

# A Parallel Task-based Approach to Linear Algebra

Ashkan Tousimojarad  
School of Computing Science  
University of Glasgow

Email: a.tousimojarad.1@research.gla.ac.uk

Wim Vanderbauwhede  
School of Computing Science  
University of Glasgow  
Email: wim@dcs.gla.ac.uk

**Abstract**—Processors with large numbers of cores are becoming commonplace. In order to take advantage of the available resources in these systems, the programming paradigm has to move towards increased parallelism. However, increasing the level of concurrency in the program does not necessarily lead to better performance. Parallel programming models have to provide flexible ways of defining parallel tasks and at the same time, efficiently managing the created tasks. OpenMP is a widely accepted programming model for shared-memory architectures. In this paper we highlight some of the drawbacks in the OpenMP tasking approach, and propose an alternative model based on the Glasgow Parallel Reduction Machine (GPRM) programming framework. As the main focus of this study, we deploy our model to solve a fundamental linear algebra problem, LU factorisation of sparse matrices. We have used the SparseLU benchmark from the BOTS benchmark suite, and compared the results obtained from our model to those of the OpenMP tasking approach. The TILEPro64 system has been used to run the experiments. The results are very promising, not only because of the performance improvement for this particular problem, but also because they verify the task management efficiency, stability, and flexibility of our model, which can be applied to solve problems in future many-core systems.

## I. INTRODUCTION

Task-based parallel programming models are evolving rapidly. With the emerge of many-core processors, they can compete with data-parallel approaches, while offering more flexibility, because of their MIMD nature. A task is any form of computation that can be run in parallel with other tasks, if their data dependencies allow. Most of these programming models provide the programmer with some keywords to express parallelism in an imperative language such as C/C++. Pure functional programming languages on the other hand provide native parallelism, but compared to mainstream languages such as C++ and Java, none of them have found widespread adoption. Even if a many-core programming language would find wide adoption, it would in the short term obviously be impossible to rewrite the vast amount of single-core legacy code libraries, nor would it be productive. Our mission is therefore to propose a programming model that can be integrated into existing codes in imperative languages, while offering native parallelism, similar to functional languages. Before going into the details of our approach, the Glasgow Parallel Reduction Machine (GPRM), we would like to briefly review some of the available models in the market, namely Clojure, Chapel, Intel Cilk Plus, Intel Threading Building Blocks (TBB), and OpenMP.

Parallel programming is not as simple as sequential programming. In addition to *what* to compute, the programmer should specify *how* to coordinate the computation. Clojure [1] –also called a modern Lisp– is a functional programming language that targets the Java Virtual Machine (JVM). Clojure’s syntax is based on S-expressions, i.e. lists where the first element represents the operation and the other elements the operands. GPRM uses a similar approach for the internal representation of its communication code.

OpenCL [2] is an industrial standard for heterogeneous architectures. It basically defines a set of core functionality that is supported by all devices, and allow vendors to expose more programming interfaces as well as hardware features. It is however not as easy to use as the following models for shared memory architectures.

Intel Cilk Plus which is based on the Cilk++ [3] is an extension to C/C++ to provide both task and data parallelism. Is has become popular because of its simplicity. It has `_Cilk_spawn` and `_Cilk_sync` keywords to spawn and synchronise the tasks. `_Cilk_for` loop is a parallel replacement for sequential loops in C/C++. Intel Cilk Plus starts a pool of threads in the beginning of the program which is analogous to the GPRM thread pool.

Intel Threading Building Blocks (TBB) is another well-known approach for expressing task-based parallelism [4]. Intel TBB is an object-oriented C++ runtime library that contains data structures and algorithms to be used in parallel programs. It abstracts the low-level threading details, which is similar to the GPRM design consideration. However, the tasking comes along with an overhead. Conversion of the legacy code to TBB requires restructuring certain parts of the program to fit the TBB templates. One of the advantages of TBB over OpenMP and Cilk Plus is that it does not require specific compiler support.

OpenMP can be called the most widely used programming standard in the shared-memory architectures. Since the release of OpenMP 3.0 [5], irregular parallelism can be expressed by means of the OpenMP *tasks*. Moreover, OpenMP works very well with predictable data parallel situations compared to Cilk Plus and TBB. This makes it a challenging competitor for new programming models such as GPRM. We have compared the performance of GPRM with that of OpenMP in 2 different scenarios: first a micro-benchmark which has structured parallelism, and second, a linear algebra problem which fits very well into less structured task-based parallelism.

It has been studied in [6]–[9] that OpenMP performs poorly for fine-grained tasks. It indicates that the programmer is responsible to figure out how a problem with specific input parameters would fit on a particular platform. As a common solution, programmers use a cutoff value when creating OpenMP tasks to avoid composing fine-grained tasks. Firstly, finding a proper cutoff value is not straightforward, and sometimes needs a comprehensive analysis of the program. It often depends on the application structure and the input data set [10]. Secondly, only in some special cases, such as recursion, the cutoff value can be controlled by the user code. Leaving the decision to the runtime system has been proposed as an alternative. The idea is to aggregate tasks by not creating some of the user specified tasks and instead executing them serially. Adaptive Task Cutoff (ATC) [6] implemented in the Nanos runtime system –a research OpenMP runtime system– is a scheme to modify the cutoff dynamically based on profiling data collected early in the program’s execution. This, however, cannot be done without any overhead at all, plus the scheme needs to be extended to find all interesting situations for cutoff. Instead, in GPRM there is a technique to specify which task would initially run on which thread. This reduces the overhead of tasks considerably. Moreover, there are some worksharing constructs that can be used along with tasks to avoid generating too many fine-grained tasks.

The trade-off between task granularity and the number of tasks in OpenMP is covered in [8]. The authors suggest that tasks should have sufficient granularity and the granularity should be increased as the number of consumer threads increases. This is to ensure that all threads are kept busy doing some useful work. They have also explored that the number of tasks has an effect on the load balance, which means programmers have to trade-off between the number of tasks and granularity in order to get a fair load balance, hence a good performance.

For the purposes of this paper, we will show how OpenMP fails to operate normally for fine-grained tasks, while the proposed model copes with such situations naturally (See Section V). Furthermore, we will introduce a hybrid worksharing-tasking approach to avoid creating too many tasks (See Section VI). In other words, a hybrid methodology for exploiting both task and data parallelism will be employed to solve an LU Factorisation problem on a homogeneous many-core processor .

Lower-Upper factorisation over sparse matrices is a fundamental linear algebra problem. Due to the sparseness of the matrix, conventional worksharing solutions are not enough, since a lot of load imbalance exists. As a well-known testcase, we have used the SparseLU benchmark from the the Barcelona OpenMP Tasks Suite (BOTS) [11]. In this problem, the matrix is organised in blocks that may not be allocated. More information along with the source code is openly available.

## II. GPRM

The Glasgow Parallel Reduction Machine (GPRM) [9] provides a task-based approach to many-core programming.

The programmer structures programs into task code, written as C++ classes, and communication code, written in a restricted subset of C++. A *task* is a list of bytecodes representing an S-expression, e.g.  $(S_1 (S_2 10) 20)$  represents a task  $S_1$  taking two arguments, the first argument is the task  $S_2$  which takes as argument the numeric constant 10, and the second argument is the numeric constant 20. GPRM executes the corresponding list of bytecodes with concurrent evaluation of function arguments. For more details on the compilation of the S-expressions and the bytecode see [9]. In our parlance, a *task node* consists of a *task kernel* and a *task manager*. A *task kernel* is typically a complex, self-contained entity offering a specific functionality to the system, which on its own is not aware of the rest of the system. The *task kernel* has run-to-completion semantics. The corresponding *task manager* provides an interface to the kernel. The computational *task kernels* are written as C++ classes. This means that the end user simply creates classes in the `GPRM::Kernel` namespace.

Conceptually, GPRM consists of a set of *tiles* connected over a network. Each *tile* consists of a *task node* and a FIFO queue for incoming packets. Every *tile* runs in its own thread and blocks on the FIFO. The system is event driven, with two possible types of events: arrival of a packet, or the events generated by the *task kernel*. The latter is either creation of a packet or modification of the local state. The reduction engine, i.e. the *task manager* evaluates the bytecode via parallel dispatch of packets requesting computations to other tiles.

Threads in GPRM are treated as execution resources. Therefore, for each processing core there is a thread with its own *task manager*. At the beginning, a pool of threads is created before the actual program starts. It has been pointed out in the Introduction that GPRM offers an effective way of combining tasks with worksharing constructs to avoid creating fine-grained tasks. For instance, instead of creating tasks in a loop which is common in the OpenMP tasking approach [5], one can create as many tasks as the concurrency level in GPRM, each of which with their own indices. These indices can be then used by a worksharing construct to specify which elements of the loop belongs to which thread. The concurrency level is defined as the number of jobs that can be theoretically run simultaneously in the system.

Normally, when the tasks are fairly equal, the best result can be obtained by choosing the concurrency level as the same as the number of threads, which is itself as the same as the number of cores in GPRM. Although in Section VI tasks are not equal, as a solution one can use the GPRM parallel loops to balance the load amongst threads. This solution, as will be shown, works very well when medium size or large sparse matrices are used.

## III. PARALLEL LOOPS

We have created a number of useful parallel loop constructs for use in GPRM. These worksharing constructs corresponds to the `for` worksharing construct in OpenMP, in the sense that they are used to distribute different parts of a work among

different threads. However, there is a big difference in how they perform the operation. In OpenMP, the user marks a loop as an OpenMP `for` with a desirable scheduling strategy, and the OpenMP runtime decides which threads should run which part of the loop; in GPRM, multiple instances of the same task –normally as many as the concurrency level– are generated, each with a different index (similar to the `global_id` in OpenCL). Each of these tasks calls the parallel loop passing in their own index to specify which parts of the work should be performed by their host thread.

The `par_for` construct is used to parallelise a single loop. It distributes the work in a Round-Robin fashion to the threads. A `par_nested_for` treats a nested loop as a single loop and follows the same pattern to distribute the work. Alternatively, the *contiguous* method gives every thread an  $m/n$  chunk, and the remainder  $m\%n$  is distributed one-by-one to the foremost threads. These methods are shown in Fig 1. The need to parallelise nested loops arises often, e.g. in situations where there are variable size loops such as the SparseLU benchmark in Section VI.



Fig. 1: Partitioning a nested  $m(3 \times 3)$  or a single  $m(9)$  loop amongst  $n(4)$  threads. a) Step size of 1, as in the `par_for` and `par_nested_for`, b) Contiguous

```

1 template<typename Tclass, typename Param1>
2 int par_for(int start, int size, int ind, int CL,
3             Tclass* TC, int (Tclass::*work_function)(int,
4             int, int, Param1), Param1 p1) {
5     /* ind: Index, CL: Concurrency Level */
6     int turn=0;
7     for(int i = start; i < size;) {
8         if(turn % CL == ind) {
9             (TC->*work_function)(i, start, size, p1);
10            i = i + CL;
11        }
12        else {
13            i++;
14            turn++;
15        }
16    }
17    return;
18 }

```

Listing 1: Implementation of the `par_for`

The `par_for` and `par_nested_for` loops in GPRM are implemented using C++ templates and member-function pointers. The implementation of these worksharing constructs are given in Listing 1 and 2. They would be our worksharing constructs by default. The Contiguous parallel loops have similar implementations. We denote Contiguous parallel loops as Contiguous GPRM approaches.

Another useful worksharing construct is a parallel nested loop. Since the GPRM `par_nested_for` is implemented with minimum overhead, it is a significantly useful worksharing construct, as we will see in the next sections.

```

1 template<typename Tclass, typename Param1>
2 int par_nested_for(int start1, int size1, int
3                   start2, int size2, int ind, int CL, Tclass*
4                   TC, int (Tclass::*work_function)(int, int, int,
5                   int, int, int, Param1), Param1 p1) {
6     int turn=0;
7     for(int i = start1; i < size1; i++) {
8         for(int j = start2; j < size2;) {
9             if((turn >= 0) && (turn % CL == ind)) {
10                (TC->*work_function)(i, j, start1, size1, start2,
11                size2, p1);
12                j = j + CL;
13                if(j >= size2) turn = size2 - j + ind;
14            }
15            else {
16                j++;
17                turn++;
18            }
19        }
20    }
21    return;
22 }

```

Listing 2: Implementation of the `par_nested_for`

#### IV. EXPERIMENTAL SETUP

The Tiler TILEPro64 Tile Processor is a 32-bit VLIW multicore with 64 tiles, interconnected via multiple  $8 \times 8$  mesh networks. It provides distributed cache-coherent shared memory by default. It has 16GB of DDR memory, but in order to use the global address space shared among all tiles, addressing is limited to 32-bit, i.e. 4GB. It has per-core L1 caches of 8KB, and L2 caches of 64KB. The union of all L2 caches across the chip comprises the distributed L3 cache. The operating frequency of the cores is 866MHz. Out of 64 tiles, one is used for the PCI communication, and the other 63 tiles have been used for our experiments. For our experiments in this study, we have used the *tile-g++* compiler provided by MDE 3.0 from the Tiler Corporation, and is based on the GCC version 4.4.3. The compiler flags `-O2` and `-std=c++0x` have been specified. It is worth stating that the TILEPro64 runs *Tile Linux* which is based on the standard open-source Linux version 2.6.36.

#### V. MATRIX MULTIPLICATION MICRO-BENCHMARK

In this section, we use a naive matrix multiplication algorithm with a triple nested loop as a micro-benchmark to evaluate the overhead of the OpenMP approach compared to the model we use to solve matrix problems. The code is given at Listing 3.

As our aim is to use this micro-benchmark to identify the most important barriers on the way, we change the interpretation of the problem to performing multiple jobs. Suppose that the product of an  $m \times n$  matrix A and an  $n \times p$  matrix B is the  $m \times p$  matrix C. We want to parallelise the first loop of the triple nested loop, which loops on  $m$ , therefore  $m$  becomes the number of jobs for this problem. The size of each job is identified by the sizes of the next two loops in the triple nested loop, i.e.  $p * n$ . We have chosen  $n = p$  to make the problem more regular. We end up with matrices with the following

specification:  $A : m \times n$ ,  $B : n \times n$ , and  $C : m \times n$ . Due to the poor data locality of this algorithm, one should not expect to see a linear speedup.

---

```

1 for (int i = 0; i < m; i++){
2   for (int j = 0; j < p; j++){
3     for (int k = 0; k < n; k++){
4       C[i*p+j] += A[i*n+k] * B[k*p+j];
5     }
6   }
7 }

```

---

Listing 3: Matrix Multiplication Micro-benchmark

Four approaches are selected for the comparison: I) The OpenMP `for` worksharing construct, II) The OpenMP `for` with *dynamic* schedule and *chunk\_size* of 1, III) The OpenMP Tasks, and IV) The GPRM `par_for` construct.

Figure 2 shows the performance measurement for different job sizes. GPRM outperforms OpenMP in all cases but especially for the small job case (even then, the job size is still not small enough to show the real overhead of having fine-grained tasks in OpenMP). To our best understanding, the performance difference is due to the overhead of thread scheduling, which is more visible in the small job cases with short execution times.

In order to investigate the effect of task granularity on the behaviour of the OpenMP’s performance, we decrease the size of the tasks even more. We also looked at the influence of the cutoff value on the performance. Since the behaviours of the OpenMP worksharing constructs were fairly similar, only the default `omp for` is used for the next experiment.

To improve the behaviour of the tasking approach, we added a cutoff value for the tasks, such that only  $m/\text{cutoff}$  tasks were created. This is similar to sequencing multiple tasks. Fig 3 compares the speedup of a tuned version of the OpenMP task-based model with the other alternatives. We believe that the regularity of GPRM in assigning tasks to its threads and the lower overhead of its tasks make it the winner.

---

```

1 for (int i = 0; i < (m/cutoff); i++){
2   #pragma omp task firstprivate(i)
3   for (int t = 0; t < cutoff; t++) {
4     for (int j = 0; j < p; j++){ // p=n
5       for (int k = 0; k < n; k++){
6         C[(i*cutoff+t)*p+j] += A[(i*cutoff+t)*n+k] *
7           B[k*p+j];
8       }
9     }
10 }

```

---

Listing 4: Matrix Multiplication Micro-benchmark with a cutoff value

Figure 3 shows that the poor behaviour of fine-grained tasks can be remedied to a considerable extent by using a proper cutoff value. The OpenMP approaches gradually becomes better when the size of the job is increased. We have chosen the first two cases to show the effect of using a cutoff in more detail, because these cases show degraded performance

compared to the sequential implementation if no cutoff is used at all.

Figure 4 shows that a good choice of the cutoff value gives the speedup of  $38.6\times$  against the case with no cutoff and  $7.8\times$  against the sequential version, for the job size of  $50 \times 50$  with 63 threads. The speedup for the job size of  $100 \times 100$  is also improved by  $10.8\times$  compared to the case with no cutoff and  $8.2\times$  compared to the sequential runtime.

## VI. SPARSE LU FACTORISATION

The SparseLU benchmark from the BOTS suite, which computes an LU factorisation over sparse matrices, is a proper example of matrix operations with load imbalance. In the OpenMP approach, a task is created for each non-empty block. The main SparseLU code from [5] (omitting the details of the OpenMP task-based programming, such as dealing with shared and private variables) is copied in Fig 5. The number of tasks and the granularity of them depends on the number of non-empty blocks and the size of each block, hence a cutoff value cannot be defined inside the user-written OpenMP code. As has also been discussed in [5], using OpenMP tasks results in better performance compared to using the `for` worksharing construct with dynamic scheduling. Therefore, we use the tasking approach for the comparison.

The code for the SparseLU benchmark in GPRM can be found in Listing 5.

---

```

1 #include <GPRM/API.h>
2 using namespace GPRM::API;
3 GPRM::Kernel::SpLU sp;
4
5 void fwd_bdiv_tasks (int kk, float** A,
6   const int CL) {
7   #pragma gprm unroll
8   for (int n = 1; n < (CL/2); n++) {
9     sp.fwd_t(kk, A, n-1, CL/2); // fwd task
10    sp.bdiv_t(kk, A, n-1, CL/2); // bdiv task
11  }
12 }
13
14 void bmod_tasks (int kk, float** A, const int CL) {
15   #pragma gprm unroll
16   for (int n = 1; n < CL; n++) {
17     sp.bmod_t(kk, A, n-1, CL); // bmod task
18   }
19 }
20
21 void GPRM::Compute_LU () {
22   #pragma gprm seq
23   { /* GPRM evaluates in parallel unless otherwise
24      stated */
25     float** A = init_task();
26     for (int kk=0, kk<NB, i++) { // NB: #Blocks
27       #pragma gprm seq
28       {
29         lu0_task(kk, A);
30         fwd_bdiv_tasks(kk, A, 63)
31         bmod_tasks(kk, A, 63)
32       }
33     }
34   }
35   return; }

```

---

Listing 5: SparseLU code in GPRM, Concurrency Level: 63

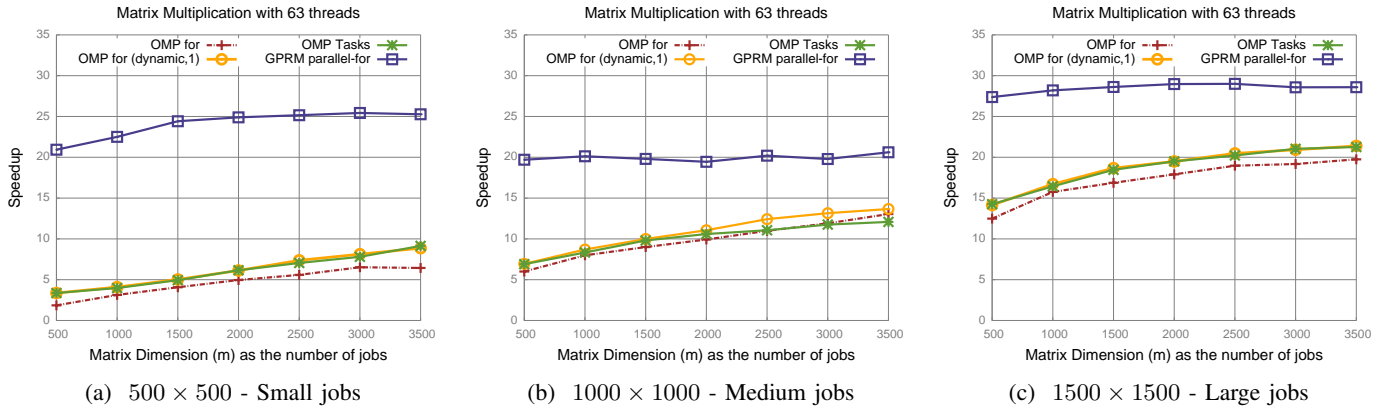


Fig. 2: Matrix Multiplication: GPRM `par_for` vs. different OpenMP approaches




Fig. 3: Speedup measurement for fine-grained jobs, Number of the jobs: 200,000




Fig. 4: Speedup improvement by using a cutoff value for the fine-grained OpenMP tasks. Number of the jobs: 200,000. Size of the jobs: 50x50 - 100x100

The `unroll pragma` results in compile-time evaluation of any control construct in which any of the variable in the argument list occurs. In the example, this means that the `for` loop will be unrolled. By default, expressions in non-kernel GPRM code are evaluated in parallel. The `seq pragma` forces

```
1 int sparseLU() {
2   int ii, jj, kk;
3   #pragma omp parallel
4   #pragma omp single nowait
5   for (kk=0; kk<NB; kk++) {
6     lu0(A[kk][kk]);
7     /* fwd phase */
8     for (jj=kk+1; jj < NB; jj++)
9       if (A[kk][jj] != NULL)
10        /* only create tasks for non-empty blocks */
11        #pragma omp task
12        fwd(A[kk][kk], A[kk][jj]);
13    /* bdiv phase */
14    for (ii=kk+1; ii < NB; ii++)
15      if (A[ii][kk] != NULL)
16        /* only create tasks for non-empty blocks */
17        #pragma omp task
18        bdiv(A[kk][kk], A[ii][kk]);
19    /* wait for previous tasks */
20    #pragma omp taskwait
21    /* bmod phase */
22    for (ii=kk+1; ii < NB; ii++)
23      if (A[ii][kk] != NULL)
24        for (jj=kk+1; jj < NB; jj++)
25          if (A[kk][jj] != NULL)
26            /* only create tasks for non-empty blocks */
27            #pragma omp task
28            {
29              if (A[ii][jj]==NULL)
30                A[ii][jj]=allocate_clean_block();
31              bmod(A[ii][kk], A[kk][jj], A[ii][jj]);
32            }
33    /* wait for all previous tasks */
34    #pragma omp taskwait
35  }
36 }
```

Fig. 5: Main code of the SprsrLU benchmark, without revealing OpenMP programming details [5]

sequential evaluation of the block it precedes.

It is worth mentioning that we have not changed the initialisation phase of generating sparse matrices in the BOTS benchmark suite. The matrices become sparser as the number of blocks increases. For example, in the case of 50x50 blocks, the matrices are 85% sparse, while for the cases with 100x100 blocks, the matrices become 89% sparse.

To obtain a fair distribution of the matrix elements amongst different threads –bearing in mind the sparseness of the matrix– we have used a `par_nested_for`, because the numbers of iterations are not fixed in this problem. The

loops become smaller as  $kk$  grows, which means that using a `par_for` would, after a few iterations when `outer_iters > concurrency_level`, lead to starvation of some of the threads. By using a `par_nested_for` the threads can get some work as long as the `outer_iters*inner_iters > concurrency_level`. Therefore, in order to implement the `fwd`, `bdiv`, and `bmod` tasks, one can use the GPRM APIs for the parallel loops, as shown in Listing 6.

```

1 #include <GPRM/API.h>
2 using namespace GPRM::API;
3
4 int SpLU::fwd_t(int kk, float** A, int ind, int CL)
5 {return par_for(kk+1,NB, ind, CL, this,
6   &SpLU::fwd_work, A);}
7
8 int SpLU::bdiv_t(int kk, float** A, int ind, int CL)
9 {return par_for(kk+1,NB, ind, CL, this,
10  &SpLU::bdiv_work, A);}
11
12 int SpLU::bmod_t(int kk, float** A, int ind, int CL)
13 {return par_nested_for(kk+1,NB, kk+1,NB, ind, CL,
14  this, &SpLU::bmod_work, A);}
15
16 /* The fwd function here is the same as fwd in
17    the cited paper */
18 int SpLU::fwd_work (int jj, int kk, int NB,
19   float** A) {
20   if (A[(kk-1)*NB+jj] != NULL)
21     fwd(A[(kk-1)*NB+kk-1], A[(kk-1)*NB+jj]);
22   return 0;
23 }

```

Listing 6: Implementation of the member functions of the SparseLU class. Work functions can also be defined similar to the phases in [5]




Fig. 6: Execution time of sparse matrices of size 4000 with variable block sizes

Figure 6 shows a sparse matrix of  $4000 \times 4000$  divided into blocks of varying size. It is again clear that with larger numbers of blocks in each dimension which results in smaller block sizes, the OpenMP's performance drops drastically. GPRM can deal with tiny  $8 \times 8$  blocks  $6.2 \times$  better than the best result obtained by OpenMP.

Table I reveals that the best results for the OpenMP approach is not obtained with the default setting, which is as many threads as the number of cores. Besides the large difference in execution times when the block size becomes

less than  $20 \times 20$ , there is a significant performance degradation if the number of threads is set to the default value of 63. For example, the execution time becomes  $12.25 \times$  worse than the best time for the last case. However, it is clear that GPRM reaches its best execution time without the need to tune the number of threads—here, the number of threads and concurrency level are the same for GPRM—. This is again due to the fact that instead of creating very small tasks, GPRM offers an efficient way of distributing the work amongst threads.

The speedup diagrams for the SparseLU benchmark have been shown in Fig 7. We increased the concurrency level up to 128 to show how regular our approach is, in the sense that it gets its best performance with the factors of the number of cores. This is not surprising, since the problem has been partitioned amongst threads regularly, and therefore they can exploit the underlying architecture more efficiently. Since the OpenMP tasking model is different from ours, we simply increased the number of threads in that case.

TABLE I: Number of threads for the best results

Number of Blocks	50	100	200	400	500
OpenMP tasks	64	63	32	16	8
GPRM	63	63	63	63	63

## VII. DISCUSSIONS

### A. OpenMP Performance Bottlenecks


We have identified a number of performance bottlenecks when programming with OpenMP. The first is the thread migration overhead. This overhead can often be removed by statically mapping (pinning) the OpenMP threads to the execution cores. Using static thread mapping (pinning) in a platform with per-core caches could be very useful, particularly for load balanced data parallel problems, in which the portion of the work to be done by each thread is fairly equal, and therefore CPU time and local caches can be effectively utilised. Our study [12] shows that for such a platform, static thread mapping is often a good practice in single-program environments, but for multiprogramming environments, in which different programs compete for the resources, it is not always efficient. We refer the readers to [12]–[14] for detailed discussions on thread mapping (pinning) for OpenMP programs.

Another barrier is to find a proper cutoff point to avoid creating overly fine-grained tasks. Programmers have to be very careful about the granularity of the tasks, otherwise the results maybe totally unexpected. It has also been observed that there is no guarantee that running an OpenMP program with the maximum number of threads—equal to the number of cores—will result in the best performance.


### B. Comparison of OpenMP and GPRM Approaches

For the SparseLU problem, although creating OpenMP tasks for non-empty blocks is a smart solution, it is not working very well for all matrices. The first reason is that a single thread





(a)  $50 \times 50$



(b)  $100 \times 25$

Fig. 7: SparseLU Factorisation: GPRM approaches vs. OpenMP Tasks, on the TILEPro64 with 63 available cores

explores the whole matrix and creates relatively small tasks for non-empty blocks, while in the proposed solution implemented in GPRM, multiple threads look into their portions of the work in parallel. The difference in performance can be noticeable especially in the *bmod* phase with a nested loop. As also reported in [15], combining the OpenMP `for` worksharing construct with tasks, as implemented in `sparselu_for` in the BOTS benchmark suite, is not a viable approach with OpenMP 3.0.

Secondly, in GPRM, every thread has specific work based on the program's *task description* file. If needed, runtime decisions will be applied to improve the performance, while OpenMP creates the tasks dynamically and all decisions are taken dynamically at runtime, which makes its task management less efficient when the numbers of tasks and/or threads become larger. Moreover, the overhead of task management becomes significant when the tasks become more fine-grained, as made clear by the Matrix Multiplication micro-benchmark.

In GPRM, using a `par_nested_for` construct in order to partition the matrix in a non-contiguous manner results in a good load balance. Furthermore, our hybrid worksharing-tasking technique is pretty much straightforward. There is no pressure on the programmer to worry about private and shared variables, in contrast to many other parallel programming models, including OpenMP.

The proposed programming toolkit, GPRM is stable, scalable, low-overhead, and flexible compared to OpenMP. Stable, as there is no need to change the default configurations, such as the number of threads in order to get the optimal performance. It scales as expected, which is showing continuous speedup as the concurrency level increases up to the number of cores. The overhead of task management is negligible, even for fine-grained tasks. It is flexible, because the number of tasks can be easily controlled by the programmer. It is also straightforward to specify which task to be run on which thread initially. If required, the runtime system can change the host thread dynamically.

## VIII. CONCLUSION

Many-core systems have emerged to solve the existing problems faster. Giving the same amount of work to more execution resources simply means having more fine-grained tasks. In this paper, we have proposed a regular task composition approach which can address both regular and irregular parallel problems. We have shown that it fits naturally to the systems with per-core caches, which consequently makes it very promising for future many-cores.

In this paper, our new model is used to solve a fundamental linear algebra problem, namely lower-upper factorisation over sparse matrices. We used a matrix multiplication micro-benchmark to highlight the differences between the proposed model and OpenMP. For the small jobs, GPRM outperformed OpenMP by  $2.8\times$  to  $11\times$ . For the medium-sized jobs, the speedup improvement ranged from  $1.5\times$  to  $3.3\times$ . As the jobs get larger the difference became less significant. For the large jobs, the speedup ranged from  $1.3\times$  to  $2.2\times$ .

As a real-world example we used the SparseLU testcase from the BOTS benchmark suite. We used a sparse matrix of  $4000 \times 4000$  divided into blocks of varying size. We demonstrated that for the larger numbers of blocks, i.e. smaller block sizes, the difference is considerable. The main advantage of GPRM is that it does not need to be tuned in terms of the number of threads. By contrast, tuning is crucial for OpenMP, otherwise for fine-grained tasks a huge drop in performance is inevitable.

We also investigated the impact of the concurrency level on speedup. Again, GPRM scaled  $2\times$  better than the best result obtained using OpenMP for both the  $50 \times 50$  and  $100 \times 100$  cases. For *concurrency level*=63, which is the default setting, the speedup improvements were respectively  $2.1\times$  and  $4.9\times$ .

As GPRM can offer a hybrid task-data parallelism, utilising it on hybrid CPU-GPU systems remains an open research area.

## ACKNOWLEDGMENT

The first author would like to thank the Scottish Informatics and Computer Science Alliance (SICSA) for supporting his research studies.

## REFERENCES

- [1] R. Hickey, “The clojure programming language,” in *Proceedings of the 2008 symposium on Dynamic languages*. ACM, 2008, p. 1.
- [2] J. E. Stone, D. Gohara, and G. Shi, “Opencl: A parallel programming standard for heterogeneous computing systems,” *Computing in science & engineering*, vol. 12, no. 3, p. 66, 2010.
- [3] C. E. Leiserson, “The cilk++ concurrency platform,” *The Journal of Supercomputing*, vol. 51, no. 3, pp. 244–257, 2010.
- [4] J. Reinders, *Intel threading building blocks: outfitting C++ for multi-core processor parallelism*. O’Reilly Media, Inc., 2007.
- [5] E. Ayguadé, N. Coptý, A. Duran, J. Hoeflinger, Y. Lin, F. Massaioli, X. Teruel, P. Unnikrishnan, and G. Zhang, “The design of openmp tasks,” *Parallel and Distributed Systems, IEEE Transactions on*, vol. 20, no. 3, pp. 404–418, 2009.
- [6] A. Duran, J. Corbalán, and E. Ayguadé, “An adaptive cut-off for task parallelism,” in *High Performance Computing, Networking, Storage and Analysis, 2008. SC 2008. International Conference for*. IEEE, 2008, pp. 1–11.
- [7] A. Podobas and M. Brorsson, “A comparison of some recent task-based parallel programming models,” in *Proceedings of the 3rd Workshop on Programmability Issues for Multi-Core Computers, (MULTIPROG’2010), Jan 2010, Pisa*, 2010.
- [8] X. Teruel, C. Barton, A. Duran, X. Martorell, E. Ayguadé, P. Unnikrishnan, G. Zhang, and R. Silvera, “Openmp tasking analysis for programmers,” in *Proceedings of the 2009 Conference of the Center for Advanced Studies on Collaborative Research*. IBM Corp., 2009, pp. 32–42.
- [9] A. Tousimojarad and W. Vanderbauwhede, “The glasgow parallel reduction machine: Programming shared-memory many-core systems using parallel task composition,” *EPTCS*, vol. 137, pp. 79–94.
- [10] A. Duran, J. Corbalán, and E. Ayguadé, “Evaluation of openmp task scheduling strategies,” in *OpenMP in a new era of parallelism*. Springer, 2008, pp. 100–110.
- [11] A. Duran, X. Teruel, R. Ferrer, X. Martorell, and E. Ayguade, “Barcelona openmp tasks suite: A set of benchmarks targeting the exploitation of task parallelism in openmp,” in *Parallel Processing, 2009. ICPP’09. International Conference on*. IEEE, 2009, pp. 124–131.
- [12] A. Tousimojarad and W. Vanderbauwhede, “An efficient thread mapping strategy for multiprogramming on manycore processors,” in *International Conference on Parallel Computing (ParCo 2013)*. IOS Press, 2013, pp. 63–71.
- [13] A. Mazouz, S. Touati, and D. Barthou, “Performance evaluation and analysis of thread pinning strategies on multi-core platforms: Case study of spec omp applications on intel architectures,” in *High Performance Computing and Simulation (HPCS), 2011 International Conference on*. IEEE, 2011, pp. 273–279.
- [14] A. Tousimojarad and W. Vanderbauwhede, “Cache-aware parallel programming for manycore processors,” *ACM SIGARCH Computer Architecture News*, vol. 41, 2013.
- [15] S. L. Olivier, A. K. Porterfield, K. B. Wheeler, and J. F. Prins, “Scheduling task parallelism on multi-socket multicore systems,” in *Proceedings of the 1st International Workshop on Runtime and Operating Systems for Supercomputers*. ACM, 2011, pp. 49–56.