`. It is clear that this gives the most direct access to the threading mechanism, as the programmer can decide explicitly anything about the threads. This model is especially useful when several threads have to share resources, as one can lock variables, etc., using the `ConditionVariable` and `Mutex` packages. We have benchmarked a version of `cmulWpar` using this direct method (using `numcpus=4` threads) and again obtained essentially the same performance results. Finding this, it seems to be advisable to use in Maple the easiest threading model available for the task at hand, as we have demonstrated.

References

1. Ablamowicz, R.: Clifford algebra computations with Maple. In: Baylis, W. E. (ed.) Clifford (Geometric) Algebras with Applications in Physics, Mathematics, and Engineering, pp. 463–502, Birkhäuser, Boston (1996)
2. ———: Computations with Clifford and Grassmann algebras. *Adv. Applied Clifford Algebras* **19**, No. 3–4, 499–545 (2009)
3. Ablamowicz, R., Fauser, B.: CLIFFORD with Bigebra – A Maple Package for Computations with Clifford and Grassmann algebras (2013), <http://math.tntech.edu/rafal/>. Cited May 28, 2013
4. ———: Maple worksheets created with CLIFFORD for verification of the results presented in this paper (2012), <http://math.tntech.edu/rafal/publications.html>. Cited June 10., 2012
5. ———: Using periodicity theorems for computations in higher dimensional Clifford algebras. Submitted (2012)

6. _____: On the transposition anti-involution in real Clifford algebras III: the automorphism group of the transposition scalar product on spinor spaces. *Linear and Multilinear Algebra*. Online version: iFirst DOI : 10.1080/03081087.2011.624093 (2011)
7. _____: On the transposition anti-involution in real Clifford algebras II: stabilizer groups of primitive idempotents. *Linear and Multilinear Algebra*, Vol. **59**, No. 12, 1359–1381 (2011)
8. _____: On the transposition anti-involution in real Clifford algebras I: the transposition map. *Linear and Multilinear Algebra*, Vol. **59**, No. 12, 1331–1358 (2011)
9. _____: Clifford and Grassmann Hopf algebras via the `Bigebra` package for Maple. *Computer Physics Communications* **170**, 115–130 (2005), math-ph/0212032
10. _____: Mathematics of `CLIFFORD` – A Maple package for Clifford and Grassmann algebras. *Adv. in Applied Clifford Algebras* **15**, No. 2, 157–181 (2005), math-ph/0212031
11. _____: Hecke algebra representations in ideals generated by q -Young Clifford idempotents. In Abłamowicz, R., Fauser, B. (eds.), *Clifford Algebras and their Applications in Mathematical Physics*, Vol. 1: Algebra and Physics, Birkhäuser, Boston, pp. 245–268 (2000)
12. Abłamowicz, R., Sobczyk, G.: Software for Clifford (geometric) algebras. Appendix in Abłamowicz, R., Sobczyk, G. (eds.), *Lectures on Clifford Geometric Algebras and Applications*, pp. 189–209, Birkhäuser, Boston (2004)
13. Bayro-Corrochano, E.: Private communication (2012)
14. Hitzer E., Helmstetter, J., Abłamowicz, R.: Square roots of -1 in real Clifford algebras. <http://arxiv.org/abs/1204.4576>. To appear (2012)
15. Lounesto, P.: *Clifford Algebras and Spinors*, 2nd ed. Cambridge University Press, Cambridge (2001)
16. Lounesto, P., Mikkola, R., Vierros, V.: *CLICAL User Manual: Complex Number, Vector Space and Clifford Algebra Calculator for MS-DOS Personal Computers*. Helsinki University of Technology, Institute of Mathematics, Research Reports **A248**, August (1987)
17. Maple 15 and 16 from Maplesoft, Waterloo Maple Inc. Waterloo, Ontario, <http://www.maplesoft.com/>. Cited June 10, 2012
18. Porteous, I.: *Clifford Algebras and the Classical Groups*. Cambridge University Press, Cambridge (1995)