Taking advantage of hybrid systems for sparse direct solvers via task-based runtimes

Xavier Lacoste, Mathieu Faverge, Pierre Ramet, Samuel Thibault, George Bosilca

Taking advantage of hybrid systems for sparse direct solvers via task-based runtimes

Xavier Lacoste, Mathieu Faverge, Pierre Ramet, Samuel Thibault, George Bosilca

Taking advantage of hybrid systems for sparse direct solvers via task-based runtimes

Xavier Lacoste[*], Mathieu Faverge[*†‡], Pierre Ramet[*‡], Samuel Thibault[*‡], George Bosilca[§]

Project-Team HiePACS

Abstract: The ongoing hardware evolution exhibits an escalation in the number, as well as in the heterogeneity, of the computing resources. The pressure to maintain reasonable levels of performance and portability, forces the application developers to leave the traditional programming paradigms and explore alternative solutions. PaStiX is a parallel sparse direct solver, based on a dynamic scheduler for modern hierarchical architectures. In this paper, we study the replacement of the highly specialized internal scheduler in PaStiX by two generic runtime frameworks: PaRSEC and StarPU. The tasks graph of the factorization step is made available to the two runtimes, providing them with the opportunity to optimize it in order to maximize the algorithm efficiency for a predefined execution environment. A comparative study of the performance of the PaStiX solver with the three schedulers – native PaStiX, StarPU and PaRSEC schedulers – on different execution contexts is performed. The analysis highlights the similarities from a performance point of view between the different execution supports. These results demonstrate that these generic DAG-based runtimes provide a uniform and portable programming interface across heterogeneous environments, and are, therefore, a sustainable solution for hybrid environments.

Key-words: Sparse linear solver, DAG based runtime, multicore, GPU

---

[*] Inria Bordeaux - Sud-Ouest, Talence, France
[†] Institut Polytechnique de Bordeaux, Talence, France
[‡] Laboratoire Bordelais de Recherche en Informatique, Talence, France
[§] Innovative Computing Laboratory – University of Tennessee – Knoxville, Tennessee, USA

Solveurs linéaires creux efficaces sur machines hétérogènes au dessus de supports d'exécution pour graphes de tâches

Résumé : Les architectures de calcul intègrent de plus en plus de coeurs de calcul partageant une même mémoire nécessairement hiérarchique. Les algorithmes, en particulier ceux relatifs à l'algèbre linéaire, nécessitent d'être adaptés à ces nouvelles architectures pour être efficaces. PaStiX est un solveur direct parallèle pour matrices creuses qui intègre un ordonnanceur dynamique pour des architectures hiérarchiques de grande taille. Dans ce papier, nous étudions la possibilité de remplacer cette stratégie interne d'ordonnancement par deux supports d'exécution génériques : PaRSEC et StarPU. Ces supports d'exécution offrent la possibilité de dérouler le graphe de tâches de la factorisation numérique sur des noeuds de calcul disposant d'accélérateurs. Nous présentons une étude comparative des performances de notre solveur supernodal avec ces trois ordonnanceurs sur des architectures multicoeurs, et en particulier les gains obtenus avec plusieurs accélérateurs GPU. Ces résultats montrent qu'une approche basée sur un DAG offre une interface de programmation uniforme pour réaliser du calcul haute performance sur des problèmes irréguliers comme ceux de l'algèbre linéaire creuse.

Mots-clés : Solveur linéaire creux, support d'exécution, graphe de tâches, multicœurs, GPU

1 Introduction

Emerging processor technologies put an emphasis on increasing the number of computing units instead of increasing their working frequencies. As a direct outcome of the physical multiplexing of hardware resources, complex memory hierarchies had to be instated to relax the memory bottleneck and insure a decent rate of memory bandwidth for each resource. The memory become divided in several independent areas, capable of delivering data simultaneously through a complex and hierarchical topology, leading to the mainstream Non Uniform Memory Accesses (NUMA) machines we know today. Together with the availability of hardware accelerators, this trend profoundly altered the execution model of current and future platforms, progressing them toward a scale and a complexity unattained before. Furthermore, with the established integration of accelerators into modern architectures, such as GPUs or Intel Xeon Phis, high-end multi-core CPUs are consistently outperformed by these novel, more integrated, architectures both in terms of data processing rate and memory bandwidth. As a consequence, the working environment of today application developers evolved toward a massively parallel environment, where computation become cheap and data movements expensive, driving up the energetic cost and performance of the algorithms.

With the advent of API for GPUs programming such as CUDA or OpenCL, programming accelerators have been rapidly evolving in the past years, permanently bringing accelerators into the mainstream. Hence, GPUs are becoming a more and more attractive alternative to traditional CPUs, particularly for their more interesting cost-per-flop and watts-per-flop ratios. However, the availability of particular programming API only partially addresses the development of hybrid algorithms capable of taking advantage of all computations resources available, including accelerators and CPUs. Extracting a satisfactory level of performance out of such entangled architectures, remains a real challenge due to the lack of consistent programming models and tools to assess their performance. In order to efficiently exploit current and future architectures, algorithm developers are required to expose a large amount of parallelism, adapt their code to new architectures with different programming models, and finally map it efficiently on the heterogeneous resources. A gargantuan task for most developers as they do not possess the architectural knowledge necessary to mold their algorithms on the hardware capabilities in order to achieve good efficiency, and/or do not want to spend new efforts with every new generation of hardware.

Solving large linear system of equations, $Ax = b$, dense or sparse, is one of the most important and time-consuming part in many scientific and engineering algorithms, a building block en route to more complex scientific application. A bubbling field where scientific advances have a lasting impact on the accuracy and the pace of scientific discoveries. A tremendous amount of work has been done on dense linear algebra side of the field, but sparse linear algebra on heterogeneous system is a work-in-progress area. The PaStiX solver, is a sparse direct solvers that can solve symmetric definite, indefinite, and general problems using Cholesky, LDL^T, and LU factorizations, respectively. The PaStiX implementation relies on a two levels approach using the POSIX Thread library within a node and the Message Passing Interface between different nodes. Historically, PaStiX scheduling strategy is based on a cost model of the tasks executed that defines during the analyze phase the execution order used at runtime. In order to complement the lack of precision of the cost model on hierarchical architectures, a dynamic scheduler based on a work-stealing strategy has been developed to reduce the idle times while preserving a good locality for data mapping [1]. Recently, the solver has been optimized to deal with the new hierarchical multi-cores architectures [2], at the levels of internal data structures of the solver, communication patterns, and scheduling strategies.

In this paper to took upon the challenge of rewriting the PaStiX solver using a task-based

programming paradigm, a novel approach with the promise of delivering the programmers from the hardware constraints. We expose the original algorithm using the concept of tasks, a self-contained computational entity, linked to the other tasks by data dependencies. Specialized task-graph description formats are used in accord with the underlying runtime (PARSEC and STARPU). In addition, we also provided specialized GPU-aware versions for some of the most compute intensive tasks, providing the runtimes with the opportunity to unroll the graph of tasks on all available computing resources. The resulting software is, to the best of our knowledge, the first implementation of a sparse direct solver with supernodal method supporting hybrid execution environments composed of multi-cores and multi-GPU. Based on these elements, we pursue with the evaluation of the usability and the appeal of using a task-based runtime as a substrate for executing this particular type of algorithm, an extremely computationally challenging sparse direct solver. Furthermore, we take advantage of the integration of accelerators (GPUs in this context) with our supporting runtimes, to evaluate and understand the impact of this drastically novel portable way of writing efficient and perennial algorithms. Since the runtime system offers a uniform programming interfaces, dissociated from a specific set of hardware or low-level software entities, applications can take advantage of these uniform programming interfaces in a portable manner. Furthermore, the exposed graph of tasks allows the runtime system to apply specialized optimization techniques to minimize application's time to solution by dynamically mapping the tasks onto computing resources as efficiently as possible.

The rest of the paper is organized as follows. After the related works section, we describe the supernodal method of the PASTIX solver in Section 3, followed by a description of the runtimes used. Section 5 explains the implementation over the DAG schedulers with a preliminary study over multi-core architectures, followed by details on the extension to heterogeneous architectures. All choices are supported and validated by a set of experiments on a set of matrices with a wide range of characteristics. Finally, section 6 concludes with some prospects of the current work.

2   Related Works

The dense linear algebra community spent a great deal of effort to tackle the challenges raised by the sharp increase of the number of computational resources. Due to their heavy computational cost, most of their algorithms are relatively simple to handle. Avoiding common pitfalls such as the "fork-join" parallelism and expertly selecting the blocking factor provide an almost straightforward way to increase the parallelism and thus achieve better performance. Moreover, due to their natural load-balance most of the algorithms can be approached hierarchically, first at the node level, and then at the computational resource level. In a shared memory context, one of the seminal paper [3] replaced the commonly used LAPACK layout with one based on tiles/blocks. Using basic operations on these tiles expose the algorithms as a graph of tasks augmented with dependencies between them. In shared memory, this approach quickly generate a large number of ready tasks, while, in distributed memory, the dependencies allow the removal of hard synchronizations. This idea lead to the design of new algorithms for various algebraic operations [4] now at the base of well known software packages like PLASMA [5].

This idea is recurrent in almost all novel approaches surrounding the many-core revolution, spreading outside the boundaries of dense linear algebra. Looking at the sparse linear algebra, the efforts were directed toward improving the behavior of the existing solvers by taking into account both task and data affinity and relying on a two-level hybrid parallelization approach, mixing multithreading and message passing. Numerous solvers are now able to efficiently exploit the capabilities of these new platforms [2, 6]. New solvers have also been designed and implemented to address these new platforms. For them the chosen approach follows the one for dense linear

algebra, fine-grained parallelism, thread-based parallelization, and advanced data management to deal with complex memory hierarchies. Examples of this kind of solvers are HSL [7] and SuperLU-MT [8] for sparse LU or Cholesky factorizations and SPQR [9] and `qr_mumps` [10] for sparse QR factorizations.

With the advent of accelerator-based platforms, a lot of attention has shifted toward extending the multi-cores aware algorithms to fully exploit the huge potential of accelerators (mostly GPUs). The main challenges raised by these heterogeneous platforms are mostly related to task granularity and data management: although regular cores require fine granularity of data as well as computations, accelerators such as GPUs need coarse-grain tasks. This inevitably introduces the need for identifying the parts of the algorithm which are more suitable to be processed by accelerators. As for the multi-core case described above, the exploitation of this kind of platforms was first considered in the context of dense linear algebra algorithms.

Moreover one constant become clear, a need for a portable layer that will insulate the algorithms and their developers from the rapid hardware changes. Recently, this portability layer appeared under the denomination of task-based runtime. The algorithms are described as tasks with data dependencies in-between, and the runtime systems are used to manage the tasks dynamically and schedule them on all available resources. These runtimes can be generic, like the two runtimes used in the context of this study (STARPU [11] or PARSEC [12]), or more specialized like QUARK [13]. These efforts resulted in the design of the DPLASMA library [14] on top of PARSEC and the adaptation of the existing FLAME library [15]. On the sparse direct methods front, preliminary work has resulted in mono-GPU implementations based on offloading parts of the computations to the GPU [16, 17, 18]. Due to its very good data locality the multifrontal method is the main target of these approaches. The main idea is to treat some parts of the task dependency graph entirely on the GPU. Therefore, the main originality of these efforts is in the methods and algorithms used to decide whether or not a task can be processed on a GPU. In most cases, this was achieved through a threshold based criterion on the size of the computational tasks.

Many initiatives have emerged in previous years to develop efficient runtime systems for modern heterogeneous platforms. Most of these runtime systems use a task-based paradigm to express concurrency and dependencies by employing a task dependency graph to represent the application to be executed. Without going in details, the main differences between these approaches are related to their representation of the graph of tasks, whether they manage data movements between computational resources, the extent they focus on task scheduling and their capabilities to handle distributed platforms.

3 Supernodal factorization

Sparse direct solvers are algorithms that address sparse matrices, mostly filled with zeroes. In order to reduce the number of operations, they consider only non-zeroes of the matrix A. During factorization, new non-zero entries – called fill-in – appear in the factorized matrix and lead to more computation and memory consumption. One of the main objectives of those solvers is to keep the fill-in to its minimum to limit the memory consumption. The first step of any sparse direct solver is the computation of a nested dissection of the problem that result in a permutation of the unknowns of A. This process is a graph partitioning algorithm applied to the connectivity graph associated to the matrix. The computed permutation reduces the fill-in that the factorization process will generate, and the elimination tree [19] is generated out of the separators discovered during the graph partitioning process. Basically, each node of the tree represents the set of unknowns that belongs to a separator and edges are connections between
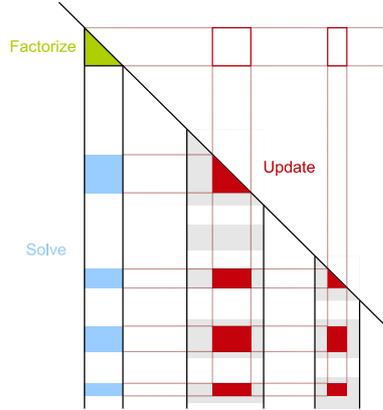
Figure 1: Decomposition of the task applied while processing one panel

those separators in the original graph. The edges of the tree connect a node to almost all nodes
in the path that connect it to the root of the tree. They represent contributions from one node
to another during the factorization. The second step of sparse solvers is the analysis stage which
predicts the non-zero pattern of the factorized matrix through a symbolic factorization. The
resulting pattern is grouped in blocks of non-zero elements to allow for more efficient BLAS
calls. Blocks can be enlarged if extra fill-in is authorized for better performance, or split to
generate more parallelism.

Once the elimination tree and the symbolic factorization are computed, two different methods
can be applied: multifrontal [20] or supernodal [21]. PaStiX solver is a supernodal method.
Each node of the elimination tree, or supernode, represents in the matrix a subset of contiguous
columns, also called panel. At each node of the elimination tree, we associate a task –called 1D
task – that performs three steps associated to the panel A as shown on the figure 1:

1. Factorization of the diagonal block,

2. Triangular solve on the off-diagonal blocks in the panel (TRSM), and

3. For each off-diagonal block A_i, apply the associated update to the facing panel C_i (GEMM)
   – We call facing panel, C_i, the panel with the diagonal block owning the same rows than
   the off-diagonal block A_i.

Figure 1 represents a lower triangular matrix used in case of symmetric problems with
Cholesky factorization, but is also valid for non symmetric cases with PaStiX. In a general
manner, PaStiX works on the matrix $A + A^T$ which produces a symmetric pattern. In non-
symmetric cases, the step 2 and 3 will then be duplicated for the L and U matrices of the LU
factorization. For all the paper, we will discuss only the Cholesky implementation. LDL^T and
LU factorizations follow the same method.

PaStiX solver relies on a cost model of this 1D task to compute a static scheduling. This
static scheduling associates ready tasks to the first available resources among the computational
units. The complexity of such an algorithm depends on the number of tasks and resources.
Hence, the 1D task is kept as a large single task to lower the complexity of the analysis part.
However, it is obvious that more parallelism could be extracted from those tasks, but would
increase the analysis step complexity.

First, the triangular solves, applied on off-diagonal blocks of each panel, are independent
computations that depend only on the diagonal block factorization. Thus, each panel is stored

as a single tall and skinny matrix, such that the TRSM granularity can be decided at runtime and are independent of the data storage. At lower levels of the elimination tree, the small block granularity might induce a large overhead if they are considered as independent tasks. On the contrary, at higher levels, the larger supernodes (Order of $N^{\frac{2}{3}}$ for a 3D problem of N unknowns, or $sqrt(N)$ for a 2D problem) might be split to create more parallelism with low overhead. That is why, supernodes of the higher levels are split vertically prior to the factorization to limit the task granularity and create more parallelism. In this study, to compare to what is existing in PASTIX solver, we keep all TRSM operations on a single panel grouped together as a single operation.

Secondly, the same remark applies to the update tasks with an order of magnitude larger. Each couple of off-diagonal blocks (A_i, A_j), with $i < j$ in a panel, generates an independent update to the trailing submatrix formed by their outer product. To adapt to the small granularity of off-diagonal blocks in sparse solvers, those updates are grouped together. Two variants exists. *Left-looking*: all tasks contributing to a same panel are associated in a single task, they have a lot of input edges and only one in-out data. *Right-looking*: all update generated by a single panel are directly applied to the multiple destination panels. This solution has a single input data, and many panels are accessed as in-out. PASTIX uses the *right-looking* variant. Therefore, the nature of the supernodal algorithm is in itself a DAG of tasks. However, since we consider a data as a panel and generic runtime takes a fixed number of dependencies per tasks, one update task will be generated per couple of panel instead of one per panel, as in PASTIX.

4   Runtimes

In our exploratory approach toward moving to a generic scheduler for PASTIX, we considered two different runtimes, STARPU, and PARSEC. They have been proven mature enough in the context of dense linear algebra, while providing two orthogonal approaches to task-based systems.

The PARSEC [12] distributed runtime system is a generic data-flow engine supporting a task-based implementation targeted toward hybrid systems. Domain specific languages are available to expose a user-friendly interface to developers and allow them to describe their algorithm using high-level concepts. This programming paradigm constructs an abridged representation of the tasks and their dependencies as a graph of tasks – a structure agnostic to algorithmic subtleties, where all intrinsic knowledge about the complexity of the underlying algorithm is extricated, and the only constraints remaining are annotated dependencies between tasks [22]. This symbolic representation augmented with a specific data distribution is then mapped on a particular execution environment. The runtime supports the usage of different types of accelerators, GPUs and Intel Xeon Phis, in addition to distributed multi-cores processors. Data are transferred between computational resources based on coherence protocols and computational needs, with emphasis on minimizing the unnecessary transfers. The resulting tasks are dynamically scheduled on the available resources following a data reuse policy mixed with different criteria for adaptive scheduling. The entire runtime targets very fine grain tasks (order of magnitude under ten microseconds), with a flexible scheduling and adaptive policies to mitigate the effect of system noise and take advantage of the algorithmic inherent parallelism to minimize the execution span.

The experiment presented in this paper takes advantage of a specialized domain specific language and interface for PARSEC, designed for linear algebra [14]. This specialized interface allows for a drastic reduction in the memory used by the runtime, as tasks do not exists until they are ready to be executed, and the concise representation of the task-graph allow for an easy and stateless exploration of the graph. The need for a window of visible tasks is then pointless, the runtime can explore the graph dynamically based on the ongoing state of the execution.

STARPU [11] is a runtime system aiming to allow programmers to exploit the computing power of clusters of hybrid systems composed of CPUs and various accelerators (GPUs, Intel Xeon Phis, . . .) while relieving them from the need to specially adapt their programs to the target machine and processing units. The STARPU runtime supports a *task-based programming model*, where applications submit computational tasks, with CPU and/or accelerator implementations, and STARPU schedules these tasks and associated data transfers on available CPUs and accelerators. The data that a task manipulates is automatically transferred among accelerators and the main memory in an optimized way (minimized data transfers, data prefetch, communication overlapped with computations, etc.), so that programmers are relieved from the scheduling issues and technical details associated with these transfers. STARPU takes particular care of scheduling tasks efficiently, by establishing performance models of the tasks through on-line measurements, and then using well-known algorithms from the literature. In addition, it allows scheduling experts, such as compilers or computational library developers, to implement custom scheduling policies in a portable fashion.

The differences between the two runtimes can be classified into two groups: conceptual and practical differences. At the conceptual level the main differences between PARSEC and STARPU are the tasks submission process, the centralized scheduling and the data movement strategy. PARSEC uses its own parameterized language to describe the DAG in comparison with the simple sequential submission loops typically used with STARPU. Therefore, STARPU relies on a centralized strategy that analyzes at runtime the dependencies between tasks and schedules these tasks on the available resources. On the contrary, through compile-time information, each computational unit of PARSEC immediately releases the dependencies of the completed task solely using the local knowledge of the DAG. At last, while PARSEC uses an opportunistic approach, STARPU scheduling strategy exploits cost models of the computation and data movements to schedule tasks to the right resource (CPU or GPU) in order to minimize overall execution time. However, it does not have data-reuse policy on CPU-shared memory systems, resulting in lower efficiency when no GPUs are used compared to the data-reuse heuristic of PARSEC. At the practical level, PARSEC support multiple streams to manage the CUDA devices allowing partial overlap between computing tasks maximizing the occupancy of the GPU. On the other hand, STARPU allows data transfers directly between GPUs without going through central memory, potentially increasing the bandwidth of data transfers when data is needed by multiple GPUs.

5 Supernodal factorization over DAG schedulers

Similarly to dense linear algebra, sparse direct factorization relies on three types of operations: the factorization of the diagonal block (POTRF), the solve on off-diagonal blocks belonging to the same panel (TRSM) and the trailing panels updates (GEMMs). Whereas the task dependency graph from a dense Cholesky factorization [4] is extremely regular, the DAG describing the supernodal method contains rather small tasks with variable granularity and less uniform ranges of execution space. This lack of uniformity makes the DAG resulting from a sparse supernodel factorization complex, accruing the difficulty to efficiently schedule the resulting tasks on homogeneous and heterogeneous computing resources.

The current scheduling scheme of PASTIX exploits a 1D-block distribution, where a task assemble a set of operations together, including the tasks factorizing one panel (POTRF and TRSM) and all updates generated by this factorization. However, increasing the granularity of a task in such a way limits the potential parallelism, and has a growing potential of bounding the efficiency of the algorithm when using many-cores architectures. To improve the efficiency
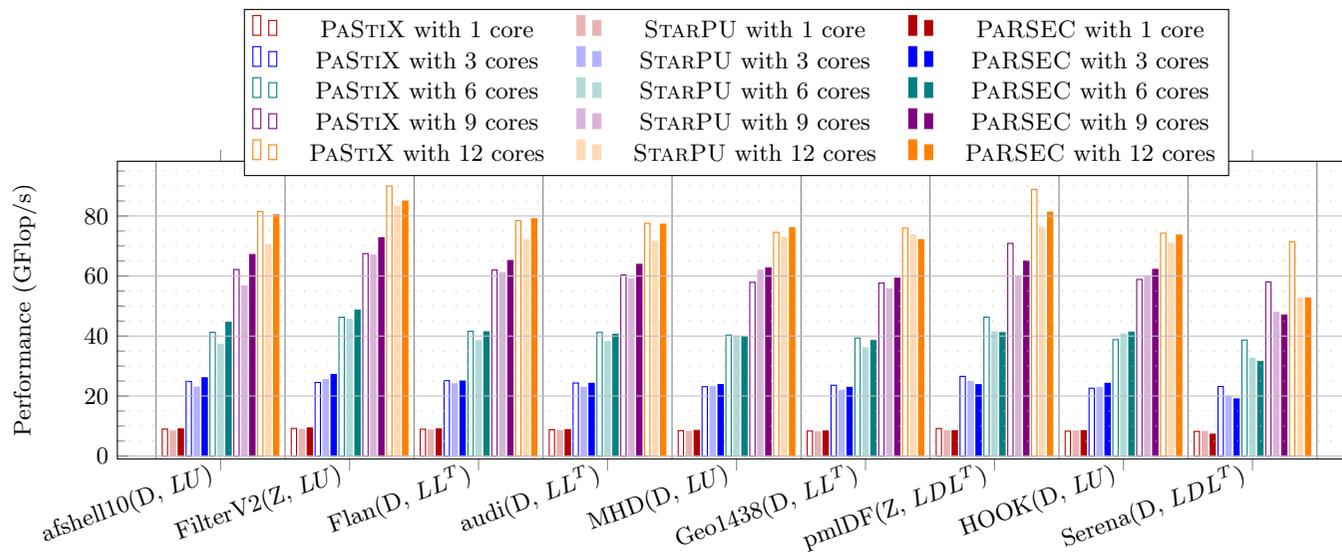
Figure 2: CPU scaling study: Flop/s reached during factorization of the different matrices, with the three schedulers.

of the sparse factorization on a multi-core implementation, we introduced a way of controlling the granularity of the BLAS operations (called ESP option for Enhanced Sparse Parallelism). This functionality dynamically splits a single task of computing the contribution to the trailing submatrix, using the current panel into subtasks, so that the critical path of the algorithm can be reduced. In this paper, for both PaRSEC and StarPU runtimes, we split PaStiX tasks into two sub-sets of tasks:

- the diagonal block factorization and off-diagonal blocks updates, performed on one panel;

- the updates from off-diagonal blocks of the panel to one other panel of the trailing submatrix.

Hence, the number of tasks is bound by the number of blocks in the symbolic structure of the factorized matrix.

Moreover, when taking in account for heterogeneous architectures in the experiments, a finer control of the granularity of the computational tasks is needed. Some references for benchmarking dense linear algebra kernels are described in [23] and show that efficiency could be obtained on GPU devices only on relatively large blocks – a limited number of such blocks can be found on a supernodal factorization only on top of the elimination tree. Similarly, the amalgamation algorithm [24], reused from the implementation of an incomplete factorization, is a crucial step to obtain larger supernodes and efficiency on GPU devices. The default parameter for amalgamation has been slightly increased to allow up to 12% more fill-in to build larger blocks while maintaining a decent level of parallelism.

In the remaining of the paper, we present the extensions to the solver to support heterogeneous many-cores architectures. These extensions were validated through experiments conducted on *Mirage* nodes from PLaFRIM cluster at INRIA Bordeaux. A *Mirage* node is equipped with two hexa-cores Westmere Xeon X5650 (2.67 GHz), 32 Go of memory and 3 Tesla M2070 GPU. PaStiX was built without `MPI` support using `GCC 4.6.3`, CUDA 4.2, `Intel MKL 10.2.7.041` and Scotch `5.1.12b`. Experiments were performed on a set of nine matrices, all part of the University of Florida sparse matrix collection [25], and are described in the table 1. These matrices are representative of different research fields and exhibit a wide range of properties (size, arithmetic, symmetry, definite problem, . . .).

5.1   Multicores Architectures

As mentioned earlier, the PaStiX solver has already been optimized for distributed clusters of NUMA nodes. We use the current state-of-the-art PaStiX scheduler as a basis, and compare the results obtained using the StarPU and PaRSEC runtimes from there. The Figure 2 reports the results from a strong scaling experiment, where the number of computing resources varies from 1 to 12 cores, and where each group represents a particular matrix. Empty bars correspond to PaStiX original scheduler, shaded bars to StarPU, and filled bars to PaRSEC. The figure is in Flop/s, a higher value on the Y-axis represents a more efficient implementation. Overall, this experiment shows that on a shared memory architecture the performance obtained with any of the above-mentioned approaches are comparable, the differences remaining minimal on the target architecture.

We can also notice that, in most cases, the PaRSEC implementation is more efficient than StarPU, especially when the number of cores increases. StarPU shows an overhead on multicore experiments attributed to its lack of cache reuse policy compare to PaRSEC and PaStiX internal scheduler. A careful observation highlight the fact that both runtimes obtain lower performance compared with PaStiX for LDL^T on both PmlDF and Serena matrices. Due to

Matrix	Prec	Method	Size	nnz_A	nnz_L	TFlop
Afshell10	D	LU	1.5e+6	27e+6	610e+6	0.12
FilterV2	Z	LU	0.6e+6	12e+6	536e+6	3.6
Flan	D	LL^T	1.6e+6	59e+6	1712e+6	5.3
Audi	D	LL^T	0.9e+6	39e+6	1325e+6	6.5
MHD	D	LU	0.5e+6	24e+6	1133e+6	6.6
Geo1438	D	LL^T	1.4e+6	32e+6	2768e+6	23
Pmldf	Z	LDL^T	1.0e+6	8e+6	1105e+6	28
Hook	D	LU	1.5e+6	31e+6	4168e+6	35
Serena	D	LDL^T	1.4e+6	32e+6	3365e+6	47

Table 1: Matrix description (Z: double complex, D: double).

its single task per node scheme, PaStiX stores in a temporary buffer the DL^T matrix which allows the update kernels to call a simple GEMM operation. On the contrary, both StarPU and PaRSEC implementations are using an less efficient kernel, that performs at each update the full operation LDL^T. Indeed, due to the extended set of tasks, the life span of the temporary buffer could cause large memory overhead. In conclusion, using these generic runtimes shows similar performances and scalability to the PaStiX internal solution on the majority of test cases while providing a suitable level of performance and a desirable portability allowing for a smooth transition toward more complex heterogeneous architectures.

5.2 Heterogeneous Architectures Implementation

While obtaining an efficient implementation was one of the goals of this experiment, it was not the major one. The ultimate goal was to develop a portable software environment allowing for a even transition to accelerators, a software platform where the code is factorized as much as possible, and where the human cost of adapting the sparse solver to current and future, hierarchical complex heterogeneous architectures remains consistently low. Building upon the efficient supernodal implementation on top of DAG based runtimes, we can exploit more easily heterogeneous architectures. The GEMM updates are the most compute intensive part of the matrix factorization, it is natural that these tasks are important to be offloaded to the GPU. We decide not to offload the tasks that factorize and update the panel to the GPU due to the limited computational load, in direct relationship with the small width of the panels. It is a common in dense linear algebra to use the accelerators for the update part of a factorization while the CPUs factorize the panel, so from this perspective our approach is conventional. However, such an approach combined with look ahead techniques, gives really good performance for a low programming effort on the accelerators [26]. The same solution is applied in this study, since the panels are split during analysis step to fit the classic look-ahead parameters.

It is a known fact that the update is the most compute intensive task during a factorization. Therefore, generally speaking, it is paramount to obtain good efficiency on the update operation in order to ensure a reasonable level of performance for the entire factorization. Due to the embarrassingly parallel architecture of the GPUs and to the extra cost of moving the data back and forth between the main memory and the GPUs, it is of greatest importance to maintaing this property also on the GPU.

As presented in Figure 1, the update task used in PaStiX solver groups together all outer products that are applied to a same panel. On CPU side, this GEMM operation is split in two steps due to the gaps in the destination panel: the outer product is computed in a contiguous temporary buffer, and upon completion, the result is dispatched on the destination panel. This
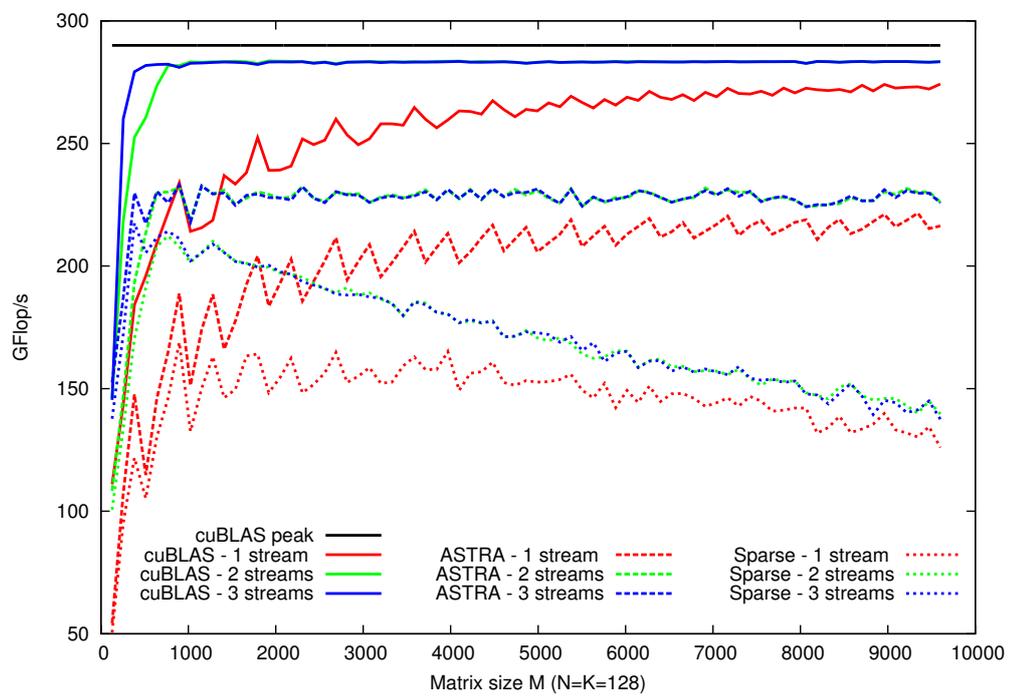
Figure 3: Multi-stream performance comparison on the DGEMM kernel for three implementations: cuBLAS library, ASTRA framework and the sparse adaptation of the ASTRA framework.

solution has been chosen to exploits the performances of vendor's provided BLAS libraries in exchange of a constant memory overhead per working thread.

For the GPU implementation, the requirements for an efficient kernel are different. First, a GPU has significantly less memory compared with what is available to a traditional processor, usually in the range of 3 to 6 GB. This force us to carefully restrict the amount of extra memory needed during the update, making the temporary buffer used in the CPU version unsuitable. Second, the uneven nature of sparse irregular matrices might limit the number of active computing units per task. As a result only a partial number of the available warps on the GPU might be active, leading to a deficient occupancy. Thus, we need the capability to submit multiple concurrent updates in order to provide the GPU driver with the opportunity to overlap warps between different tasks to increase the occupancy, and thus the overall efficiency.

Many CUDA implementations of the dense GEMM kernel are available to the scientific community. The most widespread implementation is provided by Nvidia itself in the cuBLAS library [27]. This implementation is extremely efficient and since CUDA 4.2 allows for calls on multiple streams but is not open source. Volkov developed an implementation for the first generation of CUDA enabled devices [23] in real single precision. In [28], authors propose a assembly code of the DGEMM kernel that get 20% improvement on cuBLAS 3.2 implementation. The MAGMA library implemented proposed a first implementation of the DGEMM kernel [29] for the second generation of Nvidia GPU, known as Fermi. Later, an auto-tuned framework, called ASTRA, has been presented in [30] and included to the MAGMA library. This implementation, similar to the ATLAS library for CPU, is an highly configurable skeleton with a set of scripts to tune the parameters for each precision.

As our update operation is applied on a sparse representation of the panel and matrices, we cannot exploit the a efficient vendor-provided GEMM kernel. We need to develop our own, starting from a dense version and altering the algorithm to fit our needs. Due to the source code availability, the coverage of the four floating point precisions and it's tuning capabilities, we decided to use the ASTRA-based version for our *sparse* implementation. As explained in [30] the matrix-matrix operation is performed in two steps in this kernel. Each block of threads computes the outer-product $tmp = AB$ into the GPU shared memory, and then the addition $C = \beta C + \alpha tmp$ is computed. To be able to compute directly into C the result of the update from one panel to another, we extended the kernel to provide the structure of each panel. This allows the kernel to compute directly the correct position into C during the *sum* step. This introduce a loss in the memory coalescence and deteriorates the update parts, however it prevents the requirement of an extra buffer on the GPU for each offloaded kernel.

One problem in the best parameters used in the MAGMA library for the ASTRA kernel is that it has been determined that using textures gives the best performance for the update kernel. The function `cudaBindTexture` and `cudaUnbindTexture` are not compatible with concurrent kernel calls on different streams. Therefore, the textures have been disabled in the kernel reducing by about 5% the performance of the kernel on large square matrices.

Figure 3 shows the study we made on the GEMM kernel and the impact of the modifications we did on the ASTRA kernel. This experiments is done on a single GPU of *Mirage* cluster. The experiments consists in computing a representative matrix-matrix multiplication of what is typically encountered during sparse factorization. Each point is the average performance of 100 calls to the kernel that computes: $C = C - AB^T$ with A, B and respectively C matrices of dimension M-by-N, K-by-N and M-by-N. B is taken as the first block of K rows of A as it is the case in Cholesky factorization. The plain lines are the performance of the cuBLAS library with 1 stream (*red*), 2 streams (*green*) and 3 streams (*red*). The black line represents the peak performance obtained by the cuBLAS library on square matrices. This peak is never reached with the particular configuration case studied here. The dashed lines are the performance of

the ASTRA library in the same configurations. We observe that this implementation looses already 50GFlop/s, around 15%, against the cuBLAS library and that might be caused by the parameters chosen by the auto-tuning framework that has been run only on square matrices. Finally, the dotted lines illustrate the performance of the modified ASTRA kernel to includes the gaps into the C matrix. For the experiments, C is a panel twice as taller as A in which blocks are randomly generated with size of average 200 rows. Blocks in A are also randomly generated with the constraint that the rows interval of a block of A is included in the rows interval of one block of C, and no overlap is made between two blocks of A. We observe a direct relationship between the height of the panel and the performance of the kernel: the taller the panel is the lower is the performance of the kernel. The memory loaded to do the outer product is still the same than for the ASTRA curves, but memory loaded for the C matrix growth twice faster without increasing the number of Flop to perform. The ratio Flop per memory access is dropping and explain the decreasing performance. However, when the factorization progresses and moves up the elimination trees, nodes get larger and the real number of block encountered is smaller than the one used in this experiment to illustrate worst cases.

Without regards to the kernel choice, it is important to notice how the multiple streams can have a large impact on the average performance of the kernel. For this comparison, the 100 calls made in the experiments are distributed in a round-robin manner over the streams available. One stream gives always the worst performance. Adding a second stream gives increase the performance of all implementations and especially an important one for small cases when matrices are too small to feed all resources of the GPU. The third one is an improvement for matrices with M smaller than 1000 and is similar to two streams over 1000.

This kernel is the one we provide to both runtimes to offload computations on GPUs in case of Cholesky and LU factorizations. An extension of the kernel has been made to handle the LDL^T factorization that takes an extra parameter to the diagonal matrix D and computes: $C = C - LDL^T$. This modified version decreases the performance by 5%.

5.3   Data mapping over multiple GPUs

We noticed that both runtimes encountered difficulties to compute the GEMM mapping onto the GPUs. Tasks' irregularities – size of the GEMM block, Flop in the operation – complicates the prediction model calibration of STARPU to determines which process unit fit best for a task. Dynamic tasks mapping in PARSEC could also result in large unbalanced workload. Indeed in PARSEC, the mapping of a panel to a given GPU enforces the mapping of all other tasks applied to the same panel, for the same reason as in STARPU, the irregularities in the task made it difficult to the simple model used in dense to correctly detect the actual load of each GPU. In order to help the runtime, panels that will be updated on GPU are selected at analyze time, such that the amount of data updated on a GPU does not exceed the memory of the GPU. This limit will reduce the data transfer by keeping data on the GPUs. It is necessary to sort the panels according to a criterion to decide in which ones will be mapped on a GPU. Several sorting criteria were tested: *target panel's size*, larger is panel more chances it has to receive uploads; *number of updates performed on the panel*: this correspond to the number of GEMM applied to the panel; *Flop produced by these updates*: not only the number of update is important, but larger are those updates, the more Flop will be performed; and, finally *the priority statically computed in* PASTIX: higher is the priority, sooner the result is required, such that accelerators can help providing them rapidly. Panels are marked to be updated on a specific GPU thanks to a greedy algorithm that associates at each step of the algorithm the first panel according to the selected criteria to the less loaded GPU. A check is made to guaranty we do not excess the memory capacity of the GPU to avoid excess use of the runtime LRU.
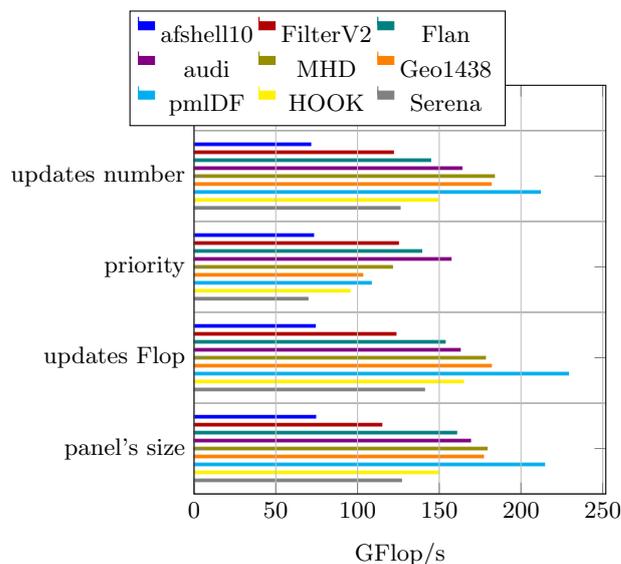
Figure 4: Sort criteria study: Flop/s obtained with PARSEC, on 3 GPUs, with 3 streams.

Figure 4 compares Flop/s obtained with the different criteria on our set of matrices. Each color correspond to a matrix, and the bars are grouped by criterion. The best performance were obtained using the Flop criterion. In future experiments, we report only results with this criterion.

5.4   Heterogeneous experiments

Figure 5 present the performance obtained on our set of matrices on the *Mirage* platform by enabling the GPUs in addition to all available cores. PASTIX run is shown as a reference. STARPU runs are empty bars, PARSEC with 1 streams shaded and PARSEC with 3 streams fully colored. This experiment shows that we can efficiently use the additional computational power provided by the GPUs using the generic runtimes. In its current implementation, STARPU has either GPU or CPU worker threads. A GPU worker will execute only GPU tasks. Hence, when a GPU is used, a CPU worker is removed. With PARSEC, no thread is dedicated to a GPU and they all might executes CPU tasks as well as GPU. The first computational threads that submit a GPU tasks, take the management of the GPU until no GPU work remains in the pipeline. Both runtimes manage to get similar performance and satisfying scalability over the 3 GPUs. In only two cases, MHD and pmlDF, STARPU outperforms PARSEC results with 3 streams. This experimentation also reveals that, as it was expected, the computation takes advantage of the multiple streams that are available through PARSEC. Indeed, the tasks generated by a sparse factorization are rather small and won't use the entire GPU. This PARSEC feature compensates the prefetch strategy of STARPU that gave it the advantage when compare to the one stream results. One can notice the poor performance obtained on the `afshell` test case: in this case, the amount of Flop produced is too small to efficiently benefit from the GPUs.
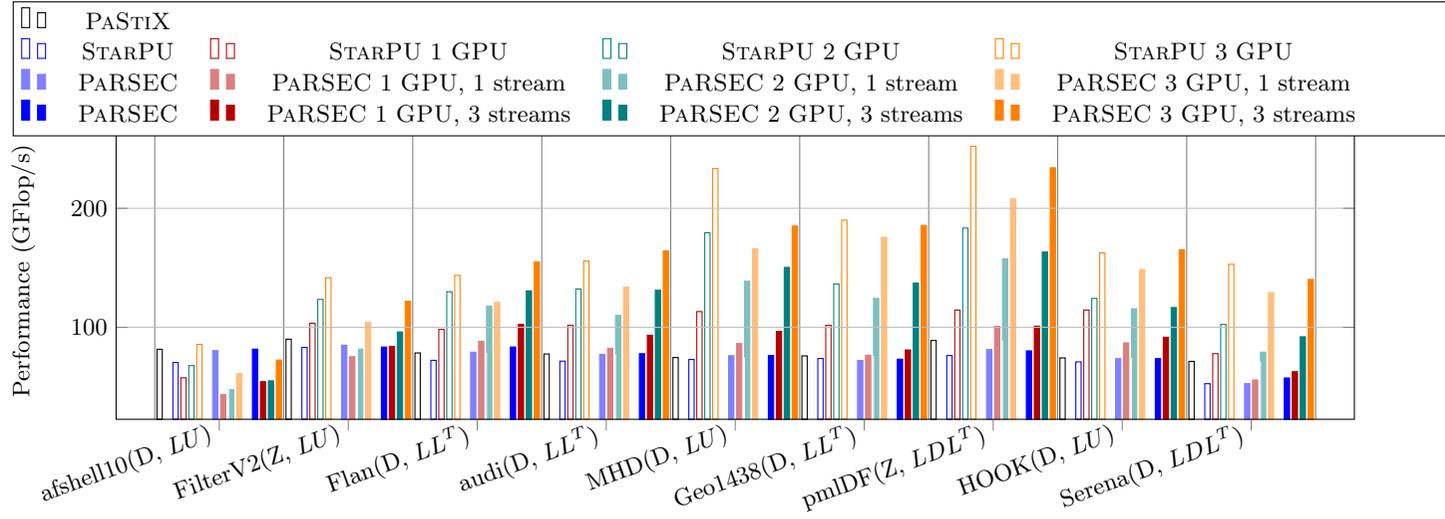
Figure 5: GPU scaling study: Flop/s reached during factorization of the different matrices, with the three schedulers, using the maximum number of cores and zero to three GPUs.

5.5 Memory study

Memory is a critical ressource for direct solver. The figure 6 compares the memory peaks obtained with the three implementations of the solver, using `eztrace`. The runs where obtained with 12 cores but the results would not be much different with an other setup.

The memory allocated cna be separated in five categories :

- the coefficients of the factorized matrix which are allocated at the begining of the compu-ataion and is the largest part of the memory;

- the structure of the factorized matrix also built and allocated before the factorization is performed;

- user's CSC[1] which is the imput given to PaStiX;

- internale block distributed `CSC` which correspond to the input matrix and is used to compute relative error or to perform an iterative refinement;

- a last part of memory that includes the scheduler overhead.

As shown in the plot, a large part of memory correspond to the first four categories and is independant from the runtime used.

The last part of the bars correspond to the overhead of the scheduler.

The values on top of the bars are the overhead ratio compared to memory overhead obtained with PaStiX native scheduler. We can see that we obtained a small overhead with ParSEC whereas StarPU allocate about 7% more memory than PaStiX.

Figure 7 shows the number of calls to `malloc()` function for each test cases. The memory overhead noticed earlier with StarPU may be due to an impressive number of small allocations. We have to investigate why StarPU require such a large number of small allocation.

Table 2: Memory consumption

MAT	PaStiX	StarPU	$\frac{StarPU}{PaStiX}$	ParSEC	$\frac{PaRSEC}{PaStiX}$
afshell10	$1.12 \cdot 10^{10}$	$1.2 \cdot 10^{10}$	1.08	$1.18 \cdot 10^{10}$	1.05
FilterV2	$9.88 \cdot 10^{9}$	$1.08 \cdot 10^{10}$	1.09	$1.02 \cdot 10^{10}$	1.03
Flan1565	$1.69 \cdot 10^{10}$	$1.8 \cdot 10^{10}$	1.06	$1.74 \cdot 10^{10}$	1.03
audi	$1.27 \cdot 10^{10}$	$1.38 \cdot 10^{10}$	1.08	$1.29 \cdot 10^{10}$	1.01
MHD	$1.38 \cdot 10^{10}$	$1.47 \cdot 10^{10}$	1.06	$1.37 \cdot 10^{10}$	0.99
Geo1438	$2.43 \cdot 10^{10}$	$2.58 \cdot 10^{10}$	1.07	$2.45 \cdot 10^{10}$	1.01
pmlDF	$1.88 \cdot 10^{10}$	$2.15 \cdot 10^{10}$	1.14	$1.92 \cdot 10^{10}$	1.02
Hook	$3.09 \cdot 10^{10}$	$3.27 \cdot 10^{10}$	1.06	$3.12 \cdot 10^{10}$	1.01
Serena	$2.59 \cdot 10^{10}$	$2.79 \cdot 10^{10}$	1.08	$2.61 \cdot 10^{10}$	1.01

Table 3: Memory consumption

MAT	coefficients	Factorization matrix structure	Symbolic factorization matrix structure
afshell10	$9.5 \cdot 10^{9}$	$1.12 \cdot 10^{7}$	$2.72 \cdot 10^{6}$
FilterV2	$8.46 \cdot 10^{9}$	$4.87 \cdot 10^{7}$	$8.35 \cdot 10^{6}$

---

[1]Compress Sparse Column

Table 3 – continued from previous page

Flan1565	$1.35 \cdot 10^{10}$	$5.25 \cdot 10^7$	$8.32 \cdot 10^6$
audi	$1.03 \cdot 10^{10}$	$8.29 \cdot 10^7$	$8.46 \cdot 10^6$
MHD	$1.23 \cdot 10^{10}$	$5.14 \cdot 10^7$	$4.46 \cdot 10^6$
Geo1438	$2.18 \cdot 10^{10}$	$2.11 \cdot 10^8$	$1.27 \cdot 10^7$
pmlDF	$1.71 \cdot 10^{10}$	$2.84 \cdot 10^8$	$3.07 \cdot 10^7$
Hook	$2.87 \cdot 10^{10}$	$1.43 \cdot 10^8$	$1.45 \cdot 10^7$
Serena	$2.33 \cdot 10^{10}$	$3.21 \cdot 10^8$	$1.34 \cdot 10^7$

Table 4: Memory consumption

MAT	User's CSC before check	CSCua	User's CSC after check
afshell10	$4.46 \cdot 10^8$	$4.46 \cdot 10^8$	$8.55 \cdot 10^8$
FilterV2	$3.16 \cdot 10^8$	$3.16 \cdot 10^8$	$6.12 \cdot 10^8$
Flan1565	$9.65 \cdot 10^8$	$9.65 \cdot 10^8$	$1.89 \cdot 10^9$
audi	$6.36 \cdot 10^8$	$6.36 \cdot 10^8$	$1.25 \cdot 10^9$
MHD	$3.91 \cdot 10^8$	$3.91 \cdot 10^8$	$5.85 \cdot 10^8$
Geo1438	$5.28 \cdot 10^8$	$5.28 \cdot 10^8$	$1.02 \cdot 10^9$
pmlDF	$2.08 \cdot 10^8$	$2.08 \cdot 10^8$	$3.82 \cdot 10^8$
Hook	$5.12 \cdot 10^8$	$5.12 \cdot 10^8$	$9.87 \cdot 10^8$
Serena	$5.39 \cdot 10^8$	$5.39 \cdot 10^8$	$1.04 \cdot 10^9$

6   Conclusion

In this paper, we have presented a new implementation of a sparse direct solver with supernodal method using a task-based programming paradigm. The programming paradigm shift insulates the solver from the underlying hardware. The runtime takes advantage of the parallelism exposed via the graph of tasks to maximize the efficiency on a particular platform, without the involvement of the application developer. In addition to the alteration of the mathematical algorithm to adapt the solver to the task-based concept, and to provide an efficient memory-contraint sparse GEMM for the GPU, contributions to both runtimes (PARSEC and STARPU) were made such that they could efficiently support tasks with irregular durations, and minimize the non-regular data movements to, and from, the devices. While the current status of this development is already significant in itself, the existence of the conceptual task-based algorithm opened astonishing new perspective toward the seamless integration of any types of accelerators. Providing computational kernels adapted to specialized architectures became now the only obstruction to a portable, efficient, and generic sparse direct solver exploiting these devices. In the context of this study, developing efficient and specialized kernels for GPUs allowed a swift integration on hybrid platforms. Globally, our experimental results corroborate the fact that the portability and efficiency of the proposed approach are indeed available, elevating this approach to a suitable programming model for application on hybrid environments.

Future work will concentrate on smoothing the runtime integration within the solver. First, in order to minimize the scheduler overhead, we plan to increase the granularity of the tasks at the bottom of the elimination tree. Merging leaves or subtrees together constructs bigger, more computationally intensive tasks. Second, we will pursue the extension of this work in distributed heterogeneous environments. On such platforms, when a supernode update another non-local
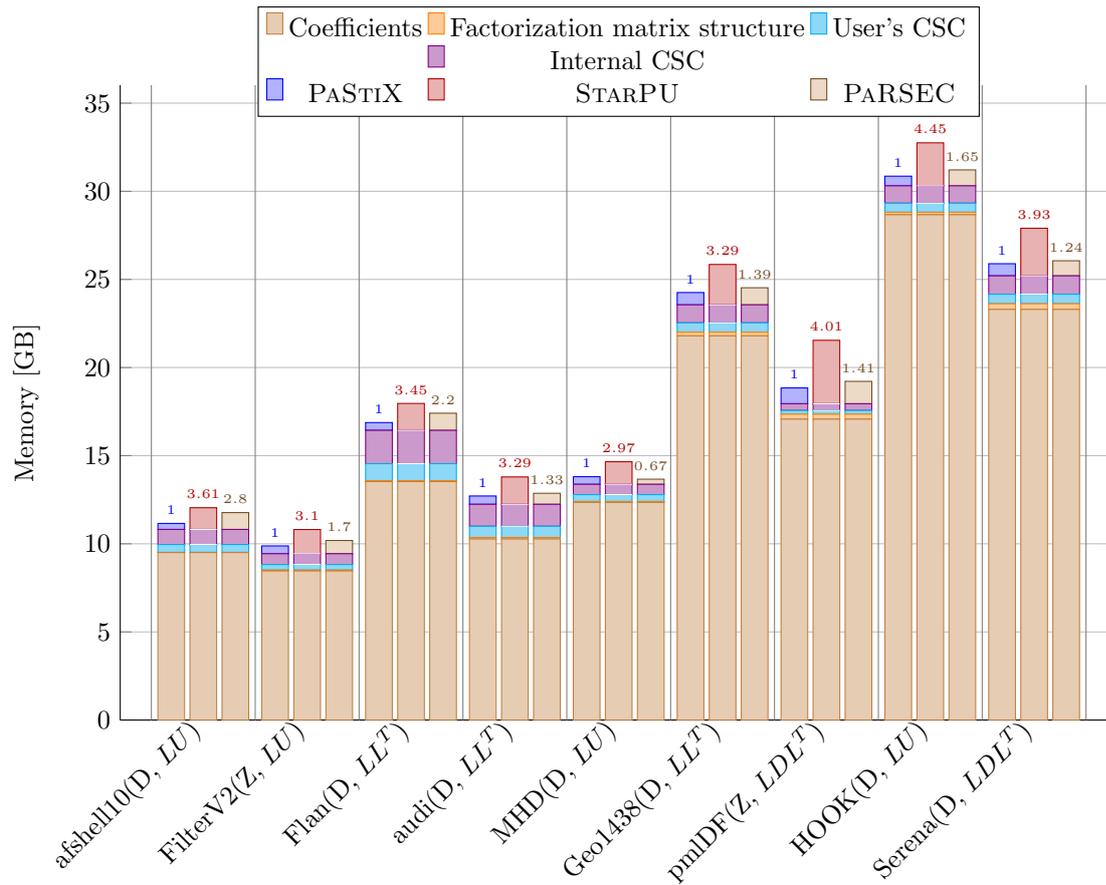
Figure 6: Memory consumption, common data structures are detailed, overhead ratio on top of bar chart
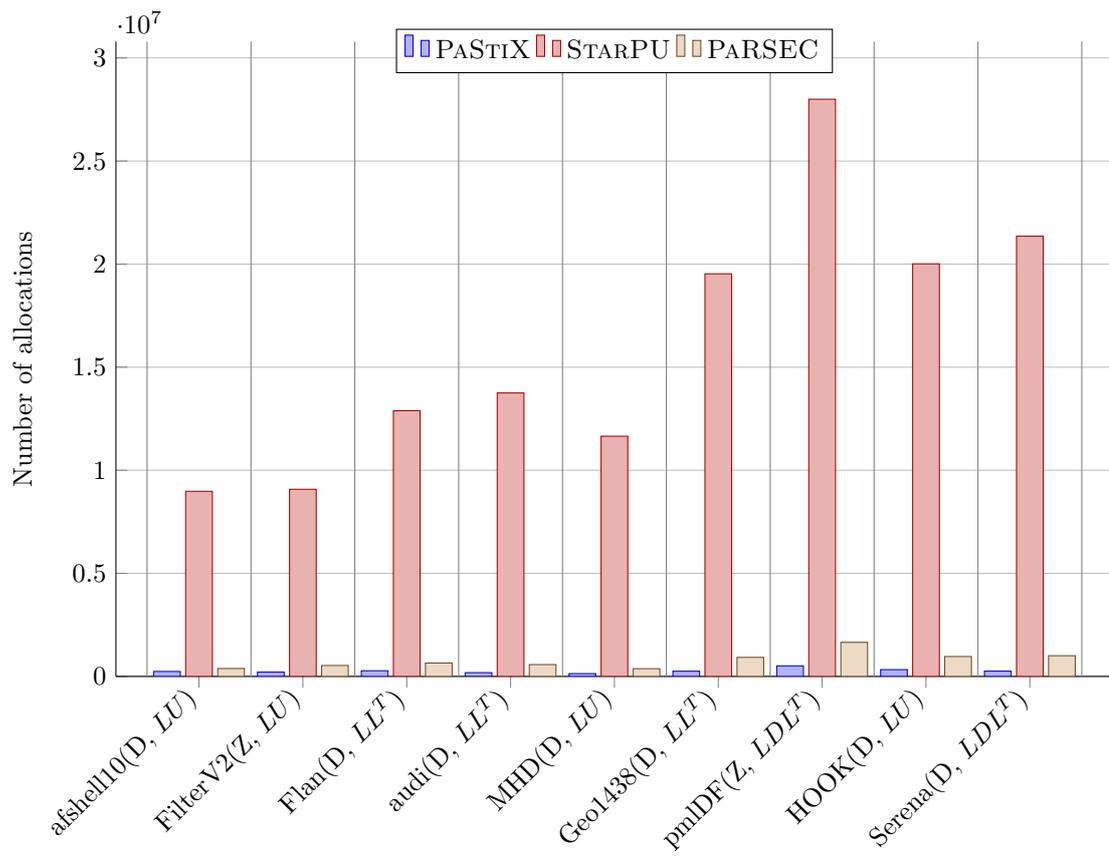
Figure 7: Number of allocation during run

supernode, the update blocks are stored in a local extra-memory space (this is called "fan-in" approach). By locally accumulating the updates until the last updates to the supernode is available, we trade bandwidth for latency. The runtime will allow to study dynamic algorithms, where the number of local accumulation has bounds discovered at runtime. Finally, the availability of extra computational resources, highlight the potential to dynamically build or rebuild the supernodal structures according to the load on the cores and the GPUs. Simultaneously, we will work on the challenging problem of refining the initial mapping of the data based on the heterogeneous capabilities of the distributed memory architectures and to dynamically rebalance the data to match the computational power of the available resources.

Acknowledgment

References

[1] M. Faverge and P. Ramet, "Fine grain scheduling for sparse solver on manycore architectures," in *15th SIAM Conference on Parallel Processing for Scientific Computing*, Savannah, USA, Feb. 2012.

[2] ——, "Dynamic Scheduling for sparse direct Solver on NUMA architectures," in *PARA'08*, ser. LNCS, Norway, 2008.

[3] A. Buttari, J. Dongarra, J. Kurzak, J. Langou, P. Luszczek, and S. Tomov, "The impact of multicore on math software," *PARA 2006*, 2006.

[4] A. Buttari, J. Langou, J. Kurzak, and J. Dongarra, "A class of parallel tiled linear algebra algorithms for multicore architectures," *Parallel Computing*, vol. 35, no. 1, pp. 38–53, 2009.

[5] E. Agullo, J. Demmel, J. Dongarra, B. Hadri, J. Kurzak, J. Langou, H. Ltaief, P. Luszczek, and S. Tomov, "Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects," vol. Vol. 180, no. 1, 2009, p. 012037.

[6] O. Schenk and K. Gärtner, "Solving unsymmetric sparse systems of linear equations with pardiso," *Future Gener. Comput. Syst.*, vol. 20, no. 3, pp. 475–487, 2004.

[7] J. D. Hogg, J. K. Reid, and J. A. Scott, "Design of a multicore sparse cholesky factorization using dags," *SIAM Journal on Scientific Computing*, vol. 32, no. 6, pp. 3627–3649, 2010.

[8] X. S. Li, "Evaluation of sparse LU factorization and triangular solution on multicore platforms," in *VECPAR*, ser. Lecture Notes in Computer Science, J. M. L. M. Palma, P. Amestoy, M. J. Daydé, M. Mattoso, and J. C. Lopes, Eds., vol. 5336. Springer, 2008, pp. 287–300.

[9] T. A. Davis, "Algorithm 915, SuiteSparseQR: Multifrontal multithreaded rank-revealing sparse QR factorization," *ACM Trans. Math. Softw.*, vol. 38, no. 1, p. 8, 2011.

[10] A. Buttari, "Fine-grained multithreading for the multifrontal QR factorization of sparse matrices," 2013, to appear in SIAM SISC and APO technical report number RT-APO-11-6.

[11] 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*, vol. 23, no. 2, 2011.

[12] G. Bosilca, A. Bouteiller, A. Danalis, T. Hérault, P. Lemarinier, and J. Dongarra, "DAGuE: A generic distributed DAG engine for High Performance Computing," *Parallel Computing*, vol. 38, no. 1-2, 2012.

[13] A. YarKhan, "Dynamic task execution on shared and distributed memory architectures," Ph.D. dissertation, dec 2012.

[14] G. Bosilca, A. Bouteiller, A. Danalis, M. Faverge, A. Haidar, T. Herault, J. Kurzak, J. Langou, P. Lemarinier, H. Ltaief, P. Luszczek, A. YarKhan, and J. Dongarra, "Flexible development of dense linear algebra algorithms on massively parallel architectures with DPLASMA," in *12th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC'11)*, 2011.

[15] F. D. Igual, E. Chan, E. S. Quintana-Ortí, G. Quintana-Ortí, R. A. van de Geijn, and F. G. V. Zee, "The FLAME approach: From dense linear algebra algorithms to high-performance multi-accelerator implementations," *J. Parallel Distrib. Comput.*, vol. 72, no. 9, pp. 1134–1143, 2012.

[16] T. George, V. Saxena, A. Gupta, A. Singh, and A. R. Choudhury, "Multifrontal Factorization of Sparse SPD Matrices on GPUs," *2011 IEEE International Parallel & Distributed Processing Symposium*, pp. 372–383, May 2011.

[17] R. F. Lucas, G. Wagenbreth, D. M. Davis, and R. Grimes, "Multifrontal computations on GPUs and their multi-core hosts," in *Proceedings of the 9th international conference on High performance computing for computational science*, ser. VECPAR'10. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 71–82.

[18] C. D. Yu, W. Wang, and D. Pierce, "A CPU-GPU Hybrid Approach for the Unsymmetric Multifrontal Method," *Parallel Computing*, vol. 37, no. 12, pp. 759–770, Oct. 2011.

[19] J. W.-H. Liu, "The role of elimination trees in sparse factorization," *SIAM J. Matrix Anal. Appl.*, vol. 11, pp. 134–172, 1990.

[20] I. S. Duff and J. K. Reid, "The multifrontal solution of indefinite sparse symmetric linear," *ACM Transactions on Mathematical Software (TOMS)*, vol. 9, no. 3, pp. 302–325, 1983.

[21] C. C. Ashcraft, R. G. Grimes, J. G. Lewis, B. W. Peyton, H. D. Simon, and P. E. Bjørstad, "Progress in sparse matrix methods for large linear systems on vector supercomputers," *International Journal of High Performance Computing Applications*, vol. 1, no. 4, pp. 10–30, 1987.

[22] M. Cosnard, E. Jeannot, and T. Yang, "Compact dag representation and its symbolic scheduling," *Journal of Parallel and Distributed Computing*, vol. 64, no. 8, pp. 921–935, 2004.

[23] V. Volkov and J. W. Demmel, "Benchmarking GPUs to tune dense linear algebra," in *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, ser. SC '08. Piscataway, NJ, USA: IEEE Press, 2008, pp. 31:1–31:11.

[24] P. Hénon, P. Ramet, and J. Roman, "On finding approximate supernodes for an efficient ILU(k) factorization," *Parallel Computing*, vol. 34, pp. 345–362, 2008.

[25] T. A. Davis, "University of Florida sparse matrix collection," *NA Digest*, vol. 92, 1994.

[26] I. Yamazaki, S. Tomov, and J. Dongarra, "One-sided Dense Matrix Factorizations on a Multicore with Multiple GPU Accelerators," *Procedia Computer Science*, vol. 9, no. Complete, pp. 37–46, 2012.

[27] C. Nvidia, "Cublas library," *NVIDIA Corporation, Santa Clara, California*, vol. 15, 2008.

[28] G. Tan, L. Li, S. Triechle, E. Phillips, Y. Bao, and N. Sun, "Fast implementation of dgemm on fermi gpu," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '11. New York, NY, USA: ACM, 2011, pp. 35:1–35:11.

[29] R. Nath, S. Tomov, and J. Dongarra, "An improved MAGMA GEMM for Fermi graphics processing units," *International Journal of High Performance Computing Applications*, vol. 24, no. 4, pp. 511–515, 2010.

[30] J. Kurzak, S. Tomov, and J. Dongarra, "Autotuning GEMM Kernels for the Fermi GPU," *IEEE Transactions on Parallel and Distributed Systems*, vol. 23, no. 11, pp. 2045–2057, 2012.