

Bridging the gap between OpenMP and task-based runtime systems for the fast multipole method

Emmanuel Agullo, Olivier Aumage, Berenger Bramas, Olivier Coulaud and Samuel Pitoiset

Abstract—With the advent of complex modern architectures, the low-level paradigms long considered sufficient to build High Performance Computing (HPC) numerical codes have met their limits. Achieving efficiency, ensuring portability, while preserving programming tractability on such hardware prompted the HPC community to design new, higher level paradigms while relying on runtime systems to maintain performance. However, the common weakness of these projects is to deeply tie applications to specific expert-only runtime system APIs. The OpenMP specification, which aims at providing common parallel programming means for shared-memory platforms, appears as a good candidate to address this issue thanks to the latest task-based constructs introduced in its revision 4.0. The goal of this paper is to assess the effectiveness and limits of this support for designing a high-performance numerical library, ScalFMM, implementing the fast multipole method (FMM) that we have deeply re-designed with respect to the most advanced features provided by OpenMP 4. We show that OpenMP 4 allows for significant performance improvements over previous OpenMP revisions on recent multicore processors and that extensions to the 4.0 standard allow for strongly improving the performance, bridging the gap with the very high performance that was so far reserved to expert-only runtime system APIs.

Index Terms—high performance computing, fast multipole method, runtime system, OpenMP, compiler, parallel programming model, priority, commutativity, multicore architecture



1 INTRODUCTION

The advent of modern computing architectures with large number of cores puts a strong pressure on parallel programming paradigms. The low-abstraction, thread-based paradigms reach their limits, due to the difficulty to handle the resulting management and synchronization complexity for a programmer. As a consequence, the High Performance Computing (HPC) community is investigating the design of new, higher level programming paradigms. Among such paradigms, task-based parallelism models have been proposed and implemented in several robust runtime systems. The successful ports of popular numerical libraries on top of these runtimes have demonstrated their effectiveness and benefit in terms of performance portability on complex platforms. Beyond numerical libraries, task-based parallelism is suitable for a wide range of parallel programming patterns [1], the main limit being the weight of task management compared to the parallelism grain. Even then, some task-based models such as Cilk [2] are specifically designed for low overhead, to cope with fine grain parallelism.

However, each task-based runtime system comes with its own, dedicated application programming interface, which may differ significantly from the API of other task-based runtimes. The result is a babel tower-like scheme where choosing a given runtime to port an application deeply ties the application to it. A major objective of the OPENMP specification is to define a common, abstract programming method for shared-memory parallel platforms. OPENMP therefore appears as a prominent candidate to address the

issue of the dedicated task-based runtime systems API profusion, since the introduction of tasks (OPENMP rev 3.x) and more recently dependent tasks (OPENMP rev 4.0) in the specification.

Pondering the use of an abstract layer such as OPENMP as the programming model for a parallel HPC application, involves questioning its benefits and performance trade-offs beyond the mere gain in portability. The goal of this paper is to explore these benefits and trade-offs for a Fast Multipole Methods (FMM) library, SCALFMM, aiming at computing pair-wise particles interactions and whose an overview is presented in Section 3. The main interest of this application is to generate very different workloads and to expose widely differing behaviours depending on the input datasets. The study is conducted both with the GNU GCC compiler, which targets its own LibGOMP runtime system, and with our KLANG compiler, for which we selected the STARPU runtime target. The GCC/LIBGOMP couple is one of the reference implementation of the OPENMP specification, and offers an efficient, lightweight task scheduler. The KLANG/STARPU couple also provides the compliance with the OPENMP `task` directive, while offering the opportunity to experiment with accessing STARPU additional features such as advanced scheduling or data management policies.

We show that OPENMP 4 allows for significant performance improvements over previous OPENMP revisions on recent multicore processors. We also identify situations where this abstraction may prevent leveraging useful runtime system features. Using pinpointed *addenda* to the OPENMP `task` directive, we show that opportunities for extensions exist to give additional driving hints to the underlying runtime system, and take advantage of its ad-

Inria Bordeaux Sud Ouest - HiePACS and Storm Inria project teams.
Max-Planck Computing and Data Facility.

vanced capabilities when they are relevant for a given case, without resorting to a specific, native runtime port of the application.

The contributions of this study are the following:

- delivering a new design of the SCALFMM library with enhanced performance on multicore architectures, compatible with the OPENMP 4 standard;
- evaluating the behaviour of a reference lightweight (LibGOMP) and a fully-featured (STARPU) back-end runtime system to OPENMP compilers;
- proposing optional extensions to the OPENMP 4 standard and assessing how they can speed-up a target state-of-the-art scientific library.

All in all, we show that we can bridge the productivity versus performance gap between OpenMP and native runtimes, ensuring the efficiency of a code tailored for a native runtime system with the compactness and elegance of OpenMP.

The rest of the paper is organized as follows. Section 2 presents the abstraction effort put in designing the OPENMP specification, and the concurrent enterprise of the community to design powerful runtime systems to efficiently execute parallel applications on high performance platforms; it also exposes related work investigating their combination. We propose a brief overview of the FMM in Section 3 to make the paper self-contained, together with guidelines for the reader in a hurry to proceed to the FMM parallelization schemes proposed in Section 4. Thanks to new data structures, we propose extremely compact schemes, based on either OPENMP directives or STARPU routines. We implemented them in the SCALFMM library for the purpose of the present study. Section 5 presents a performance analysis, before Section 6 concludes this paper and presents future research directions.

2 BACKGROUND

2.1 OPENMP

The OPENMP ARB (Architecture Review Board) published the first revisions of the OPENMP specification [3] by the end of the nineties and the beginning of years 2000, with the aim to provide an abstract, portable, programming model and make parallel programming on shared-memory machines a straightforward and user-friendly process. It provided basic constructs to define parallel regions in a fork-join fashion, and build parallel loops as well as parallel sections on top of such regions. At this time, with the exception of a few programming environments such as Cilk [2], programming shared-memory machines necessitated tricky, error-prone manipulations of thread objects provided by the various POSIX threads compliant libraries or custom multi-threading library flavors. The specification of OPENMP has since been revised and enriched several times, especially introducing constructs for independent-tasks programming model (rev. 3.0/3.1 [4]), and dependent tasks (rev. 4.0 [3]), besides many other features.

2.2 Task-based runtime systems

The job of managing the execution and of mapping the parallelism of an application onto computing units has been

the subject of numerous works. It led to the proposal of many runtime systems to implement scheduling algorithms based on theoretical scheduling researches. Among runtime systems, and especially since the emergence of multicore processors, task-based execution models have become popular parallel application foundations. As the number of cores increases in processors, and may greatly vary from platform to platform, the ability of task execution models to flexibly map computations on available resources is a key reason of their success [2], [5], [6], [7], [8], [9]. However, each task-based runtime system comes with its dedicated programming interface. Some environments extend or revisit task-based models with new, richer parallel object constructs and capabilities [10], [11]. However, these extended models are available only within their environments' own ecosystem. Thus, an application ported on a given runtime will necessitate additional programming and refactoring code to run on another runtime, resulting in reduced portability and preventing comparisons between different runtime systems.

2.3 OPENMP/ runtime systems relationships

All OPENMP-compliant compilers underneath rely on a runtime system to support the parallel execution of the compiled code. Yet, reference implementations tend to prefer lightweight engines in order to limit the overhead that advanced policies could induce. For instance, GCC relies on the LibGOMP lightweight runtime system, and the GCC/-LIBGOMP pair got its success in being one of the reference implementation of the OPENMP specification, because this design allows users to achieve competitive performance, for embarrassingly or moderately complex parallel cases, with an extremely high programming productivity.

Those users, instead, who need to achieve even higher performance on test cases exhibiting complex parallelism structure, and endeavour to exploit the most of modern platforms' power, the option to resort to fully-featured task-based runtime systems — designed for that very purpose — is impeded by productivity issues. The learning curve may be very steep and comes with the additional drawback of tying the application to a specific tool. As a result, several proposals have been made to help combine the abstraction of the OPENMP programming model with the benefits of fully-featured task-based runtime systems. These frameworks have in common to rely on a source-to-source OPENMP compiler responsible for translating abstract OPENMP constructs into calls to a runtime system dedicated API. The ROSE compiler [12] developed at LLNL is designed to support multiple back-end runtime libraries. The OPENUH compiler [13], developed at the University of Houston is based on the OPEN64 compiler framework and provides both source-to-source compilation as well as binary generation. The MERCURIUM [7], [14] source-to-source compiler developed at Barcelona Supercomputing Center targets the NANOS++ runtime system, and supports OPENMP compliance as well as OMPs extensions to the OPENMP specification. In particular, MERCURIUM has been used as a prototype for designing the OPENMP tasking model during the latest specification revision processes. Commercial compilers also come with their own runtime system, such as the IBM XL compiler [15], or the INTEL compiler (runtime ABI shared with LLVM [16]).

These proposals are made possible by the abstract, implementation independent architecture of OpenMP. The language defines concepts such as *device*, *thread*, *thread team*, *implicit task* (a task implicitly arising from encountering an OPENMP `parallel` construct) and *explicit task* (a task arising from encountering an OPENMP `task` construct), to name the most useful here. Conceptually, a host device may launch teams of threads, where each thread can run implicit and explicit tasks. Tasks may create new nested tasks, new nested threads teams and so on (see OPENMP's specification [17]). OPENMP also defines rules to create corresponding objects, as well as rules enabling to collapse such objects. This enables different possible implementation strategies, accommodating different hardware capabilities (as an extreme case, these objects can be entirely collapsed to a sequential execution). In particular, implicit tasks are often collapsed with their corresponding thread, in which case no dedicated structure are allocated for them by the runtime system following this implementation strategy. This is for instance the case for GCC/LIBGOMP and for NANOS++.

In this paper, we rely on the KLANG C/C++ source-to-source OPENMP compiler together with the STARPU runtime system [5]. The KLANG compiler is based on the LLVM framework and on INTEL's CLANG-OMP front-end. It translates OPENMP directives into calls to task-based runtime system APIs such as the STARPU runtime. KLANG supports legacy fork-join OPENMP constructs such as `parallel` regions, `parallel for` loops and sections. It also supports independent tasks as defined by the OPENMP specification revision 3.1 as well as dependent tasks introduced with OPENMP 4.0. From an implementation point of view, all *explicit task regions* (arising from both dependent and independent tasks in OPENMP terminology) are directly mapped on STARPU tasks. Legacy fork-join constructs are implemented by mapping their *implicit task regions* on STARPU tasks. All STARPU task scheduling algorithms (priority, work-stealing, ...) available to "native" STARPU applications are available as well for OPENMP programs running on top of STARPU. We further discuss these implementation considerations in Section 4.3

3 BRIEF OVERVIEW OF THE FMM

The FMM is a hierarchical algorithm originally introduced in [18]. It aims to reduce the quadratic complexity of pairwise interactions to a linear or a linearithmic complexity. The FMM is now used in a large range of applications such as vortex methods, boundary element methods (BEM) or radial basis functions. In order to make the paper self-contained, we present a brief overview of the FMM algorithm, but **the reader in a hurry may proceed to Section 4 knowing that the issue we address in Section 4 will be to parallelize Algorithm 1 with a concise OPENMP code, in order to exploit the potential parallelism it provides as expressed in the example DAG in Figure 2b.**

The key-point of the FMM algorithm is to approximate the far-field - the interactions between far-apart particles - while maintaining a desired accuracy, exploiting the property that the underlying mathematical kernel decays with the distance between particles. While the interactions between close particles still remain computed with a direct

Particle to Particle (P2P) method, the far-field is processed using a tree-based algorithm instead. A recursive subdivision of the space is performed in a preprocessing symbolic step (see Figure 1). This recursive subdivision is usually represented with a hierarchical tree data structure and we call the height of the tree h the number of recursions. The type (quadtree, octree, ...) of the tree is related to the




Fig. 1: 2D space decomposition (Quadtree). Grid view and hierarchical view.

dimension of the problem. However, in the current study we use the term *octree* in a generic manner to refer to the FMM tree for any dimension. Figure 1 is an example of an octree showing the relationship between the spatial decomposition and the data structure, and we see that each cell represents its descendants composed of its children and sub-children. The multipole (M) of a given cell represents the contribution of its descendants. On the other hand, the local part (L) of a cell c represents some contributions that will be applied to the descendants of c . These local contributions in a cell c come from the potential of particles/cells that are not included by c . Relying on those recursive data structures, the FMM algorithm proceeds in four steps, namely the upward pass, the transfer pass, the direct pass and the downward pass.

In an upward pass of the FMM, the physical values of the particles are aggregated from bottom to top using the Particle to Multipole (P2M) and Multipole to Multipole (M2M) operators. After this operation, each cell hosts the contributions of its descendants. In the transfer pass, the Multipole to Local (M2L) operator is applied between each cell and its corresponding interaction list at all levels. The interaction list for a given cell c at level l is composed by the children of the neighbors of c 's parent that are not direct neighbors/adjacent to c . After the transfer pass, the local part of all the cells are filled with contributions. The downward pass aims to apply these contributions to the particles. In this pass, the local contributions are propagated from top to bottom with the Local to Local (L2L) operator, and applied to the particles with the Local to Particle (L2P) operator. After these far-field operations, the particles have received their respective far contributions.

Algorithm 1 shows the resulting sequential pseudo-code (for a matter of conciseness, we provide the detailed code of the M2M operator only). The dependencies between the operations occurring within this algorithm can be encoded with a Directed Acyclic Graph (DAG) where vertices represent operators and edges dependencies between them. Figure 2b for instance represents the DAG of the FMM associated with the tiny grid provided in Figure 2a.


Algorithm 1: FMM Sequential Algorithm

```

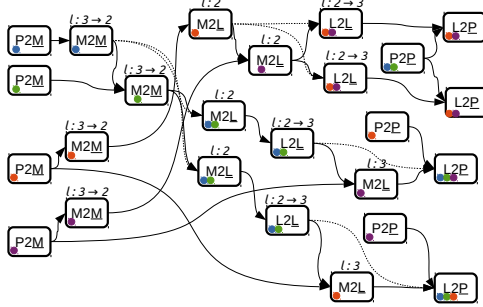
1 function FMM(tree, kernel)
2   // Near-field
3   P2P(tree, kernel);
4   // Far-field
5   P2M(tree, kernel);
6   for l = tree.height-2 → 2 do
7     M2M(tree, kernel, l);
8   for l = 2 → tree.height-2 do
9     M2L(tree, kernel, l);
10    L2L(tree, kernel, l);
11  M2L(tree, kernel, tree.height-1);
12  L2P(tree, kernel);

13 function M2M(tree, kernel, level)
14   foreach cell cl in tree.cells[level] do
15     kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
16     level).multipole);

```



(a) Particles layout, tree height $h = 4$.



(b) Corresponding FMM DAG. The colored circles represent the contributions, such that the L2P contains the far-field potential to apply on the leaves.

Fig. 2: Example: Interaction between 4 particles.

4 BRIDGING THE PERFORMANCE GAP BETWEEN OPENMP-BASED AND NATIVE RUNTIME SYSTEMS-BASED FMM

In a previous study [19], we showed that elaborated and complex FMM parallelization schemes based on OPENMP 3.0 were not competitive against task-based designs natively exploiting runtime systems. On the contrary, we now propose concise OPENMP-based FMM implementations compliant not only with the OPENMP standard but also with the original idea of the OPENMP ARB to make parallel programming on shared-memory machines a user-friendly process. The resulting codes will then be used in Section 5 to assess whether the new revision of the OPENMP standard allows for achieving both high performance and programming productivity.

4.1 Fork-Join schemes

A straightforward FMM parallelization scheme consists in performing a level by level parallelization of each inner loop with a `#pragma omp parallel for` directive. Such a parallelization can be implemented with the sequential algorithm (Algorithm 1) within which the main loop of each operator implementation is prepended with this directive [20], [21], [22]. Algorithm 2 shows how the M2M operator is readily adapted. We name `ff-omp3#for-dyn` this fork-join (ff) approach based on the `#pragma omp parallel for` dynamic loop scheduling directive.

Algorithm 2: `ff-omp3#for-dyn`

```

1 function M2M(tree, kernel, level)
2   #pragma omp parallel for schedule(dynamic, 10)
3   foreach cell cl in tree.cells[level] do
4     kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
5     level).multipole);
6   // Implicit barrier from omp parallel

```

We consider a second variant based on static loop scheduling, where we aim at assigning the same amount of work to each thread. We use the number of interactions per elements as the balancing metric, and a greedy pre-processing to find the correct interval. We name `ff-omp3#for-stat` this static variant in the sequel whose M2M code is provided in Algorithm 3.

Algorithm 3: `ff-omp3#for-stat`

```

1 function M2M(tree, kernel, level)
2   #pragma omp parallel
3   s = thread_interval[level][omp_get_threadnum()].start;
4   e = thread_interval[level][omp_get_threadnum()].end;
5   foreach cell cl in tree.cells[level] from s to e do
6     kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
7     level).multipole);

```

The `#pragma omp task` directive from the revision 3 of the standard allows one to design schemes based on explicit tasks. For that, a single thread inserts the tasks and explicitly performs a `#pragma omp taskwait` call before moving to the next loop such as illustrated in Algorithm 4. Using this

Algorithm 4: `ff-omp3#task`

```

1 function M2M(tree, kernel, level)
2   #pragma omp parallel
3   #pragma omp single
4   foreach cell cl in tree.cells[level] do
5     #pragma omp task
6     kernel.M2M(cl.multipole,
7     tree.getChildren(cl.mindex,
8     level).multipole);
9   #pragma omp taskwait

```

primitive still induces many synchronizations. The barriers after each parallel section indeed require the main thread to wait for all the others before proceeding and creating the next section. These global barriers guarantee the coherency of the algorithm by ensuring that when a level/operator

is computed, all the required data are ready and no race-conditions are possible. However, they still lead to a fork-join approach. We hence call *ff-omp3#task* (fork-join scheme based on the *#pragma omp task* directive from the revision 3 of the standard) this algorithm in the sequel.

4.2 Task-based schemes

The main strength of task-based paradigms is to remove global synchronizations and potentially execute a task as soon as its predecessors (the tasks it depends on) are completed (and, of course, that a resource is available to process it). From this point of view, the introduction of the *#pragma omp task* directive in the version 3.1 of the standard can be viewed as a partial support for task-based scheme that the *depend* clause completed only with the revision 4.0. In the sequel, we will therefore refer to task-based schemes only those either relying on OPENMP 4.0 *task* and *depend* constructs or equivalent low-level native runtime directives (STARPU in our case). We name this scheme *tb-omp4#task#dep* for short, whether or not it is implemented with OPENMP directives or native STARPU constructs.

4.2.1 Task-based scheme specific granularity control




Fig. 3: Group tree example for $n_g = 3$.

We showed in [19] that the overhead of task management however highly penalizes task-based FMM approaches (at least when using a fully-featured runtime support such as STARPU) and we proposed to increase the granularity of tasks in order to limit their number. To do so, we introduced a new data structure, the group tree, which is an octree where a number n_g (the *group size*) of consecutive leaves or cells following the Morton index are allocated together and seen as a single element. Therefore, the same cells/particles exist in an octree and a group tree, but the tasks and subsequent dependencies are not on the cells/leaves but on groups of cells/leaves. A simplified representation of a group tree is shown in Figure 3.




Fig. 4: Group tree memory block view, where C and P refer to cells and particles respectively.

Since the scope of the present study aims at discarding complex parallel implementations of the FMM, we furthermore introduce a new design of the group tree from which we can naturally extract extra parallelism. In this new design, the symbolic and numerical data associated with the particles are split. The symbolic parts are composed of the values that remain constant during the computation as the physical values and particles' positions. In the case of cells, the numerical data associated with the downward and upward passes are also split. For the cells, the symbolic data are their coordinates and Morton indexes. Figure 4 illustrates this new design of the group data structures we have implemented in SCALFMM for the purpose of the present study. We have implemented two versions of this

Algorithm 5: *tb-omp4#task#dep* scheme with OPENMP 4.0 directives

```

1 function FMM(tree, kernel)
2   #pragma omp parallel
3     #pragma omp single
4       // Near-field
5       P2P_taskdep(tree, kernel);
6       // Far-field
7       P2M_taskdep(tree, kernel);
8       for l = tree.height-2 → 2 do
9         M2M_taskdep(tree, kernel, l);
10      for l = 2 → tree.height-1 do
11        M2L_taskdep(tree, kernel, l);
12      for l = 2 → tree.height-2 do
13        L2L_taskdep(tree, kernel, l);
14      // Merge
15      L2P_taskdep(tree, kernel);
16      #pragma omp taskwait

17 function M2M_taskdep(tree, kernel, level)
18   foreach cell cl in tree.cells[level] do
19     #pragma omp task depend(inout:cl.multipole) \
20
21       depend(in:tree.getChildren(cl.mindex,
22                                level).multipole)
23       kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
24                                level).multipole);

```

proposed *tb-omp4#task#dep* scheme. Algorithm 5 shows a version written for OPENMP while Algorithm 6 shows the equivalent code written for STARPU.

4.3 Runtime support

The question we address in this paper is whether OPENMP or OPENMP-like FMM codes can achieve a performance competitive with an optimized code natively written with a modern runtime system (STARPU in our case). For that, we considered some OPENMP codes written in the spirit of historical OPENMP constructs (Section 4.1) as well as new compact task-based algorithms based on modern directives (Section 4.2). We now present the runtime support for those schemes.

One option for executing those algorithms is to rely on a lightweight dedicated OPENMP runtime system such as the reference LibGOMP library provided with GCC (GCC / LIBGOMP). All fork-join schemes and task-based OPENMP schemes discussed in sections 4.1 and 4.2, respectively, can be executed with this support. The paths going through

Algorithm 6: *tb-omp4#task#dep* scheme with native STARPU directives

```

1 function FMM(tree, kernel)
2   // Near-field
3   P2P_starpu(tree, kernel);
4   // Far-field
5   P2M_starpu(tree, kernel);
6   for l = tree.height-2 → 2 do
7     M2M_starpu(tree, kernel, l);
8   for l = 2 → tree.height-1 do
9     M2L_starpu(tree, kernel, l);
10  for l = 2 → tree.height-2 do
11    L2L_starpu(tree, kernel, l);
12  // Merge
13  L2P_starpu(tree, kernel);
14  starpu_task_wait_for_all()

15 function M2M_starpu(tree, kernel, level)
16  foreach cell cl in tree.cells[level] do
17    starpu_insert_task(READ_WRITE, cl.multipole,
18                      READ, tree.getChildren(cl.mindex,
19                      level).multipole,
19                      &kernel.M2M);

```

the arrows in figures 5a and 5b show the corresponding software stacks. Another option for executing the task-based algorithms proposed in Section 4.2 is to execute their implementations written with native runtime directives (such as Algorithm 6) directly on top of the task-based STARPU runtime, as shown with the path going through the red arrow in Figure 5b.

We furthermore propose to bridge the gap between those two cases by executing OPENMP codes using a native task-based runtime system (STARPU in our case). For that, we rely on the source-to-source KLANG/STARPU compiler¹ which translates an OPENMP code into a STARPU code. The most straightforward translation is the one of task-based algorithms. Indeed, we can readily transform OPENMP tasks (called *explicit* tasks in the OPENMP specification) into STARPU tasks. For instance, Algorithm 5 can be converted by the compiler to be strictly equivalent to the native STARPU source code proposed in Algorithm 6. The path going through the yellow arrow in Figure 5b shows the corresponding software stack. However, the OPENMP execution model specification differs slightly with the native STARPU execution model with respect to the *main* thread. In the native STARPU execution model, the *main* thread only executes the sequential part of the application: it inserts tasks and eventually waits for their completion, but it is not bound to any core and is not involved in computing any tasks. Therefore, this thread may freely move over the cores and, though the cost of submitting tasks is often low in regard of their execution cost, it still may visibly alter the execution of the *worker* threads computing light workload tasks, in regard of which the submission time is not negligible. The OPENMP execution model specifies that the *main* thread also acts as a *worker* thread: it takes an active part in executing tasks and is bound to a core. Thus, we extended STARPU such that an OPENMP compliant execution model is enforced when STARPU is used

1. KLANG/STARPU is presented for the first time in the present article, see <http://kstar.gforge.inria.fr/> for details.

through the KLANG/STARPU compiler, with respect to the *main* thread. The execution model of STARPU when used through its native directives has been left unmodified however, which may in some cases account for slightly different execution behaviour between native STARPU programs and KLANG/STARPU programs. An example of such a slightly differing behaviour will indeed be observed on the task-based schemes detailed timing study in Sec. 5.4.3.

The fork-join *ff-omp3#task* scheme based on OPENMP 3 independent tasks (Algorithm 4) is also readily implemented by mapping those explicit tasks onto STARPU tasks while furthermore ensuring the `taskwait` synchronizations with `starpu_task_wait_for_all()` barriers. The language specification also states that fork-join parallel regions are logically expressed as one task (called *implicit* task) for each participating thread. KLANG/STARPU takes advantage of this statement to express OPENMP fork-join models as a direct mapping of these implicit tasks on STARPU tasks, following the path through the yellow arrow in Figure 5a. The loop scheduling attribute of an OPENMP parallel for region further specifies how the iterations of the parallel loop are assigned to the participating implicit tasks. The *static* scheduling version, as in Algorithm 3, statically assigns chunks of the iteration range to each participating task. The *dynamic* scheduling version, as in Algorithm 2, instead lets the participating tasks contend at runtime to pick iteration range chunks. Thus, the *static* loop scheduling advantage is to incur less processing overhead, while the advantage of the *dynamic* scheduling is to enable better load balancing if the workload is not uniform across the iteration range.

4.4 Enhancing task-based schemes with priority and commutativity

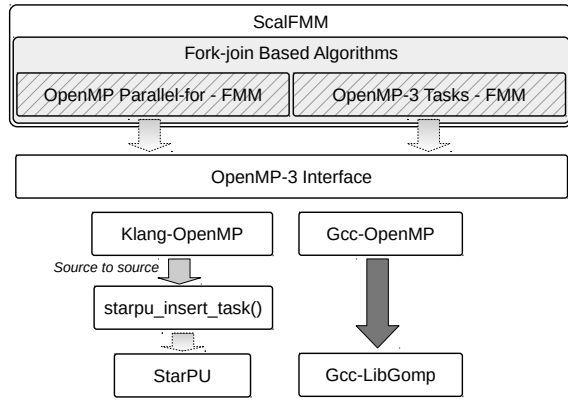
We [19] showed that expressing fine-grain dependencies such as the ones encoded by the DAG in Figure 2b and relying on a group scheme could achieve high performance in general, but could be further accelerated with additional hints. The first required hint, to maintain high performance on a large range of configurations, consists in ensuring a swift progress along the critical path. For that, we attribute priorities to the tasks as defined in Table 1. In the native STARPU case, this leads to slight changes when creating tasks as shown in Algorithm 7 for the M2M operator. Note that we distinguish the priorities between computation within (-inner) and between (-outer) groups (see [19] for more details).

Operator/Task	P2M	M2M	P2P	M2L-inner*	M2L-outer*
Priority	0	1	2	$\ell + 1$	$\ell + 2$
Operator/Task	L2L*	P2P-outer	M2L-inner($h-1$)	M2L-outer($h-1$)	L2P
Priority	$\ell + 3$	$(h-3) * 3 + 2$	$(h-3) * 3 + 3$	$(h-3) * 3 + 4$	$(h-3) * 3 + 5$

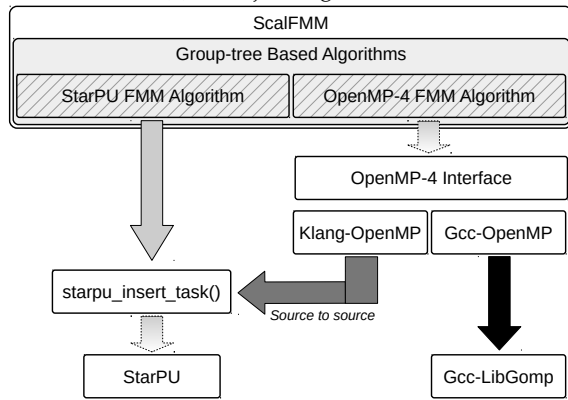
TABLE 1: Priorities for the different operators/tasks.

In KLANG/STARPU, we follow the OPENMP specification to support priorities in the task declaration. Algorithm 8 shows the use of *priority* clause so that the original code is only marginally modified. We name *tb-omp4#task#dep-P* the resulting code and we emphasize that it is compatible with the OPENMP 4.5 standard.

Another limit highlighted in [19] is that the order in which the tasks are created in Algorithm 5 is extremely



(a) Fork-join algorithms.



(b) Task-based algorithms.

Fig. 5: Software layers.

Algorithm 7: *tb-omp4#task#dep-P* scheme with native STARPU directives

```

1 function M2M_starpu(tree, kernel, level)
2   foreach cell cl in tree.cells[level] do
3     starpu_insert_task(PRIORITY, M2M_PRIO,
4                       WRITE, cl.multipole,
5                       READ, tree.getChildren(cl.mindex,
6                       level).multipole,
7                       &kernel.M2M);

```

Algorithm 8: *tb-omp4#task#dep-P* scheme with OPENMP 4.5 directives

```

1 function M2M_taskdep(tree, kernel, level)
2   foreach cell cl in tree.cells[level] do
3     #pragma omp task priority(M2M_PRIO)
4     depend(inout:cl.multipole) \
5     depend(in:tree.getChildren(cl.mindex,
6     level).multipole)
7     kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
8     level).multipole);

```

sensitive. Different valid task creation orderings, for this algorithm, lead to different instantiations of the DAG. The fundamental reason is that the *inout* dependence mode specified for data element being updated by a task, indeed not only forces mutual exclusion of tasks with respect to

that dependence data, it, furthermore, forces sequential consistency as well. For instance, if a task T_1 and a task T_2 are submitted in that order to update a piece of data A , and if T_2 becomes ready to run before T_1 , T_2 will nevertheless *not* be allowed to start before T_1 is completed. This results in tasks accessing that piece of data being forcibly executed in the task creation order. For many algorithms, this is the expected behaviour. For FMM Algorithm 5, however, this overconstrains the set of valid task schedules: The Algorithm accumulates — for a given cell — the incoming contributions from the cell's neighbour particles or cells. This accumulation is inherently commutative, since incoming contributions may update a cell in any order, meaning that multiple execution orders are valid. As a consequence, forcing an arbitrary execution order may unnecessarily delay the execution of ready tasks, wasting parallelism.

To achieve high performance in [19], despite these overconstraining *inout* dependencies, we resorted to implementing complex loop unrolling right in the application code for tuning the tree traversal and subsequently the obtained DAG. This method would however not fit the programming productivity objective of the present study, to achieve high performance, while preserving a concise and elegant application code.

Moreover, while one might think about using a reduction construct, this would not be the best option here, since parallel reductions are implemented using per-thread privatized buffers and a parallel tree-shaped reduction scheme. Parallel reductions are a win when the overall application parallelism is *limited*, that is, when the memory footprint and cost of the extra privatized buffer management is hidden by the extra parallelism offered by the reduction. In the FMM case, the large number of particles and cells being concurrently updated brings a large amount of parallelism by itself. Thus, using a parallel reduction would result in large memory footprint and processing costs from the privatized buffer management, for little parallelism gain and significant performance penalty.

Therefore, we propose to introduce a new data dependence mode, to express commutative update operations. The commutative update mode still enforces mutual exclusion on the piece of data being updated, but relaxes the ordering of the incoming updates. From the application point of view, this involves modifying the task insertion step as show, for instance, with a M2M task on Algorithm 9 (example given for the native STARPU API. Although the

Algorithm 9: *tb-omp4#task#dep-C* scheme with native STARPU directives

```

1 function M2M_taskdep(tree, kernel, level)
2   foreach cell cl in tree.cells[level] do
3     starpu_insert_task(COMMUTE, cl.multipole,
4                       READ, tree.getChildren(cl.mindex,
5                       level).multipole,
6                       &kernel.M2M);

```

OPENMP board decided to introduce a reduction mode in future revisions, there is (at the time of this writing) no clause expressing the commutative update dependence mode for the time being. We propose to add a *commute* data dependence mode to the *depend* clause, for expressing

commutative update operations in OPENMP. An alternative approach [23] relies on using fine grained OPENMP *atomic* based synchronizations. However, it suffers from the drawback to be break the flow of data dependencies, while our approach seamlessly integrates with the data dependence model by simply adding a new dependence mode. Moreover, using OPENMP *atomic* constructs requires a higher level of expertise from the programmer than merely replacing an *inout* mode with a *commute* dependence mode.

Algorithm 10: *tb-omp4#task#dep-C* scheme with OPENMP extensions

```

1 function M2M_taskdep(tree, kernel, level)
2   foreach cell cl in tree.cells[level] do
3     #pragma omp task depend(commute:cl.multipole) \
4
5     depend(in:tree.getChildren(cl.mindex,
        level).multipole)
        kernel.M2M(cl.multipole, tree.getChildren(cl.mindex,
        level).multipole);

```

Algorithm 10 shows the concise impact on the resulting code. We also implemented a support of the proposed *commute* data access mode within KLANG/STARPU to assess the resulting *tb-omp4#task#dep-C* (or *tb-omp4#task#dep-CP* when combined with *priority*) code. As a result, the




Fig. 6: Symbolic meta-DAG for the FMM associated with Figure 2b where tasks with an *OR* block can be executed in any order instantiated dynamically.

underneath runtime system is not provided with a static DAG as the one from Figure 2b but can instantiate it dynamically depending on the order in which tasks are completed. Figure 6 illustrates the dynamic opportunities that are encoded by the data structure that substitutes that DAG. The resulting compact code can bridge the gap between OPENMP and native task-based runtime systems, allowing for maximizing both productivity (discussed in this section) and performance (discussed below).

5 EXPERIMENTAL STUDY

5.1 Experimental setup

We illustrate our discussion with two types of particle distributions assessed on two multicore platforms. The first machine is a 24 cores platform (referred to as 24C in the sequel) composed of 2 dodeca-core Haswell Intel Xeon E5-2680 processors running at 2.8GHz and with 128GB of shared memory. It can be considered has a common modern multicore processor; we will therefore mainly focus on this platform. The second one is a 96 cores platform (referred to as 96C in the sequel) of the SGI UV2000 series, composed

of 4×2 dodeca-core Haswell Intel Xeon E5-2680 running at 2.8GHz and with 132GB of shared memory, interconnected in an all-to-all topology via SGI's NUMalink. This second machine can be viewed as an extreme multicore processor but is interesting for assessing the validity of our claims, especially for the task-based approach which is expected to deliver a higher level of scalability. We use STARPU 1.3 (5/4/2016), KLANG/STARPU 1.0 (1/4/2016), SCALFMM 1.4 (5/10/2016), the GNU compiler GCC 4.9, the Intel MKL Blas & FFTW 11.2 on both platforms.



(a) Cube (volume). (b) Ellipsoid (surface).

Fig. 7: Particle distributions considered in this study.

We consider two types of particle distributions. *Cube (volume)* distributions (Figure 7a) are composed of particles uniformly distributed in a unit box, leading to a regular octree and a high and well balanced amount of work for each cell/leave. *Ellipsoid (surface)* distributions (Figure 7b) are composed of particles distributed on the surface of an ellipsoid with a high density at the poles, leading to an irregular octree and highly variable amount of work associated with the nodes of the octree.

We tuned the FMM as follows. For all our simulations, we use the interpolation based FMM on equispaced grid points (UFMM) from [24] on the Laplacian kernel with an order $l = 5$ corresponding to an intermediate accuracy of 10^{-5} (see [24] for details). The choice of the height h of the octree balances the amount of work between the near and far fields. We select the height h that minimizes the sequential execution time. In the sequel, we focus on the parallel behavior, but we provide here some sequential execution times using *fj-omp3#for-dyn* to solve the FMM for different numbers of particles (N) to allow the reader for having some orders of magnitude in mind and better understand the challenge their parallelization represents. The FMM solution for a cube distribution is obtained in about 57 s and 1100 s in sequential for a number N of particles equal to $N = 10^6$ and $N = 10^8$, respectively. In the ellipsoid case, the solution requires about 6 s and 980 s in sequential for a number $N = 10^6$ and $N = 10^8$ of particles, respectively. The *fj-omp3#for-dyn* algorithm is parametrized with a chunk size equal to 10 (see Section 5.3). For task-based approaches, the group size n_g is tuned to minimize the parallel execution time on the 24C machine (see Section 5.4.1) and we use the same value on the 96C platform. Given the high number of test combination cases studied, and the substantial cumulated execution time of the whole testing campaign, each measure sample is taken only once. Thus, some isolated measured samples may exhibit experiment bias.

5.2 Performance metrics

5.2.1 Normalized efficiency

A convenient metric commonly employed to assess the success of a parallelization scheme is the *parallel efficiency* [25] (or *efficiency* for short) noted $e(p)$ in the sequel. Given a reference sequential execution time t_1 and a parallel execution time t_p with p processing units, it is computed as follows: $e(p) = \frac{t_1}{p \times t_p}$. The reference time t_1 must correspond to the fastest sequential algorithm [25]. When considering a fully-featured numerical library, there are however many factors involved and it might be cumbersome to guarantee that the sequential reference is optimal. As mentioned above, we rely on a highly optimized kernel [24] and we tune the height h of the octree so as to minimize the sequential execution time. In the task-based approach, the group size n_g may however impact the sequential performance. We tuned n_g so that it minimizes the parallel execution time (see Section 5.4.1). Large group sizes could slightly improve the sequential performance on our smaller test cases but we decided to run the sequential and the parallel algorithm with the same group size to make the analysis clearer. Because our goal is to compare multiple parallelization schemes that can be supported with multiple execution supports, we consistently use the same sequential reference over all the study. However, we distinguish the reference used when comparing fork-join schemes with each other (Section 5.3) from the one used for comparing task-based schemes with each other (Section 5.4.2). The main difference is that task-based schemes rely on group data structures and benefit from a better locality, which tends to improve both the sequential and parallel overall execution time. We compute the normalization reference t_1 using GCC/LIBGOMP on *ff-omp3#for-dyn* for Section 5.3, and using GCC/LIBGOMP on *tb-omp4#task#dep* for Section 5.4.2. We refer to the resulting normalized $e(p)$ metrics as the “normalized parallel efficiency,” or simply as the *normalized efficiency* when there is no ambiguity.

5.2.2 Detailed timings

While the parallel efficiency (or in our case the normalized efficiency defined above) is a very convenient measure for assessing the overall success of a considered parallelization, one may want to analyze in more details the behavior of a method. Indeed, multiple effects can lead to a non optimal parallel efficiency. In a task-based code running on top of a runtime system, the time spent over all processes, or *cumulated time* $\mathcal{T}(p)$, can be cast into time spent in performing actual tasks ($\mathcal{T}^{task}(p)$), in runtime management ($\mathcal{T}^{runtime}(p)$) or scheduling ($\mathcal{T}^{scheduling}(p)$) operations, and idle time ($\mathcal{T}^{idle}(p)$). The overall cumulated time $\mathcal{T}(p)$ is homogeneous to *processing unit* \times *second* and satisfies

$$\begin{aligned} \mathcal{T}(p) &:= t_p \times p \\ &:= \mathcal{T}^{task}(p) + \mathcal{T}^{runtime}(p) + \mathcal{T}^{scheduling}(p) + \mathcal{T}^{idle}(p). \end{aligned}$$

Due to potential contention (caches, buses, ...), the actual computation may be slowed down, leading to a higher time spent in tasks ($\mathcal{T}^{task}(p)/\mathcal{T}^{task}(1) > 1$). Furthermore, the runtime system must consume time in order to ensure the overall progress (task insertion, data consistency, ...). In

our case, we furthermore distinguish the time spent within the runtime from the time spent in scheduling. We want to assess whether those costs are reasonable with respect to the ideal time spent in tasks ($\mathcal{T}^{runtime}(p)/\mathcal{T}^{task}(1) \ll 1$ and $\mathcal{T}^{scheduling}(p)/\mathcal{T}^{task}(1) \ll 1$). Finally, idle time results from a lack of concurrency possibly combined with suboptimal scheduling decisions. We again assess this effect with respect to the ideal time spent in tasks ($\mathcal{T}^{idle}(p)/\mathcal{T}^{task}(1)$). Consequently, STARPU has been instrumented to separately record the time spent in application tasks, scheduler algorithm, runtime management and in the idle loop, respectively so as to report such detailed timings for both the KLANG/STARPU and the STARPU supports.

5.3 Performance of fork-join schemes

5.3.1 Normalized efficiency

Figure 8 presents the parallel efficiencies of the *fork-join* schemes discussed in Section 4.1 normalized by the GCC/-LIBGOMP *ff-omp3#for-dyn* sequential reference on the 24C platform for the cube and ellipsoid distributions, respectively. We assess all three *ff-omp3#for-dyn*, *ff-omp3#for-stat* and *ff-omp3#task* fork-join parallelization schemes proposed in algorithms 2, 3 and 4. Both the GCC/LIBGOMP and KLANG/STARPU source-to-source OPENMP compiler / runtime system frameworks are considered. The main observation is that the KLANG/STARPU support for the *-#pragma omp parallel for* directive is competitive against the lightweight GCC/LIBGOMP support for both the *ff-omp3#for-dyn* and *ff-omp3#for-stat* schemes, as long as the number of particles is not extremely low ($N = 10^6$).

The independent tasks support of KLANG/STARPU shows more overhead than GCC/LIBGOMP due to the heavier weight of the underlying STARPU task covers compared to the extremely lightweight GCC/LIBGOMP ones. This extra cost especially appears on the *ff-omp3#task* scheme, but is also visible on the sequential data point for most schemes. However, this overhead is compensated by a better scalability and smoother behavior of KLANG/STARPU task support. These figures also show the interest of the *ff-omp3#for-stat* balanced scheme on such small cases, both for GCC/LIBGOMP and KLANG/STARPU. On the contrary, the *ff-omp3#for-dyn* scheme is inefficient on these small cases (top/left plots on the cube) since the participating threads heavily contend on work-sharing the small iteration ranges of the parallel for loops.

A large number of particles (bottom plot row on the cube distribution), or a more complex structure prone to load imbalance (all plots on the ellipse distribution), makes the *ff-omp3#for-stat* scheme redundant, changing its benefits observed for small cube cases into a penalty here. Conversely the *ff-omp3#for-dyn* scheme generally performs better: the larger amount of particles reduces the contention on parallel for iteration ranges, while the dynamic loop scheduling offers better load balancing by design. The *ff-omp3#task* scheme shows mediocre to bad results for GCC/LIBGOMP and worse results for KLANG/STARPU. A large number of independent tasks generates runtime processing overhead without having the opportunity to offer some benefit in return. The phenomenon is further emphasized by increasing the height of the tree, thus, increasing the number of tasks by

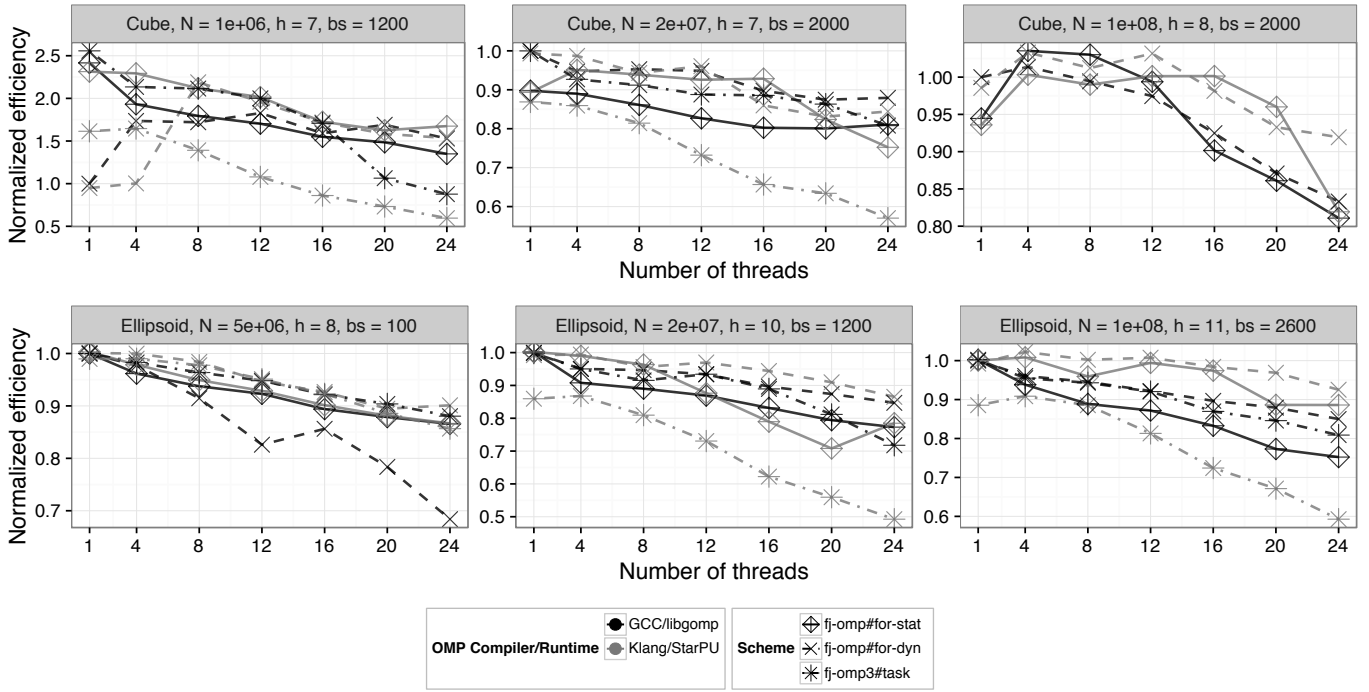


Fig. 8: Parallel efficiency normalized by the GCC/LIBGOMP *fj-omp3#for-dyn* sequential reference for the cube distribution (volume), above, and the ellipse (surface), below, on the 24C machine. The different plots do not use the same scale for the *Normalized efficiency* axis. For the largest cases (row 2, col 2–3), the *fj-omp3#task* scheme does not run to completion, due to the overwhelming amount of tasks being created.

creating lower levels with much more cells. Moreover, the *fj-omp3#task* schemes using STARPU and KLANG/STARPU supports fail to compute the largest simulations ($N = 5 \cdot 10^7$ and $N = 10^8$) only on the cube, because the number of generated tasks is too important and lead to out-of-memory issues. Indeed, for $N = 5 \cdot 10^7$ and $h = 8$ there are more than 5M tasks just for the *M2L* at leaf level.

5.3.2 Detailed timings

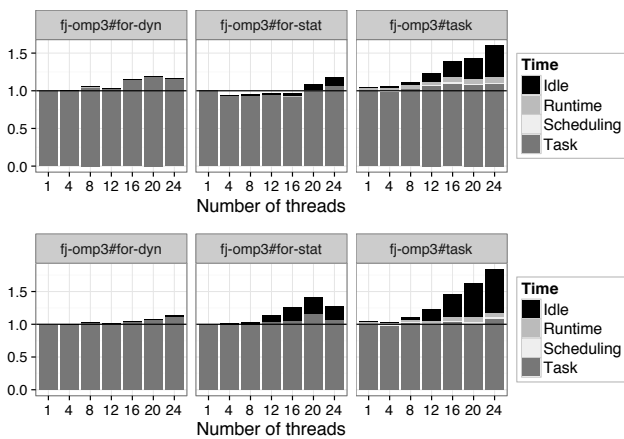


Fig. 9: Detailed *fork-join* timings (cumulated time with respect to the ideal time spent in tasks) for the cube (volume) distribution, above ($N = 2 \cdot 10^7$, $h = 7$), and the ellipsoid (surface), below ($N = 2 \cdot 10^7$, $h = 10$), with KLANG/STARPU on platform 24C.

Figure 9 shows the resulting detailed timings (see Section 5.2.2) for running all fork-join schemes with the KLANG/STARPU support. We see that the *fj-omp3#for-dyn* scheme incurs small runtime overhead and idle times. However, the time to compute the tasks increases with the number of threads, since all the threads have to dynamically compete for parallel loops work-sharing. On the contrary, the work-sharing division is done statically with the *fj-omp3#for-stat* scheme.

The *fj-omp3#for-stat* scheme shows an increasing idle time, though, which is correlated to the $\mathcal{T}^{task}(p)$ task self execution time improvement, as the number of threads increases. This acceleration is mainly due to the resulting increased data locality and reuse, because the work is split into sub-trees where most of parents/children and neighbors are included. Yet, this acceleration in the computation of the tasks also makes the static iteration space split effectively more unbalanced in terms of execution time, which increases the idle time.

Finally, the *fj-omp3#task* scheme shows significant runtime overhead and dramatic increasing $\mathcal{T}^{task}(p)$ due to the small granularity. We remind that we create one task per operation, which is equivalent to a dynamic parallel for with a chunk size of one. Meanwhile, it confirms the intuition that it is more costly in for the runtime to manage additional tasks rather than sharing work based on an index increment.

5.4 Performance of task-based schemes

5.4.1 Granularity

We introduced the group-tree and the associated granularity parameter n_g in Section 4.2. Plots on figure 10 show the

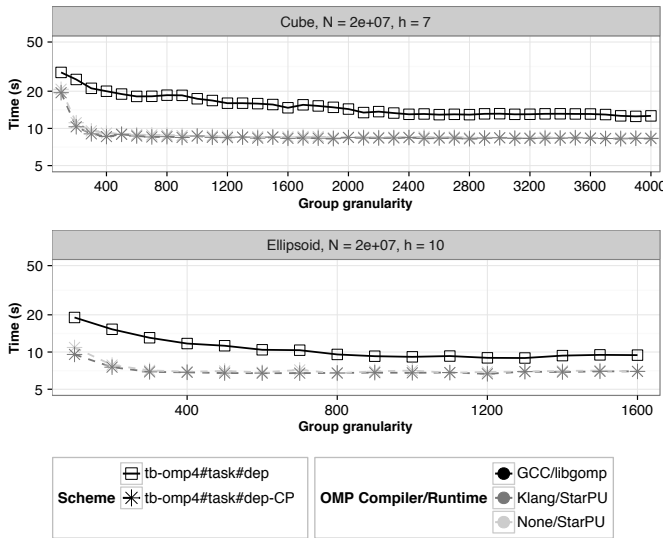


Fig. 10: Execution time for the cube (volume) distribution, above, and the ellipsoid (surface), below, for different values of the granularity n_g on platform 24C using 24 threads.

measured impact of n_g on the performance of parallel executions relying on GCC/LIBGOMP, KLANG/STARPU and STARPU. As expected when the granularity parameter is too small, all three scheme/support combinations tested perform poorly because the task computational weight is lower and the total amount of tasks generated is higher. Each runtime system introduces per-task managing costs such as allocating data structures, queuing the task, resolving its dependencies and scheduling it; thus, when the task granularity is small, this per-task overhead becomes significant in regard of the task execution time. Moreover, the larger number of tasks means that this per-task overhead adds-up to a higher penalty on the application execution time. On the other hand, increasing the granularity improves the execution time up to a given limit related to the test case sensitiveness to load imbalance: This limit can therefore especially be observed, for GCC/LIBGOMP, on the ellipsoid distribution plot around $n_g = 1400$, while for the cube distribution, the execution time still remains at the plateau level at a much larger $n_g = 4000$. STARPU appears less sensitive to load imbalance at high n_g values for both particle distributions. For each considered particle distribution, we select the granularity that leads to the lowest execution time over all scheme/support combinations on the maximum number of cores available. All executions related to that particle distribution are then applied with this granularity. The execution time of the ellipsoid test case is lower than the uniform equivalent with the same number of particles because in the case of surface we have a sparse octree and thus the far-field is then done at a lower cost.

5.4.2 Normalized efficiency

Figure 11 presents parallel efficiencies of the *tb-omp4#task#dep* scheme using GCC/LIBGOMP, KLANG/STARPU and STARPU supports, normalized by the GCC/LIBGOMP *tb-omp4#task#dep* sequential reference on the 24C platform, for the cube and ellipsoid distributions, respectively. GCC/-

LIBGOMP, KLANG/STARPU and STARPU perform very similarly on the *tb-omp4#task#dep* scheme. Again, for a single thread, GCC/LIBGOMP is faster than the other supports, showing that the GCC/LIBGOMP tasks are lighter. The figure also demonstrates the interest of the commutative dependencies extension (see Algorithm 10) on the *tb-omp4#task#dep-C* scheme using either KLANG/STARPU or STARPU. Since the test is performed on a uniform grid distribution of particles, the potential incoming contributions of a particle or group are numerous. As the number of threads increases, the ability to integrate these contributions on the fly, in no particular order, leads to a high benefit for the commutative dependencies extension. The *tb-omp4#task#dep* scheme is penalized by forcing the integration of incoming contributions in the sequential data dependence order, regardless of their availability.

If the particle distribution is less uniform, the advantage of the commutative dependencies may however eventually get canceled by the processing overhead it incurs. For instance, when using the ellipsoid distribution and a deep tree, collecting contributions for lots of small, highly-clustered tasks constitutes most of the work. Thus a task does not have to wait a lot for contributions from other tasks to be ready, even when forcing the seemingly overkill sequential data dependence order.

We further evaluated the benefit of prioritizing tasks, using the *tb-omp4#task#dep-CP* scheme, to favor tasks from the critical path (see Algorithm 8). As shown on figure 11, the extra benefit of priorities is not as convincing as the benefit of commutative dependencies was on the uniform particle distribution. Nevertheless, the fact that both the commutative dependencies and the task priorities are supported by KLANG/STARPU makes it straightforward to explore their effectiveness. This is especially valuable in the case of an application such as SCALFMM whose execution behavior highly depends on the dataset characteristics.

Figure 12 shows results obtained on a large shared-memory 96-core machine. For the uniform grid particle distribution (cube), the *tb-omp4#task#dep-C* scheme's commutative dependencies shows a much better scalability than *tb-omp4#task#dep*. The very large number of threads stresses the necessity to integrate incoming contributions as soon as possible, which commutative dependencies enable. For the ellipsoid particle distribution, the large number of threads also calls for using commutative dependencies. Moreover, the priorities of scheme *tb-omp4#task#dep-CP* improves the efficiency slightly over *tb-omp4#task#dep-C*. Indeed, while it is important to allow integrating contributions in no particular order, the deep tree structure of this test case however requires processing tasks in a sensible, critical path-aware order for maximizing parallelism. These results show the excellent strong scalability of the *tb-omp4#task#dep-CP* scheme up to 96 cores.

5.4.3 Detailed timings

From KLANG/STARPU detailed timing results in Figure 13, it appears that for all configurations, the overhead of the parallelization compared to the $\mathcal{T}^{task}(p)$ task self execution time is negligible, overall. The idle time showing up beyond 12 threads on the baseline *tb-omp4#task#dep* cases is

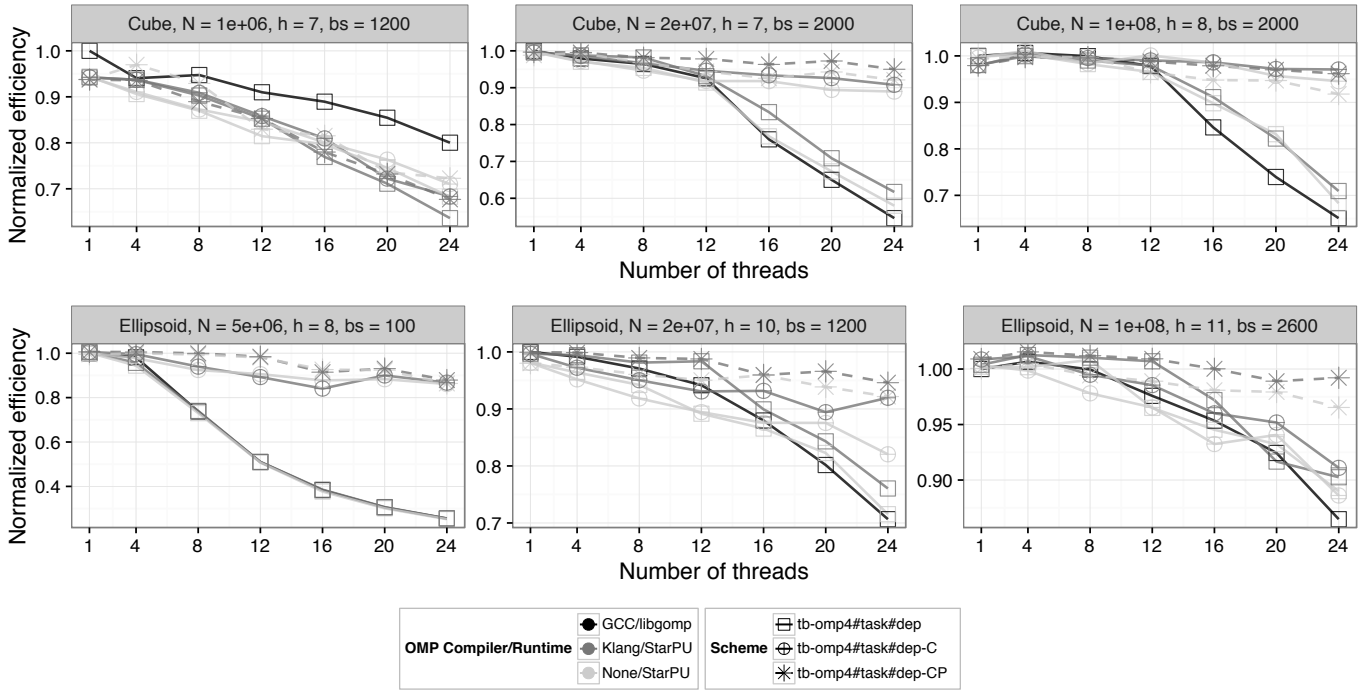


Fig. 11: Normalized efficiency for cube (volume, above) and ellipse (surface, below) on platform 24C from 1 to 24 threads. The different plots do not use the same scale for the *Normalized efficiency* axis.

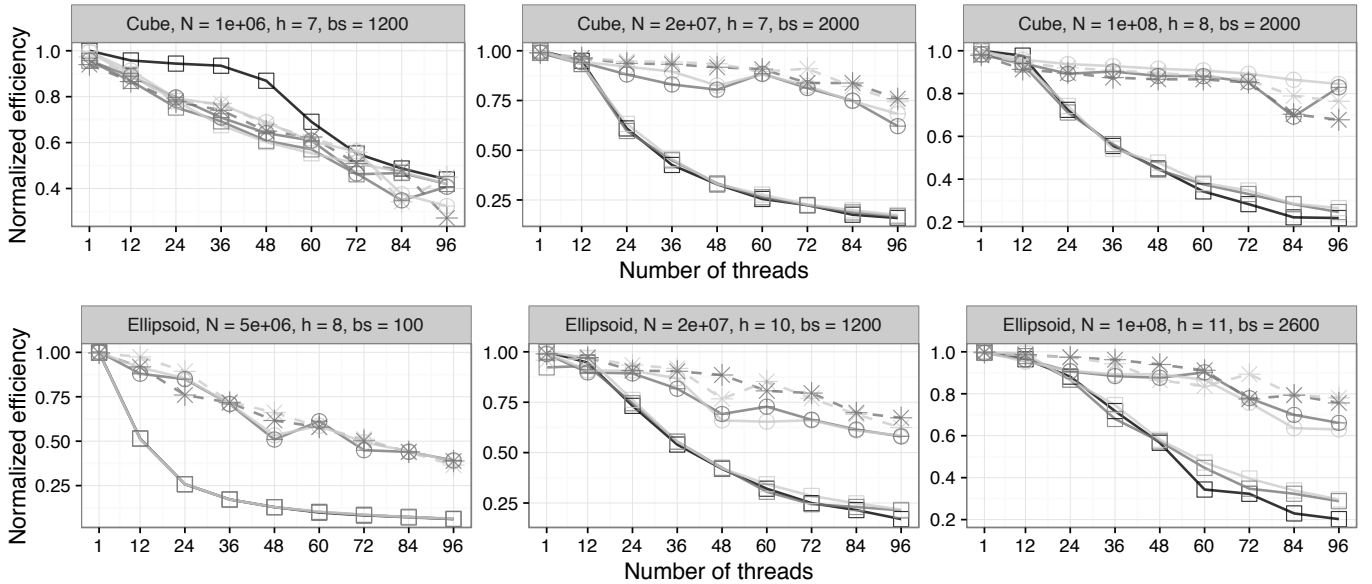


Fig. 12: Normalized efficiency for cube (volume, above) and the ellipse (surface, below) on 96C from 1 to 96 threads. The different plots do not use the same scale for the *Normalized efficiency* axis. Legend is the same as Fig. 11

mainly due to the load imbalance resulting from the over-constrained sequential consistency enforced on concurrent particle/cell updates discussed in Sec. 4.4. In contrast, the effectiveness of the commutative dependencies extensions on reducing idleness is clearly emphasized, especially when used in cooperation with task priorities.

The STARPU detailed timing results in figure 14 may appear surprising at first sight, since the idle time overhead slightly increases from the one obtained with KLANG/-STARPU, without a seemingly logical scheme. This comes from a behavioural difference of the native execution model

of STARPU and the OPENMP compliant execution model provided through KLANG/STARPU with respect to the *main* application thread, as discussed in Section 4.3. In the OPENMP model, the *main* thread takes part both in submitting and in executing tasks, while in the STARPU native execution model, the *main* thread only submits tasks, is not part of the worker threads team, and is not bound to any core by default. In that later case, the *main* thread may therefore compete with worker threads for core access.

Both KLANG/STARPU and STARPU detailed results also emphasize the benefit of exploiting the commutativity (*tb-*

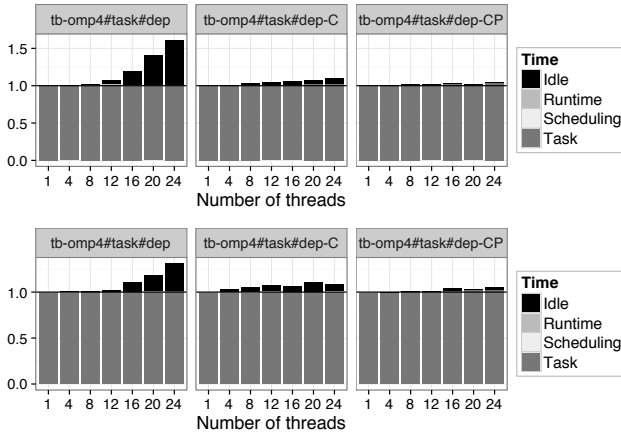


Fig. 13: Detailed *task-based* timings (cumulated time with respect to the ideal time spent in tasks) for the cube (volume) distribution, above ($N = 2 \cdot 10^7$, $h = 7$), and the ellipsoid (surface), below ($N = 2 \cdot 10^7$, $h = 10$), with KLANG/-STARPU on platform 24C.

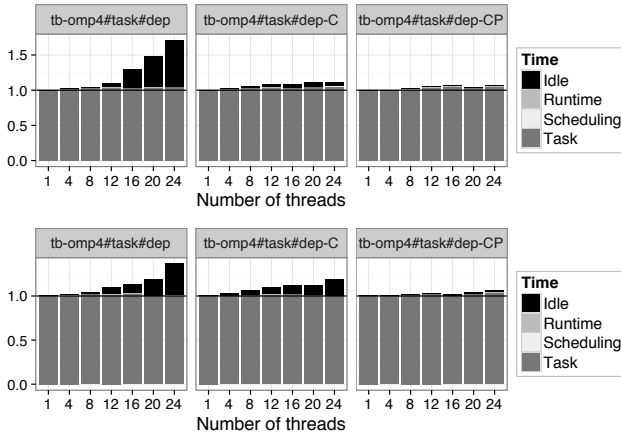


Fig. 14: Detailed *task-based* timings (cumulated time with respect to the ideal time spent in tasks) for the cube (volume) distribution, above ($N = 2 \cdot 10^7$, $h = 7$), and the ellipsoid (surface), below ($N = 2 \cdot 10^7$, $h = 10$), with STARPU on platform 24C.

omp4#task#dep-C scheme) for reducing the idle time, further enhanced using priorities (*tb-omp4#task#dep-CP* scheme), to the point where it becomes almost negligible.

6 CONCLUSION

In this paper, we presented a study of the SCALFMM library port on OPENMP using first a *legacy fork-join* parallelizing approach and then a OPENMP 4 era *dependent tasks* parallelizing approach. Our aim was to explore the benefits and potential performance trade-offs of using an abstract programming layer for improved application programming productivity. We used the GNU GCC/LIBGOMP and the KLANG/STARPU compiler and runtime couples, as well as the STARPU runtime used natively through its dedicated API, to conduct an extensive campaign of performance evaluations. We explored using two vastly different numerical setting inputs, both on a 24-core platform and a large 96-core platform. The results first showed the interest

of the OPENMP 4 dependent tasks over the legacy fork-join model, especially in the case of the ellipsoid particle distribution. In contrast to the cube distribution — where parallelism is readily abundant, in a largely unstructured form — the ellipsoid distribution indeed stresses the ability of the underlying runtime system to harvest and exploit complex parallelism. In particular, it demonstrated the high benefit of pinpointed OPENMP task extensions to access additional runtime system features, such as the support for commutative task dependencies, or to a lesser extent the ability to specify task priorities, as made possible by the KLANG/STARPU compiler. On the other hand, GCC+LIBGOMP proved to be very much competitive for datasets with very small tasks. By providing the abstraction of its programming layer, OPENMP makes it possible to get the best of both worlds. In the future, we intend to explore the coupling of OPENMP with task-based runtime systems further, in particular focusing on the support for heterogeneous CPU+accelerator platforms that could be illustrated with a task-based FMM for heterogeneous machines [26].

Note that the support for priorities, compliant with the version 4.5 of the OPENMP standard has been incorporated into GCC 6.1 when finalizing our (long started) experimental campaign. We expect that GCC+LIBGOMP will benefit similarly as KLANG/STARPU of this hint. Indeed, we showed that both suites are extremely robust and comparable concerning performance without priorities, and we expect that it shall be the same with this additional hint. All in all, indeed, the goal of the paper is not to provide a full comparison of the GCC+LIBGOMP versus KLANG/STARPU, but rather to show that, because they are both competitive suites, their use is credible for assessing the discussed programming models. As a consequence, we expect that the provided experimental results will provide a snapshot of the progress made by the OPENMP standard in bridging the gap over fully-feature runtime systems used natively. OPENMP now allows to exploit modern multicore architectures as efficiently, while preserving both runtime independence and ease of programming. It is succeeding in converging towards a well-balanced standard for writing HPC scientific libraries. On that account, we hope this work will motivate the developers of task-based scientific libraries ([27], for instance, in the case of FMM) to reconsider the potential of OPENMP in achieving high-performance while relying on task-based programming. We also believe that the present analysis might be a valuable feed-back to the OPENMP developers and community, and supports the motivation of integrating the proposed extensions. In particular, the *commute* extension proposal has been submitted to the OpenMP ARB, since it can certainly benefit to large classes of applications where many items (cells, particles, objects, ...) receive many unordered contributions from each other.

Nonetheless, it has recently been shown [28] that deeper static preliminary analysis could enhance the performance of FMM. Providing all the information of such a thorough static pre-processing through OPENMP directives still remains a challenge, that certainly needs to be co-addressed in future work by the OPENMP community and the developers of scientific, high-performance libraries.

ACKNOWLEDGMENT

Experiments presented in this paper were carried out using the PLAFRIM experimental testbed, being developed under the Inria PlaFRIM development action with support from Bordeaux INP, LABRI and IMB and other entities: Conseil Régional d'Aquitaine, Université de Bordeaux and CNRS (and ANR in accordance to the programme d'investissements d'Avenir). The KLANG/STARPU compiler has been developed with the contribution of the ADT kStar development action from Inria. The authors thank Mathieu Faverge and Luka Stanisic for reviewing a preliminary version of this manuscript.

REFERENCES

- [1] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick, "The landscape of parallel computing research: A view from Berkeley," EECS, UCB, Tech. Rep., 2006.
- [2] M. Frigo, C. E. Leiserson, and K. H. Randall, "The implementation of the cilk-5 multithreaded language," in *Conference on Programming Language Design and Implementation*, 1998.
- [3] OpenMP Architecture Review Board, "OpenMP application program interface version 4.0," 2013. [Online]. Available: <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>
- [4] E. Ayguadé, N. Copt, A. Duran, J. Hoeftlinger, Y. Lin, F. Massaioli, X. Teruel, P. Unnikrishnan, and G. Zhang, "The design of openmp tasks," *Transactions on Parallel and Distributed Systems*, 2009.
- [5] C. Augonnet, S. Thibault, R. Namyst, and P.-A. Wacrenier, "StarPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures," *Concurrency and Computation: Practice and Experience, Special Issue: Euro-Par 2009*, 2011.
- [6] G. Bosilca, A. Bouteiller, A. Danalis, T. Hérault, P. Lemerminier, and J. Dongarra, "Dague: A generic distributed dag engine for high performance computing," *Parallel Computing*, 2012.
- [7] J. Bueno, X. Martorell, R. M. Badia, E. Ayguadé, and J. Labarta, "Implementing ompss support for regions of data in architectures with multiple address spaces," in *International conference on Supercomputing*, 2013.
- [8] T. Gautier, J. V. F. Lima, N. Maillard, and B. Raffin, "Locality-aware work stealing on multi-cpu and multi-gpu architectures," in *Workshop on Programmability Issues for Heterogeneous Multicores (MULTIPROG)*, 2013.
- [9] Open Source Technology Center, "Open Community Runtime," 2014. [Online]. Available: <https://01.org/open-community-runtime/>
- [10] D. M. Kunzmann and L. V. Kalé, "Programming heterogeneous clusters with accelerators using object-based programming," *Scientific Programming*, 2011.
- [11] E. Slaughter, W. Lee, S. Treichler, M. Bauer, and A. Aiken, "Regent: a high-productivity programming language for hpc with logical regions," in *International Conference for High Performance Computing, Networking, Storage and Analysis*, 2015.
- [12] C. Liao, D. J. Quinlan, T. Panas, and B. R. de Supinski, "A rose-based openmp 3.0 research compiler supporting multiple runtime libraries," in *Beyond Loop Level Parallelism in OpenMP: Accelerators, Tasking and More*. Springer, 2010.
- [13] C. Addison, J. Lagrone, L. Huang, and B. Chapman, "OpenMP 3.0 tasking implementation in OpenUH," in *Open64 Workshop*, 2009.
- [14] R. Ferrer, S. Royuela, D. Caballero, A. Duran, X. Martorell, and E. Ayguadé, "Mercurium: Design decisions for a s2s compiler," *Cetus Users and Compiler Infrastructure Workshop*, 2011.
- [15] X. Teruel, P. Unnikrishnan, X. Martorell, E. Ayguadé, R. Silvera, G. Zhang, and E. Tiotto, "Openmp tasks in IBM XL compilers," in *Conf. of the center for advanced studies on collaborative research*, 2008.
- [16] "LLVM OpenMP Runtime Library," 2015. [Online]. Available: <http://openmp.llvm.org/Reference.pdf/>
- [17] OpenMP Architecture Review Board, "OpenMP application program interface version 4.5," 2015. [Online]. Available: <http://www.openmp.org/mp-documents/openmp-4.5.pdf>
- [18] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *Journal of computational physics*, 1987.
- [19] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner, and T. Takahashi, "Task-based FMM for multicore architectures," *SIAM Journal on Scientific Computing*, 2014.
- [20] "Optimizing and tuning the fast multipole method for state-of-the-art multicore architectures," in *2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)*. IEEE, 2010.
- [21] "Diagnosis, Tuning, and Redesign for Multicore Performance: A Case Study of the Fast Multipole Method," in *Proceedings of the 2010 ACM/IEEE conference on Supercomputing*, 2010.
- [22] R. Yokota and L. A. Barba, "A tuned and scalable fast multipole method as a preeminent algorithm for exascale systems," *International Journal of High Performance Computing Applications*, 2012.
- [23] B. Lange and P. Fortin, "Parallel dual tree traversal on multicore and many-core architectures for astrophysical n-body simulations," UPMC, ICS, LIP6, Tech. Rep., 2014.
- [24] P. Blanchard, O. Coulaud, and E. Darve, "Fast hierarchical algorithms for generating Gaussian random fields," Inria Bordeaux Sud-Ouest, Research Report 8811, Dec. 2015.
- [25] H. Casanova, A. Legrand, and Y. Robert, *Parallel algorithms*. CRC Press, 2008.
- [26] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner, and T. Takahashi, "Task-based FMM for heterogeneous architectures," *Concurrency and Computation: Practice and Experience*, 2016.
- [27] H. Ltaief and R. Yokota, "Data-driven execution of fast multipole methods," *Concurrency and Computation: Practice and Experience*, 2013.
- [28] J. Choi, A. Chandramowlishwaran, K. Madduri, and R. Vuduc, "A cpu: Gpu hybrid implementation and model-driven scheduling of the fast multipole method," in *Proceedings of Workshop on General Purpose Processing Using GPUs*. ACM, 2014.

Emmanuel Agullo holds a permanent researcher position as part of the HiePACS Inria project team in Bordeaux since 2010. His main research interest is the design of scientific, high performance, scientific libraries for modern supercomputers relying on innovative programming paradigms.

Olivier Aumage is permanent researcher as part of the STORM Team, Inria Bordeaux, since 2003. His research topics include runtimes, networks, compilers and HPC tools. He is involved in the development of StarPU and Klang-omp. He represents Inria at OpenMP ARB.

Bérenger Bramas received his MS degree in computer science and software engineering in 2009 from the ISIMA. In 2016, he received his Ph.D. degree in computer science from University of Bordeaux, and worked mainly on ScalFMM and time-domain BEM. His research focuses on HPC scientific applications, software tools for HPC, parallel algorithms, optimizations and heterogeneous computing. He is currently working as an HPC Expert at Max-Planck Computing and Data Facility.

Olivier Coulaud holds a senior researcher position at Inria Bordeaux as part of the HiePACS Inria project team. His main research interest is the design of hierarchical methods (FMM, H-matrices, ...) for material physics simulations.

Samuel Pitoiset holds a full-time engineer position as part of the STORM Team at Inria Bordeaux since 2014. He is part of the H2020 INTERTWinE European project and member of the OpenMP Language Committee since July 2016. His main contribution is the Klang-OpenMP compiler and he is also an active developer of the StarPU runtime.