RSCHED: An Effective Heterogeneous Resource Management for Simultaneous Execution of Task-Based Applications

Etienne Ndamlabin, Bérenger Bramas
Inria Nancy – Grand Est, CAMUS Team,
Villers-lès-Nancy, France
ICPS Team, ICube, Illkirch, France

***Abstract*—Modern parallel architectures have heterogeneous processors and complex memory hierarchies, offering up to billion-way parallelism at multiple hierarchical levels. Their exploitation by HPC applications greatly boosts scientific discoveries and advances, but they are still not fully utilized, leading to proportionally high energy consumption. The task-based programming model has demonstrated promising potential in developing scientific applications on modern high-performance platforms. This work introduces a new framework for managing the concurrent execution of task-based applications, RSCHED. The framework aims to minimize the overall time spent executing a set of applications and maximize resource utilization. RSCHED is a two-level resources management framework: resource distribution and task scheduling, with sharable and reusable resources on the fly. A new model of Gradient Descent has been proposed, among other strategies for resource distribution, due to its well-known speedy convergence event in fast-growing systems. We implemented our proposal on StarPU and evaluated it on real applications. RSCHED demonstrated the potential to speed up the overall makespan of executed applications compared to consecutive execution with an average factor of 10x and the potential to increase resource utilization.**

***Keywords*—Heterogeneous resource management; scheduling; task-based applications; gradient descent; StarPU**

I. Introduction

High-performance computing (HPC) is crucial to making discoveries and advances in several scientific domains (astrophysics, climatology, epidemiology, biology, geology, etc.). HPC offers the ability to perform complex calculations and massive data processing at very high speed by aggregating the power of several thousand processing units, called supercomputers. Supercomputers rely on a complex, heterogeneous, and hierarchical hardware organization. The largest supercomputers are mostly composed of central processing units (CPUs) and graphical processing units (GPUs)[1], or even Field Programmable Gate Arrays (FPGAs). As parallel systems, they can process several jobs at the same time by scheduling their execution on the available resources.

In the current HPC paradigm, there are schedulers at multiple levels, that all have the same aim: distributing the workload over hardware resources. At the higher level, the batch-scheduler, like Slurm [2], OAR [3], or OpenPBS [4], manages the hardware resources of an entire supercomputer by deciding the order of execution of the jobs submitted by the users. Submitted jobs are treated by the batch scheduler as black boxes. This is advantageous because the batch scheduler can run applications implemented with any technology, giving freedom to the programmers. However, this approach might lead to resource wastage and increasing energy consumption. Such a situation can happen when an application inefficiently uses the allocated resources, when an adjustment of resources is required during different phases of execution, or when the resource manager cannot adapt the resources to the workload of newly submitted jobs. Especially since the dominant scheme for scheduling parallel jobs on parallel computers is known as variable partitioning [1], in which scheduled jobs have partitioned assigned processors they keep and use throughout their lifetime. However, executing one job after another is likely counterproductive, since HPC applications are often composed of interdependent executing kernels, and therefore cannot fully use resources due to their precedence constraints. In the current study, we aim to improve the batch scheduler by using a dynamic resource allocation strategy. Our objective is to improve the executions at the scale of the supercomputer, i.e. to reduce the overall makespan of an application set, and not to focus on a single application only. In addition, our work is tied to the task-based method, as we consider that each application is composed of tasks and that some of them can be executed on CPU, GPU or both.

In this paper, we present a resource manager for heterogeneous environments considering task-based model applications called RSCHED, aiming at optimizing their usage. In RSCHED, we propose strategies for resource distribution between concurrent task-based applications while orchestrating their execution. We have implemented our proposal within StarPU [2] and analyzed the performance of our proposal via diverse experiments. Our contributions can be summarized as follows:

- We propose a framework for managing the execution of concurrent task-based applications.
- We propose strategies for dynamically distributing resources between concurrent task-based applications.

[1]https://www.top500.org/

[2]https://www.schedmd.com/

[3]https://oar.imag.fr/

[4]https://www.openpbs.org/

- We propose a new model of Gradient Descent in a (three-dimensional) discrete space.
- We propose a strategy to automatically create and configure a scheduling context for a given task-based application.
- We present performance analysis and results showing the effectiveness of our proposals.

The remaining sections of the paper are organized as follows. In Section II, we introduce the notion of task-based application and present the state-of-the-art concerning the scheduling of task-based application under StarPU. Then, in Section III, we present our proposed resource management for simultaneous execution of task-based applications. Finally in Section IV, we evaluate the performance of our proposals.

II. Background

A. Task-Based Application

Several strategies to parallelize applications on heterogeneous computing nodes aim at maximizing resource usage. The task-based model has demonstrated high potential in various fields [3], [4], [5]. This method allows obtaining hardware-independent algorithm descriptions while developing efficient HPC applications. The level of abstraction and encapsulation relieves the users by shifting the complexity to the runtime systems, where researchers can invest the effort to create generic and efficient optimization solutions. The HPC community has at its disposal highly documented and maintained runtime systems supporting the task-based model, such as Parsec [6], and StarPU [2] a runtime system library developed at Inria Bordeaux.

Among other runtime systems supporting the task-based model, StarPU has a plus in that it has a component – *hypervisor* – allowing concurrent execution of task-based applications with minimal interference [7]. StarPU hypervisor provides confined execution environments – *scheduling contexts* – which can be used to partition computing resources. StarPU scheduling contexts can be dynamically resized and linked to a well-designed scheduler to optimize the allocation of computing resources among concurrent task-based applications/libraries. A scheduler can be chosen for each application via its linked scheduling context. Since task-based applications have various types and structures, a scheduler can be effective only for some applications, or a specific type of machine architecture [8], [9], [10], [11], [12]. StarPU also proposes basic strategies for resizing scheduling contexts and a platform for implementing additional custom ones.

Different task-based frameworks have been used to develop efficient HPC applications, such as the Lattice-Boltzmann method [13], [14], the fast-multipole method (FMM) [3], [4], [15], N-body simulations [16], linear algebra solvers [17], [18], [19], H-matrix solvers [20], the particle-in-cell method [21], the polar decomposition method [22], seismic imaging [23], [24], [25], Galerkin solver [26], to mention a few. This demonstrates that at least at a moderate scale and when used by experts, the existing task-based runtime systems can be efficient for various classes of algorithms.

1) Task-based parallelization: The task-based method divides an application into interdependent sections, called tasks. The dependencies between the tasks ensure valid parallel executions and task execution orders without race conditions. This can be likened to a graph, where the nodes represent the tasks and the edges represent the dependencies. We consider a task-based application as a Directed Acyclic Graph (DAG) $G(V, E)$ where $V = \{t_1, t_2, ..., t_n\}$ is the set of nodes and $E = \{e_{i,j} = (t_i, t_j) | 1 \leq i, j \leq n, i \neq j\}$ the set of edges representing the existing data dependencies between tasks. An edge $(t_i, t_j) \in E$ if there is a precedence constraint between t_i and $t_j \in V$, such that t_j can be executed only after the task t_i is over, and the data made available.

A task t_i is a computational element executable on one or (potentially) several types of hardware and incorporates different interchangeable kernels, each targeting a specific architecture. For instance, a matrix-matrix multiplication task in linear algebra could be either a call to cuBLAS and executed on a GPU, or a call to Intel MKL and executed on a CPU, but both kernels return equivalent results.

B. Task Scheduling and Related Work

The scheduling problem on heterogeneous computing systems has been proven NP-complete [27], whether in static or dynamic situations. Dealing with the first situation requires prediction models which are not always accurate, and a knowledge of the complete view of the task graph [28], which need expensive analysis mechanisms and incur significant overhead. The latter one is the most used [29], [30], [31], [25], [32], [22], [4] and has demonstrated its ability to deliver high performance with reduced overhead.

The two main steps of a scheduler are task selection or prioritization, and resource selection. This action can be static or dynamic according to the scheduling situation.

Due to the evolution of computing architectures, task scheduling in heterogeneous computing is an aged but hot topic. Various strategies for task scheduling have been proposed, with Heterogeneous Earliest Finish Time (HEFT) being one of the most widely used. [33] In HEFT, tasks are prioritized using a heuristic based on a prediction of the processing length of the tasks and the data transfer time between them. Whereas, resource selection is based on a heuristic that determines the resource providing the best finish time for the tasks according to the scheduling decision of previous tasks. Several variances of this approach with more advanced ranking and resource selection models have been proposed [34], [35], [36], [37]. These schedulers have in common the limitations of static schedulers previously argued, and therefore rely on greedy algorithms. In a changing environment, re-prioritizing the tasks could be necessary, which can add more overhead. A larger spectrum of task schedulers can be found in the literature [38], [39].

1) Task scheduling in StarPU: Our scheduling process follows the terminology of StarPU. In StarPU, the user first splits the problem into smaller computational tasks. Afterwards, the tasks are implemented into codelets, which are simple C functions. One task can be implemented differently in several codelets according to the targeted hardware, allowing the user to harness special accelerators, such as vectorial CPU cores

or OpenCL devices. In StarPU terminology, these devices are called workers. For each task, the user also has to describe precisely the input data, in read mode, and the output data, in write or read/write mode. StarPU considers that a scheduler has an entry point where the ready tasks are pushed, and it provides a request method where workers pop the tasks to execute, as depicted in Fig. 1.

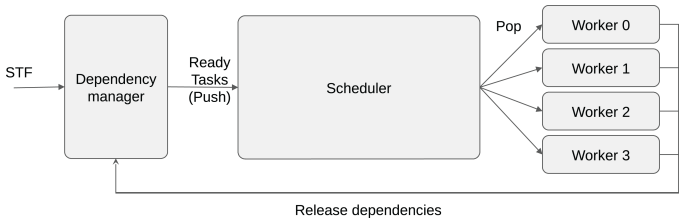

Fig. 1. Schematic view of task-based runtime system organization. A program as a sequential task flow (STF) model and converted into tasks/dependencies by the RS. New ready tasks pushed on resolved dependencies. Any idle worker calls the scheduler pop function to request a task to execute.

In StarPU, both pop/push methods are directly called by the workers that either release the dependencies or ask for a task. Consequently, assigning a task to a given worker means returning this task when the worker calls the pop method. During the execution of a StarPU program, it is possible to choose among several schedulers. The DMDA (deque model data-aware) scheduler is one of the most famous and sophisticated. It uses a HEFT-like strategy and tries to minimize the makespan by using a look-ahead strategy and data transfer costs. Another effective StarPU scheduler is Heteroprio [40], [4], a semi-automatic scheduler designed for heterogeneous machines where users must provide task priorities. A fully automatic version of the Heteroprio scheduler that computes efficient priorities is proposed [8]. Another extension of Heteroprio is the MulTreePrio [9] scheduler based on a set of balanced trees data structure, in which assignment of tasks to available resources is done according to priority scores per task for each type of processing unit. MulTreePrio makes overall good scheduling results thanks to its fast and efficient heuristics, despite the considerable variety of DAG structures from one application to another.

In all of the above, one task-based application is considered for execution/scheduling.

2) Scheduling concurrent task-based applications: In general, there is contention for the usage of resources in an HPC environment. Each user's application requests a number of processing resources (CPU/GPU for instance), an entire node or part of it. However, in both cases, it is up to the user to determine the number and type of resources, therefore this might often lead to resource wastage.

The problem presented here has a completely distinct parallelization and resources management approaches in cloud computing and Big-data frameworks, such as Spark[5] or Apache Hadoop[6]. In cloud computing, the scheduler orchestrates different executions from different types of applications, such as Big data programs, over given hardware resources. It has a view on the different operations that compose the executions, hence it can schedule and interleave them finely. However, there is a gap between the programming model and resource management.

In the basic HPC context, there is not yet a dynamic solution for concurrent job execution on the same resource as it is done in a cloud environment. In Slurm for instance, there is the notion of Job Array which consists of submitting and managing collections of similar jobs such that they may run in parallel with different input parameters or data on a node. However, it is the responsibility of the programmer to orchestrate the execution of the tasks over the resources. The same problem is encountered when several jobs are submitted separately, but now at the level of the scheduler.

The RECIPE project [41], [42] attempts to control the resources more efficiently without bridging the gap with the applications, which will end as being oriented to cloud computing instead of HPC.

Most recently, other researchers proposed a method to build scalable containerized HPC clusters in the Cloud [43], [44], by containerizing three batch schedulers, namely SLURM, OpenPBS, and OAR. They attempt to solve the problem of scaling, dynamically adding or removing containerized HPC nodes, without altering neither the Cloud orchestrator nor the HPC scheduler. This works presented promising preliminary results in that direction, scaling jobs do not impact running or pending jobs. However, it is still in the direction of Cloud Computing, where some researchers are trying to answer the question "Is the Cloud able to encompass all the categories of scientific issues in a unified way?". In addition, this does not consider task-based applications, but MPI-based applications only.

In most of the works presented above, it is either cloud oriented, or based on basic batch scheduling-like approaches using the variable partitioning scheme [1]. In both case, we can have more idle resources, which could have been used by starved applications, or could have helped in fastening the overall completion of applications and therefore in optimizing energy consumption. However, StarPU [7] offers a good platform to dynamically partition computing resources into contexts that can be used to execute several applications, and with the possibility to share resource between them. Nevertheless, there is no means to launch several applications nor dynamically distribute the workers among them.

III. Presentation of RSCHED

In this section, we present our concurrent execution approach, named RSCHED. Let us consider that we have n concurrent users' applications requesting resources for execution in a node with p workers (nb_cpus the number of CPUS and nb_gpus the number of GPUS). The main objective of RSCHED is to minimize the overall makespan for all the n applications. A secondary objective of RSCHED is maximizing resource utilization, which is a well-known energy consumption reduction approach (Fig. 2).

The RSCHED framework has a two-level scheduling model: resource distribution and task scheduling. The scheduling at each level assumes the resources are sharable between

[5]https://spark.apache.org/
[6]https://hadoop.apache.org/

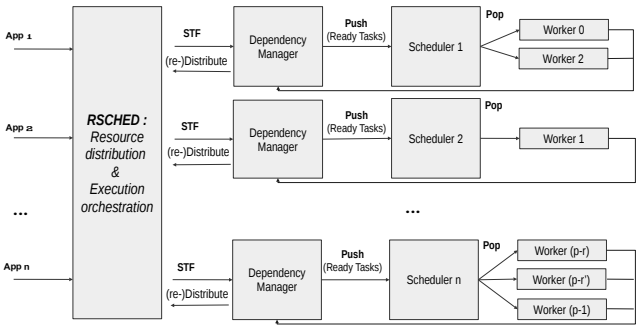

Fig. 2. Schematic view runtime system organization with n concurrent task-based applications and p workers.

the applications, and reusable on the fly (for instance when an application ends), even without a new resource subscription. The resource distribution aims at distributing the available resource among the application, and task scheduling to effectively map each task unto a given compute resource.

That said, our model is intended to be more flexible for efficient resource usage than basic batch scheduling-like approaches. Idle resources can be reused by starved applications (even before the completion of the application to which they have been assigned). For shared resources, the end of one application can help in fastening the completion of others.

Before the execution of an application, a task scheduler is associated with a separate context linked to it. The required resource distribution can be done before any task is pushed, or after all tasks have been pushed but no task has been executed. In the beginning, the lack of information on the applications makes it difficult to process an advanced distribution strategy. That is why in the first case, a naive strategy can be used to distribute the workers between the contexts, and afterward a more advanced distribution strategy. In the second case, we can have sufficient information on the applications to process an advanced distribution strategy before any task is executed.

A compromise between having all information before starting and starting sooner is to process an advanced distribution at an arbitrary time after the tasks have started to be pushed and executed. That time could be for instance when any first application completes its execution, or after all tasks have been pushed. The execution of tasks starts as soon as possible, and afterward, we re-distribute workers to the rest of the applications to balance the load and therefore minimize the overall makespan. While distributing resources, two contexts may share some workers, but with a possible performance penalty, due to context switches.

We assume all the n applications were submitted before the distribution. However, we propose early investigations of continuous application arrivals (mimicked by the proposed resizing options) that will be properly and intensively investigated in our future work.

A. RSCHED API

To distribute resources among task-based applications, the RSCHED's API requires their performances on the targeted hardware. We assume in this work that we have two types of workers, CPUs and GPUs. For each application, the required information are the following:

- CPU_W: Total (sequential) CPU workload of the graph on the targeted CPU.
- GPU_W: Total (sequential) GPU workload of the graph on the targeted GPU.
- CPU_{PW}: Total (sequential) pure CPU workload of the graph on the targeted CPU.
- GPU_{PW}: Total (sequential) pure GPU workload of the graph on the targeted GPU.

By pure CPU (or GPU) workload, we mean the workload of tasks that can only be executed by a CPU (or GPU). Providing those pieces of information implies knowing (an approximation of) the processing time of each task in the application per type of worker. There are several means of obtaining such information, like Machine Learning, or history-based performance models as it exists in StarPU.

The constraints over this information are given by the Rule 1.

Rule 1. *Given the required information as defined above, either the total workload is equal to the pure one for all the types of workers, or strictly greater than the pure one for all (see Eq. 1 and 2).*

$$CPU_W \geq CPU_{PW} \ \boldsymbol{AND} \ GPU_W \geq GPU_{PW} \tag{1}$$

$$(CPU_W = CPU_{PW} \ \boldsymbol{AND} \ GPU_W = GPU_{PW}) \ \boldsymbol{OR} \\ (CPU_W > CPU_{PW} \ \boldsymbol{AND} \ GPU_W > GPU_{PW}) \tag{2}$$

Proof: The proof of Eq. 1 is obvious. For Eq. 2, there are two cases to have equality: all the tasks are either pure CPU or pure GPU. In both cases, the other type of worker will have zero workload. ■

1) RSCHED resource distribution: Given the n applications with the required information described in the last section, a distribution strategy should produce the following information for each graph:

- L_{CPUS}: List of CPUs assigned.
- L_{GPUS}: List of GPUs assigned.
- $\#CPUS$: Number of distinct CPUs assigned ($\#CPUS = |L_{CPUS}|$).
- $\#GPUS$: Number of distinct GPUs assigned ($\#GPUS = |L_{GPUS}|$).
- $CPUS_{PR}$: The power rate of the assigned CPUs ($0 \leq CPUS_{PR} \leq \#CPUS$).
- $GPUS_{PR}$: The power rate of the assigned GPUs ($0 \leq GPUS_{PR} \leq \#GPUS$).

Rule 2. *Each application must receive at least one worker, and applications can share all the workers.*

$$(\#CPUS + \#GPUS > 0.0) \ \boldsymbol{AND} \\ (0.0 \leq \textstyle\sum_{0 \leq r \leq n} \#CPUS_r \leq n \times nb\_cpus) \ \boldsymbol{AND} \\ (0.0 \leq \textstyle\sum_{0 \leq r \leq n} \#GPUS_r \leq n \times nb\_gpus) \tag{3}$$

Rule 3. *Each application must receive at least one worker per type of pure workload.*

$$(CPU_{PW} = 0.0 \ \boldsymbol{OR}\ (CPU_{PW} \neq 0.0 \ \boldsymbol{AND}\ \#CPUS > 0.0)) \ \boldsymbol{AND}\ (GPU_{PW} = 0.0 \ \boldsymbol{OR}\ (GPU_{PW} \neq 0.0 \ \boldsymbol{AND}\ \#GPUS > 0.0)) \tag{4}$$

An estimation of the makespan of each application is used as a building block of our strategies, given a set of workers (CPUs/GPUs) assigned to the applications. We proposed an estimation called "Ideal Makespan", and more details are given in Appendix (see Algorithm **??**). The following metrics are used in the evaluation of our "Ideal Makespan".

When an application has a pure workload for a given type of worker, there is a minimum length constraint over its makespan. We denote them as $tgpu_{minM}$ and $tcpu_{minM}$ (see Eq. 5 and 6).

$$tgpu_{minM} = \frac{GPU_{PW} \times coef\_par\_eff^{\#GPUS+\#CPUS-1}}{\#GPUS} \tag{5}$$

$$tcpu_{minM} = \frac{CPU_{PW} \times coef\_par\_eff^{\#GPUS+\#CPUS-1}}{\#CPUS} \tag{6}$$

In the case no CPUs (or GPUs) are assigned unto the application, $tcpu_{minM}$ (or $tgpu_{minM}$) is equal to zero.

The general formulation of our makespan estimation is given by Eq. 7.

$$\begin{aligned} ideal\_makespan \ &= MAX(tcpu_{minM}, tgpu_{minM}) \\ &+ \frac{cpu\_rem\_wl}{\#CPUS + \#GPUS \times \frac{cpu\_rem\_wl}{gpu\_rem\_wl}} \end{aligned} \tag{7}$$

Where cpu_rem_wl (resp. gpu_rem_wl) is the exceeding CPU (resp. GPU) workload compared to the gap between $tgpu_{minM}$ and $tgpu_{minM}$, and the CPU/GPU (resp. GPU/CPU) speedup.

In this work, we present four distribution strategies: Lp-Solve, MinMaxWL (Min-Max Workload balancing), DSR-CLUS (Dedicated plus Shared Resource with Clustering) and DSR-GD (Dedicated plus Shared Resource with Gradient Descent).

a) LpSolve: In this strategy, we rely on the linear programming model presented in a previous study [4]. Originally, this model was employed to compute an ideal makespan (a theoretical lower bound) for tasks executed on heterogeneous architectures. The model is given by:

$$\begin{cases} Objective\,function : min(T) \\ \sum\limits_{\omega\ in\ \Omega} \alpha_1^\omega\, t_1^\omega = t_1 \leq T \\ \sum\limits_{\omega\ in\ \Omega} \alpha_2^\omega\, t_2^\omega = t_2 \leq T \\ \quad ... \\ \sum\limits_{\omega\ in\ \Omega} \alpha_P^\omega\, t_P^\omega = t_P \leq T \end{cases} \tag{8}$$

$$\begin{cases} \sum\limits_{p=1}^{P} \alpha_p^1 = 1 \\ \sum\limits_{p=1}^{P} \alpha_p^2 = 1 \\ \quad ... \\ \sum\limits_{p=1}^{P} \alpha_p^{|\Omega|} = 1 \end{cases} \tag{9}$$

Here, P denotes the number of processing units and $|\Omega|$ is the total number of tasks. The coefficient α_p^ω indicates the proportion of task ω processed by unit p, and t_p^ω represents the time taken to complete task ω on unit p, given that this duration varies based on the type of the processing unit. Accordingly, the first system determines the computation duration for each unit, with T being the longest duration. The second part ensures that each task is computed at 100%.

While this model provides an upper bound for ideal performance, it doesn't account for the dependencies between the task order and consider that tasks can be divided among various processing units.

In our adaptation, we assume that a single application consists of three tasks: one for exclusive CPU work, another for exclusive GPU work, and the last one for work that can be executed on either CPU or GPU. Given this characterization, the aforementioned LP model remains applicable.

However, our focus isn't on the ideal makespan T, but on the α coefficients as we aim to find the most efficient way to distribute the application across processing units. Although the LP provides an optimal distribution for an ideal system, it also guides us on the proportion of each application that should be allotted to each processing unit type. However, this distribution might be inefficient in real-world scenarios, leading to a task being fragmented across all units or a skewed allocation, like 99% on one unit and 1% on another, which may not always be practical. Consequently, it's essential to transform these coefficients into practical distribution values to derive a feasible scheduling strategy.

To achieve this, we use a two-step method. First, we sum the distribution coefficients per unit type to determine the fraction of each application designated for every processing unit type. For instance, if the LP solution suggests distributing an application as 0.1 and 0.4 across two CPUs, and 0.2, 0.2, and 0.1 across three GPUs, we infer it should be equally split (0.5 for CPU and 0.5 for GPU). Subsequently, we decide on the application distribution based on these values. In the next step, we compute the processing time for each unit type by multiplying the number of a given unit type with T. For instance, with two CPUs and a makespan of 10s, we have 20s of total CPU time to distribute. Each application is then assumed to use a fraction of this time proportional to its distribution coefficient. Our greedy algorithm identifies the application with the highest use proportion and allocates it to the processing unit with the least utilization, continuing until every application has been entirely mapped to processing units.

b) DSR (Dedicated plus Shared Resource) strategies: For illustrations, let us consider for instance that we have three applications, four CPUs, and two GPUs. The DSR strategies

proceed in two steps to distribute the workers among the applications:

The first step consists in assigning dedicated (unshared) workers (GPUs or CPUs) to each application, proportionally to their GPU/CPU workload compared to the sum of all the applications' workloads. Supposed the proportions of CPU workloads (cpu_pwl) are 0.31, 0.56, and 0.13, the numbers of dedicated CPUs (given by $\lfloor nb\_cpus \times cpu\_pwl \rfloor$) will be respectively $\lfloor 4 \times 0.31 \rfloor = \lfloor 1.24 \rfloor = 1$, $\lfloor 4 \times 0.56 \rfloor = \lfloor 2.24 \rfloor = 2$, and $\lfloor 4 \times 0.13 \rfloor = \lfloor 0.52 \rfloor = 0$. In that case, the number of remaining CPUs is 1. The same is similarly done for GPUs. In the case of hybrid workloads, and if the standard deviation between the GPU/CPU speedups is above a certain threshold, we proceed to the barter which consists of exchanging GPU against CPUs to accelerate the most GPU-optimized applications.

The second step involves sharing the remaining workers to the applications, using a given technique. Here we proposed two DSR strategies, based on two different techniques for workers sharing, the clustering and the gradient descent: DRS-CLUS (Dedicated plus Shared Resource with Clustering) and DSR-GD (Dedicated plus Shared Resource with Gradient Descent). The sharing process is done based on the remaining cpu_pwl (here in the case of CPUs, we have: 0.24, 0.24, and 0.52).

As for the DRS-CLUS strategy, resource sharing is done as follows. The number of clusters is equal to the number of remaining workers. The workers/applications mapping is done to balance the load over the workers as much as possible.

In the DSR-GD strategy, the second step is done as follows (see Algorithm **??** in Appendix, from lines 18 to 24). The remaining workers are shared between applications using a new model of gradient descent (GD) in a (three-dimensional) discrete space.

The Gradient Descent is an efficient strategy well-known for its speedy convergence in convex and smooth optimization problems (if well-tuned), even in low memory, and computational loading environments [45], [46].

Our GD strategy is modelled as follows:

Given the CPUs and GPUs IDs, {0,1,2,3} and {0,1} respectively, our research space is modelled as a three-dimensional discrete space (X, Y, Z):

- X: The numbers of assigned CPUs ($\#CPUS$) per graph. Each application can have from zero to the number of CPUs (Ex. {1,2,1}, the first app has 1 CPU, the second one 2 CPUs, and the third one has 1 CPU).
- Y: The numbers of assigned GPUs ($\#GPUS$) per graph. Each application can have from zero to the number of GPUs (Ex. {1,1,1}, the first app has 1 GPU, the second one 1 GPU, and the third one has 1 GPU).
- Z: The possible (graph-to-gpu/cpu) mappings given X and Y. An example of mapping related to the Y and Y ones above is $\{\{0\},\{0\}\}$, $\{\{1,2\},\{0\}\}$, $\{\{3\},\{1\}\}$. In this example, the first app has the CPU id=0 and the GPU id=0, app 2 has the CPU ids=1,2 and the GPU id=0 and app 3 has the CPU id=3 and the GPU id=1. We can see that app 1 and app 2 share one GPU together.

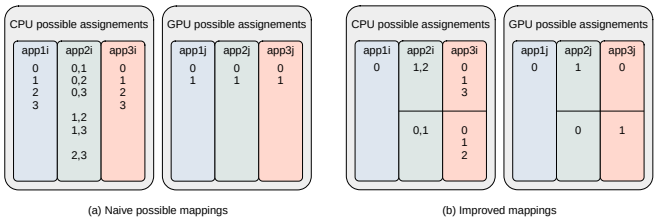

Fig. 3. Illustration of determination of possible mappings with three applications, four CPUs, and two GPUs, and given that X={1,2,1} and Y={1,1,1}.
The improved version has fewer duplicates in terms of obtained mappings 24 vs 768).

Iteration in the axis is done as follows:

- The indexes in X and Y can be seen as permutations with repetition of a possible number of assigned workers in nb_graphs positions:
  - $X_i \in$ {{0,0,0}, {0,0,1}, ... , {0,0,4}, {0,1,0}, ... ,{4,4,4}}. By looking carefully, we can see that each X_i is likened to "i" in base $(nb\_cpus + 1)$
  - $Y_j \in$ {{0,0,0}, {0,0,1}, {0,0,2}, {0,1,0}, ... ,{2,2,2}}. In like manner, Yi is likened to "j" in base $(nb\_gpus + 1)$
- Z_k (or Zijk): The k-th possible mapping, given X_i and Y_j. One naive way to get the Z_k's values is to generate the different 2-uplets, of the possible CPUs assignment given X_i and the possible GPUs assignment given Y_j. For our example, Fig. 3(b) illustrates the determination of possible mappings and how to find the Z_k. We have $Z = (app1i \times app2i \times app3i) \times (app1j \times app2j \times app3j)$. While the naive version leads to $((4 \times 6 \times 4) \times (2 \times 2 \times 2) = 96 \times 8 =)$ 768 possible mappings, the improved one gives $((1 \times 2 \times 3) \times (1 \times 2 \times 2) = 6 \times 4 =)$ 24 possible mappings. The improved version has far fewer duplicates in obtained possible mappings, and this is for only three applications.

This model presents the whole research space. However, during our research process, the search space is circumscribed around the convex zone containing our global optimum. This is done by setting for each application the minimum number of workers (CPUs and GPUs) obtained in step one, and the maximum by adding the number of workers not yet assigned. Thus, we reduce the search space and speed up the search.

Our gradient function is evaluated as a symmetric linear interpolation [47]. Our search process follows the pattern direction [48], first, we search towards the X direction, then the Y direction, and finally the Z direction. To ensure process time scaling, the learning rates for the different axes are $tau_x = 0.001 \times (\lfloor nb\_graphs/\sqrt{nb\_cpus} \rfloor + 1)$, $tau_y = 0.001 \times (\lfloor nb\_graphs/\sqrt{nb\_gpus} \rfloor + 1)$ and $tau_z = 0.0005 \times (\lfloor nb\_graphs/\sqrt{nb\_cpus + nb\_gpus} \rfloor + 1)$.

c) MinMaxWL (Min-Max Workload balancing): The MinMaxWL algorithm is a load-balancing strategy that dis-

tributes workers among applications by minimizing the maximum ideal makespan. The strategy is depicted in Algorithm 2 (in Appendix), and has four main steps.

First of all, it assigns one worker to each application having a pure workload according to the type of worker (from lines 3 to 13). While trying to assign a worker to the current application, if there are no remaining workers of the type, the application shares one with the under-loaded application related to that type. A backpropagation is employed in the case of sharing to ensure load-balancing when less-loaded applications are treated after more-loaded ones.

The second step is to ensure all the applications have at least one worker, by assigning a worker to applications without a pure workload (from lines 15 to 28). Since those applications are hybrid, the type of worker to assign is the fastest on the application. A similar sharing process is also employed, but this time the sharing is made on the worker, leading to the smallest makespan at the point.

Finally, while there are remaining workers (per type of worker), it assigns a worker to the application that will minimize the maximum ideal makespan among all the applications.

2) Distribution options: If the distribution of resources to the applications is accurate, the applications will end almost at the same time, and so will the workers. Otherwise, some resources could be idle for a long, while remaining applications may need them. To deal with that situation, a redistribution of the resources might be necessary. One crucial aspect to consider for it is the condition of resizing, the when.

The condition of resizing we employed is the following:

- An application just ended, and there remain applications to run.
- There is a significant standard deviation between the progress rate of the applications.

The estimation process time (workload) for a CPU or GPU may differ from the effective processing time during the execution. For instance, let us suppose an application with a CPU workload of 100s, and that has been assigned 10 CPU workers. Suppose 5s after executions start, there is a need for redistribution, and it remains at an overall 20s processing time for the workload. We would have expected having executed $5 \times 10 = 50$s for the application, whereas we have $100 - 20 = 80$s. The progress rate in this case is therefore equal to $80/50 = 1.6$; which means the application is running faster than expected. Now we know that possibly the application may end in $(20/1.6)/10 = 1.2$s instead of $20/10 = 2$s.

Before the redistribution, we adjust the workload of each application according to its progress rate. We have proposed two resource redistribution options.

a) One Distribution: This is the default behavior, where the distribution is done once and for all.

b) Multiple distributions: Here we do the distribution as initially, but considering the adjusted workloads of the remaining applications.

c) Inherit released workers: Here we distribute the released workers (by the just-ended application) to the remaining ones, with high privilege to those that were delaying.

B. RSCHED Implementation in StarPU

StarPU offers a platform to dynamically construct, delete, and modify Scheduling Contexts, which are used to execute several parallel kernels in an isolated way and without interference. This allows the users to assign workers to the contexts, at their creation time, or resize them during program execution. However, this is subject to the knowledge of the number of workers needed for each scheduling context. StarPU proposes online performance tools to monitor the execution of tasks, to make execution time estimations.

1) Multiple task-based applications: There are several applications implemented in StarPU. However, there is no mechanism to launch or orchestrate the execution of concurrent applications. For the sake of simplicity, instead of using several different applications, we have exploited the implementation of Cholesky factorization to have several independent applications. The Cholesky application in StarPU is implemented with a performance model for each codelet. We have added a parameter to specify the number of applications to create, and for each application, we gave the possibility to specify the size and the number of blocks via environment variables.

2) Context creation and workload determination: For each application, a separate context is created and a task scheduler is associated with it. In this work, we have chosen to use DMDA as a scheduler for all the applications. DMDA relies on a historical performance model to be able to estimate in advance the duration of a codelet on each kind of processing unit. Using StarPU historical performance model, we have been able to compute the different workloads of the concurrent applications.

IV. Performance Study

A. Experiments Setup

1) Hardware: We conducted our experiments on three configurations with different GPU models as follows:

- A100: Composed of two 32-core AMD Zen3 EPYC 7513 @ 2.60 GHz, and 2 NVIDIA A100 (40GB). We use 30 CPU cores and 16 CUDA streams per GPU.
- Quadro: Composed of 2 Icosa-core Cascade Lake Intel Xeon Gold 5218R CPU @ 2.10 GHz, and 2 NVIDIA Quadro RTX8000 (48GB). We use 30 CPU cores and 16 CUDA streams per GPU.
- K40M: Composed of 2 Dodeca-cores Haswell Intel Xeon E5-2680 v3 2.5 GHz, and 4 K40m GPUs (12GB). We use 20 CPU cores and 8 CUDA streams per GPU.

We have configured StarPU as follows. For each configuration, we set the environment variables STARPU_NCPU to the number of CPU cores, STARPU_NCUDA to the number of GPU, and STARPU_NWORKER_PER_CUDA to the number of CUDA streams. Therefore, for all the configurations, we have more GPU workers than CPU ones.

2) Task-based applications: We implemented Cholesky factorization in StarPU and tested it across twelve different configurations as follows:

- app_0: Matrix size of 3.200, with 5 blocks
- app_1: Matrix size of 3.200, with 10 blocks
- app_2: Matrix size of 6.400, with 10 blocks
- app_3: Matrix size of 6.400, with 20 blocks
- app_4: Matrix size of 9.600, with 10 blocks
- app_5: Matrix size of 9.600, with 30 blocks
- app_6: Matrix size of 19.200, with 20 blocks
- app_7: Matrix size of 19.200, with 30 blocks
- app_8: Matrix size of 25.600, with 40 blocks
- app_9: Matrix size of 25.600, with 80 blocks
- app_{10}: Matrix size of 76.800, with 80 blocks
- app_{11}: Matrix size of 76.800, with 120 blocks

3) Software configuration: For each application, we have made different affinities related to the types of compute units (CPU or GPU). By default, all the tasks of the Cholesky application have two codelets, one for CPU and one for GPU. Overall, we have used the three following affinities:

- Default (*affinity0*): each task has one CPU codelet and GPU codelet.
- Only CPU (*affinity1*): all the tasks have only a CPU codelet.
- Only GPU (*affinity2*): all the tasks have only a GPU codelet.

To analyze the influence of the number of concurrent applications, and of the workload, we have made experiments with 3, 6, and 12 concurrent applications. To be in accord with a realistic scenario, we have shuffled the list of applications in each experiment. Then we took consecutive applications to form the groups. For instance, in the case of three applications, we have executed concurrently the applications at the first, second third positions, then the fourth, fifth, and sixth positions, and so on.

B. Metrics

In our experiments, we have compared our four distribution strategies against the concurrent execution using a unique context with all the workers (DMDA_CONC), and against sequential execution (i.e. one application after the another) using a unique context with all the workers (DMDA_SEQ).

As metrics, we have considered the speedup, the data transfer, and the resource utilization efficiency (RUE). We also compared the distribution processing time of our strategies.

The RUE is a new metric hereby introduced and defined as follows:

Definition 1. *We define the RUE as the ability to maximize the utilization of the resource, that is, using the adequate number and types of resources for the execution of each application. The RUE is given by Eq. 10, which is the product of the resource utilization and the efficiency [49].*

$$RUE = \frac{\sum_{p \in USED\_WK}\{processing\_time\ of\ worker\ p\}}{\sum_{p \in USED\_WK}\{total\_active\_time\ of\ worker\ p\}} \times \frac{speedup}{|USED\_WK|} \tag{10}$$

$USED\_WK$ is the list of distinct workers (CPU or GPU) used for the execution of the applications, whether concurrently or sequentially (i.e. in DMDA_SEQ). We normalized the RUE such that the values lie between 0 and 1.

C. Experiments Results and Analysis

1) Default experiments: We first present the performance of our strategies (LpSolve, MinMaxWL, DSR-GD, and DSR-CLUS), and of DMDA_CONC, against DMDA_SEQ, in terms of Speedup, then in terms of data transfer, and finally in terms of RUE.

a) Speedup: The big picture of the speedup realized by the different strategies compared to DMDA_SEQ is presented in Fig. 4. For all the configurations and affinities, the LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, and DMDA_CONC can significantly accelerate the execution DMDA_SEQ (Fig. 4b). Moreover, our strategies (except MinMaxWL) perform better even than DMDA_CONC with the increase in the number of applications. We observe in this study an outperformance over DMDA_CONC in more than 50% of cases for LpSolve, and more than 75% of cases for DSR-GD and DSR-CLUS (Fig. 4b).

DSR-GD and DSR-CLUS reach an acceleration of 40× compared to DMDA_SEQ. However, DSR-GD outperforms DSR-CLUS in more than 50% of situations, observed while we have an increase in applications. This means that conceptually, the Gradient Descent performs better than the clustering since the two strategies have the same building block. We observe that the speedup of the strategies increases with the number of concurrent applications (Fig. 4a). The study of the variation of the speedup according to the workload and of the GPU/CPU acceleration (see Fig. 5 and Fig. 6) reveals that DSR-GD and DSR-CLUS perform better when the percentage of the standard deviation of GPU/CPU acceleration among the application increases. This is explained by the employed bartering technique that gives GPU in preference to more accelerated applications in exchange for CPU to others.

The study of the variation of the speedup according to the number of applications over the different hardware configurations (Fig. 7) reveals that the strategies perform better on recent architectures (Quadro and A100) which have more accelerated GPU than on older ones (K40M). More specifically, DSR-GD and DSR-CLUS perform better than the other strategies, due to the same reasons as previously.

Globally, the DSR-GD realizes better speedup and in more of the situations than the others, then DSR-CLUS followed by LpSolve.

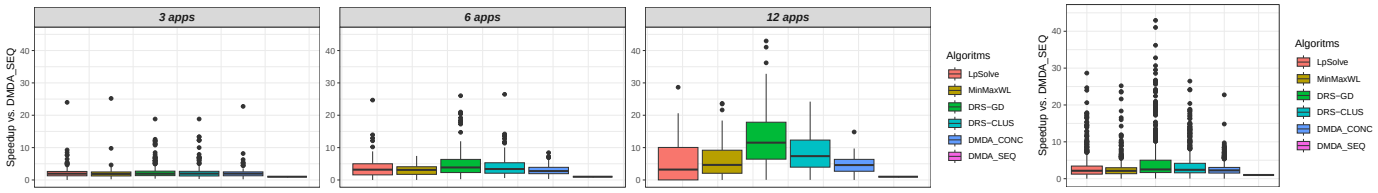

(a) Speedup obtained for the 3, 6, and 12 concurrent applications.

(b) Speedup summary

Fig. 4. Speedup of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ for all the affinities (affinity0, affinity1, affinity2) and all the hardware configurations (K40M, Quadro, and A100).

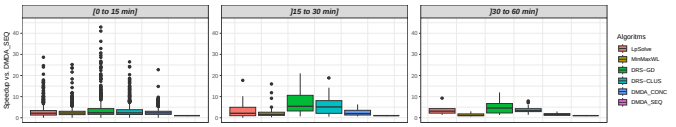

Fig. 5. Speedup of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ according to the average workload in minute, on K40M, Quadro or A100.

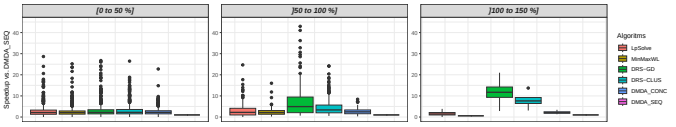

Fig. 6. Speedup of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ according to the standard deviation percentage of GPU/CPU speedup between the applications, on K40M, Quadro or A100.

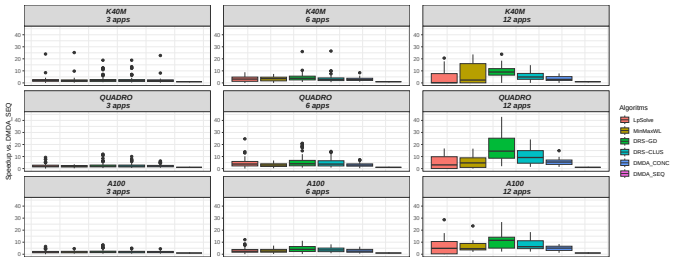

Fig. 7. Speedup of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ for 3, 6, and 12 concurrent applications with all the affinities, on K40M, Quadro or A100.

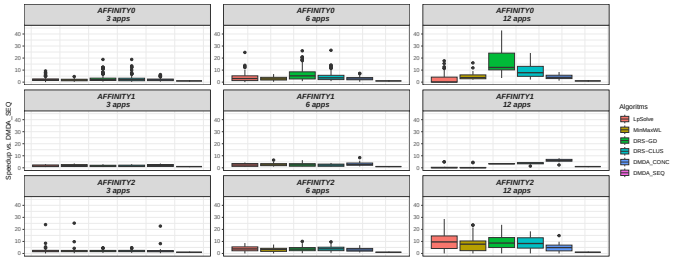

Fig. 8. Speedup of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ for 3, 6, and 12 concurrent applications with each affinity, on K40M, Quadro or A100.

b) Data transfer: The total amount of memory transfer obtained with the different strategies are provided in Fig. 9. All the strategies for concurrent execution used in this study (LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, and DMDA_CONC) significantly reduce the total memory transfer compared to the sequential execution (Fig. 9b), DSR-CLUS been the best one.

c) Resource utilization efficiency: The Normalized RUE obtained with the different strategies are provided in Fig. 10. We observe in this study that the concurrent execution of applications leads to more effective resource usage than the sequential one. The DSR-GD and DSR-CLUS strategies are more efficient in terms of resource utilization than the other (Fig. 10b), DSR-GD been the best one as the number of applications increases. Executing the applications sequentially one after the other leads to more resource wastage, which is known as a major cause of energy consumption in data centers [10], [11]. Moreover, we observe that the improvement of our strategies (DSR-GD, DSR-CLUS, and LpSolve) in terms of RUE correspond with the situations in which they are speeding up compared to the sequential execution (Fig. 10a $\equiv$ Fig. 4a). Therefore, succeeding in speeding up the sequential execution of tasks-based applications using our strategies might also help to reduce energy consumption, thanks to the effectiveness of the scheduler used.

2) Distribution processing time and options:

a) Distribution processing time: In a dynamic situation where applications arrive continuously (as we will study in the future), the decision processing has to be fast. Fig. 11 presents the evolution of processing time for each of our proposed strategies according to the acceleration of GPU compared to the CPU, and the number of executed applications.

The MinMaxWL and DSR-CLUS strategies are faster than LpSolve and DSR-GD which are meta-heuristics. However, the processing times of all the strategies are relatively small to expect good behavior even in a dynamic and computational loading environment. Moreover, even though we add the decision process time to the overall makespan, we will still have almost the same results as presented above.

Furthermore, we realize that DSR-GD scale better than LpSolve given their processing time (Fig. 11) and speedup compared to DMDA_SEQ (Fig. 4, 5, 6, 7, 8). This performance of DSR-GD is due in part to the choice of learning rate that helped speedily converge towards the optimal solution

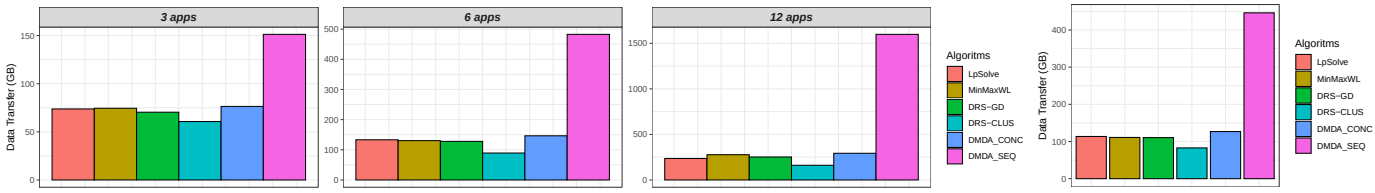

(a) Memory transfer obtained for the 3, 6, and 12 concurrent applications.

(b) Memory transfer summary

Fig. 9. Memory transfer of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ for all the affinities and all the hardware configurations (K40M, Quadro, and A100).

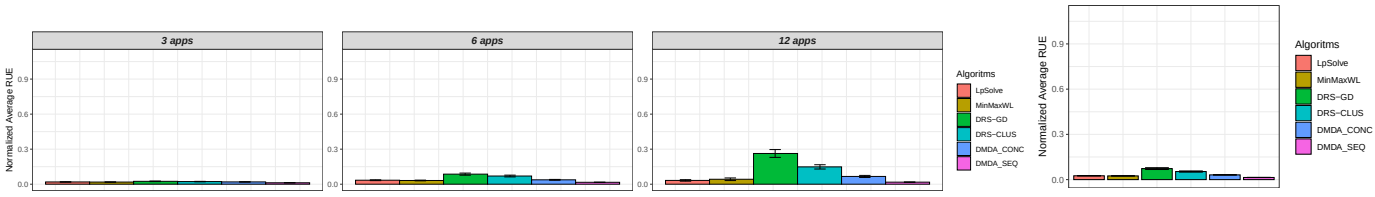

(a) Average RUE obtained for the 3, 6, and 12 concurrent applications.

(b) Average RUE summary

Fig. 10. Average RUE of LpSolve, MinMaxWL, DSR-GD, DSR-CLUS, DMDA_CONC against DMDA_SEQ for all the affinities and all the arch configurations (K40M, Quadro, and A100).

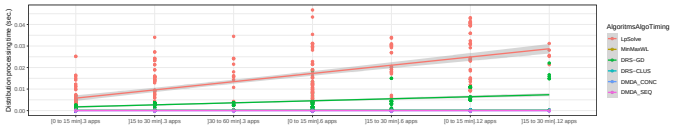

Fig. 11. Distribution processing time of LpSolve, MinMaxWL, DSR-GD, and DSR-CLUS according to the standard deviation percentage of GPU/CPU speedup between the applications, and the number of concurrent applications

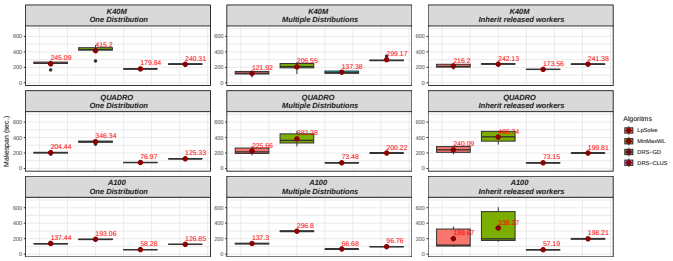

Fig. 12. **Analysis of distribution options** for LpSolve, MinMaxWL, DSR-GD, DSR-CLUS with 12 concurrent applications. *The averages makespan are displayed in red.*

no matter the number of applications and resources, and also due to the bartering technique employed (see Section III-A1). This is a promising property for a scalable system.

b) Distribution options: We carried out a study on the effectiveness of the re-distribution options ("Multiple Distributions", and "Inherit released workers") presented in Section III-A2 comparatively to the default one ("One Distribution") when combined with each of our four strategies (LpSolve, MinMaxWL, DSR-GD, DSR-CLUS). Fig. 12 presents the makespan obtained in each case, which reveals that the effectiveness of the re-distribution options ("Multiple Distributions", and "Inherit released workers") depends on the strategies and the configuration.

The combination DSR-GD/"Inherit released workers" always produces a gain for all the configurations. For the other cases, the combination strategy/option leads to a gain only for some configurations. We notify significant improvement in some cases, proving that there is hope for improving the results obtained above by using resource redistribution. However, it is imperative to do more investigations on configuration and strategy sensitivity to achieve this.

V. Conclusions

As computing resources are getting more complex and powerful, there is little doubt that we need methods to reduce the waste from the users' choices, bad application optimization, or heterogeneous workloads during executions. This is where the task-based model grants more opportunities by exposing a dynamic degree of parallelism with execution environments able to use this information in the most constructive and thus efficient way. To minimize the overall makespan and maximize resource utilization while executing multiple task-based applications, we introduce RSCHED, a two-level resource management framework that allows 1) dynamic resource distribution for concurrent execution of task-based applications, and 2) dedicated task scheduling for each application. We proposed strategies for resource distribution and implemented our proposal on the StarPU runtime system, proposing schedulers on which we rely for the second level. A new model of Gradient Descent has been proposed, among other strategies for resource distribution. We evaluated our

proposal using real applications based on the StarPU implementation of Cholesky factorization. RSCHED demonstrated the potential to speed up the overall makespan compared to consecutive execution with an average factor of 10x, and a factor of 5x when compared against the concurrent execution without resource distribution using DMDA. RSCHED also demonstrated the potential to increase the rate of resource utilization as the number of applications increases. Moreover, the decision time of our strategies and our initial study on resizing options indicate promising scalability.

In our future work, we will investigate continuous application arrivals and resizing options to have more adaptive solutions due to the variability and complexity of both the applications and the computing resources. That being said, we would like to consider different applications (instead of just Cholesky), explore multiple nodes, and improve RSCHED decisions by analyzing the structures of task graphs.

ACKNOWLEDGMENTS

This work is supported by the TExas, an Inria exploratory project. 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 (see http://www.plafrim.fr/).

REFERENCES

[1] N. P. Drakenberg, "Multi-objective processor-set selection for computational cluster-systems," in *Job Scheduling Strategies for Parallel Processing: 16th International Workshop, JSSPP 2012, Shanghai, China, May 25, 2012. Revised Selected Papers 16*. Springer, 2013, pp. 56–75.

[2] C. Augonnet, S. Thibault, R. Namyst, and P.-A. Wacrenier, "Starpu: a unified platform for task scheduling on heterogeneous multicore architectures," in *Euro-Par 2009 Parallel Processing: 15th International Euro-Par Conference, Delft, The Netherlands, August 25-28, 2009. Proceedings 15*. Springer, 2009, pp. 863–874.

[3] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner, and T. Takahashi, "Task-based fmm for multicore architectures," *SIAM Journal on Scientific Computing*, vol. 36, no. 1, pp. C66–C93, 2014.

[4] ——, "Task-based fmm for heterogeneous architectures," *Concurrency and Computation: Practice and Experience*, vol. 28, no. 9, pp. 2608–2629, 2016.

[5] J. M. Couteyen Carpaye, J. Roman, and P. Brenner, "Design and analysis of a task-based parallelization over a runtime system of an explicit finite-volume cfd code with adaptive time stepping," *Journal of Computational Science*, vol. 28, pp. 439–454, 2018.

[6] G. Bosilca, A. Bouteiller, A. Danalis, M. Faverge, T. Hérault, and J. J. Dongarra, "Parsec: Exploiting heterogeneity to enhance scalability," *Computing in Science & Engineering*, vol. 15, no. 6, pp. 36–45, 2013.

[7] A. Hugo, A. Guermouche, P.-A. Wacrenier, and R. Namyst, "Composing multiple starpu applications over heterogeneous machines: a supervised approach," *The International journal of high performance computing applications*, vol. 28, no. 3, pp. 285–300, 2014.

[8] C. Flint, L. Paillat, and B. Bramas, "Automated prioritizing heuristics for parallel task graph scheduling in heterogeneous computing," *PeerJ Computer Science*, vol. 8, p. e969, 2022.

[9] H. Tayeb, B. Bramas, A. Guermouche, and M. Faverge, "Multreeprio: Scheduling task-based applications for heterogeneous computing systems," in *COMPAS 2022-Conférence francophone d'informatique en Parallélisme, Architecture et Système*, 2022.

[10] J. E. Ndamlabin Mboula, V. C. Kamla, M. H. Hilman, and C. Tayou Djamegni, "Energy-efficient workflow scheduling based on workflow structures under deadline and budget constraints in the cloud," *arXiv preprint arXiv:2201.05429*, 2022.

[11] J. E. Ndamlabin Mboula, V. C. Kamla, and C. Tayou Djamegni, "Dynamic provisioning with structure inspired selection and limitation of vms based cost-time efficient workflow scheduling in the cloud," *Cluster Computing*, pp. 1–25, 2021.

[12] ——, "Cost-time trade-off efficient workflow scheduling in cloud," *Simulation Modelling Practice and Theory*, p. 102107, 2020.

[13] E. Calore and S. F. Schifano, "Porting a lattice boltzmann simulation to fpgas using ompss," in *Parallel Computing: Technology Trends*. IOS Press, 2020, pp. 701–710.

[14] J. V. Lima, G. Freytag, V. G. Pinto, C. Schepke, and P. O. Navaux, "A dynamic task-based d3q19 lattice-boltzmann method for heterogeneous architectures," in *2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*. IEEE, 2019, pp. 108–115.

[15] M. AbdulJabbar, R. Yokota, and D. Keyes, "Asynchronous execution of the fast multipole method using charm++," *arXiv preprint arXiv:1405.7487*, 2014.

[16] M. Pericàs, X. Martorell, and Y. Etsion, "Implementation of a hierarchical n-body simulator using the ompss programming model," in *Proceedings of the 1st Workshop on Irregular Applications: Architectures and Algorithms*, 2011, pp. 23–30.

[17] E. Agullo, L. Giraud, and S. Nakov, "Task-based sparse hybrid linear solver for distributed memory heterogeneous architectures," in *Euro-Par 2016: Parallel Processing Workshops: Euro-Par 2016 International Workshops, Grenoble, France, August 24-26, 2016, Revised Selected Papers 22*. Springer, 2017, pp. 83–95.

[18] X. Lacoste, M. Faverge, G. Bosilca, P. Ramet, and S. Thibault, "Taking advantage of hybrid systems for sparse direct solvers via task-based runtimes," in *2014 IEEE International Parallel & Distributed Processing Symposium Workshops*. IEEE, 2014, pp. 29–38.

[19] E. Agullo, A. Buttari, A. Guermouche, and F. Lopez, "Multifrontal qr factorization for multicore architectures over runtime systems," in *Euro-Par 2013 Parallel Processing: 19th International Conference, Aachen, Germany, August 26-30, 2013. Proceedings 19*. Springer, 2013, pp. 521–532.

[20] R. Carratalá-Sáez, M. Faverge, G. Pichon, G. Sylvand, and E. S. Quintana-Ortí, "Tiled algorithms for efficient task-parallel h-matrix solvers," in *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2020, pp. 757–766.

[21] R. Arias Mallo, "Particle-in-cell plasma simulation with ompss-2," Master's thesis, Universitat Politècnica de Catalunya, 2019.

[22] D. Sukkari, H. Ltaief, M. Faverge, and D. Keyes, "Asynchronous task-based polar decomposition on single node manycore architectures," *IEEE Transactions on parallel and distributed systems*, vol. 29, no. 2, pp. 312–323, 2017.

[23] L. Boillot, G. Bosilca, E. Agullo, and H. Calandra, "Task-based programming for seismic imaging: Preliminary results," in *2014 IEEE Intl Conf on High Performance Computing and Communications, 2014 IEEE 6th Intl Symp on Cyberspace Safety and Security, 2014 IEEE 11th Intl Conf on Embedded Software and Syst (HPCC, CSS, ICESS)*. IEEE, 2014, pp. 1259–1266.

[24] V. Martínez, D. Michéa, F. Dupros, O. Aumage, S. Thibault, H. Aochi, and P. O. Navaux, "Towards seismic wave modeling on heterogeneous many-core architectures using task-based runtime system," in *2015 27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. IEEE, 2015, pp. 1–8.

[25] S. Moustafa, W. Kirschenmann, F. Dupros, and H. Aochi, "Task-based programming on emerging parallel architectures for finite-differences seismic numerical kernel," in *Euro-Par 2018: Parallel Processing: 24th International Conference on Parallel and Distributed Computing, Turin, Italy, August 27-31, 2018, Proceedings 24*. Springer, 2018, pp. 764–777.

[26] B. Bramas, P. Helluy, L. Mendoza, and B. Weber, "Optimization of a discontinuous galerkin solver with opencl and starpu," *International Journal on Finite Volumes*, vol. 15, no. 1, pp. 1–19, 2020.

[27] P. Brucker and S. Knust, "Complexity results for scheduling problems," 2009.

[28] P. Baptiste, C. Le Pape, and W. Nuijten, *Constraint-based scheduling: applying constraint programming to scheduling problems*. Springer Science & Business Media, 2001, vol. 39.

[29] M. L. O. Salvana, S. Abdulah, H. Huang, H. Ltaief, Y. Sun, M. G. Genton, and D. E. Keyes, "High performance multivariate geospatial statistics on manycore systems," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 11, pp. 2719–2733, 2021.

[30] Q. Cao, Y. Pei, K. Akbudak, A. Mikhalev, G. Bosilca, H. Ltaief, D. Keyes, and J. Dongarra, "Extreme-scale task-based cholesky factorization toward climate and weather prediction applications," in *Proceedings of the Platform for Advanced Scientific Computing Conference*, 2020, pp. 1–11.

[31] K. Akbudak, H. Ltaief, A. Mikhalev, A. Charara, A. Esposito, and D. Keyes, "Exploiting data sparsity for large-scale matrix computations," in *European Conference on Parallel Processing*. Springer, 2018, pp. 721–734.

[32] J. M. C. Carpaye, J. Roman, and P. Brenner, "Design and analysis of a task-based parallelization over a runtime system of an explicit finite-volume cfd code with adaptive time stepping," *Journal of Computational Science*, vol. 28, pp. 439–454, 2018.

[33] H. Topcuoglu, S. Hariri, and M.-y. Wu, "Performance-effective and low-complexity task scheduling for heterogeneous computing," *IEEE transactions on parallel and distributed systems*, vol. 13, no. 3, pp. 260–274, 2002.

[34] Y. Xu, K. Li, J. Hu, and K. Li, "A genetic algorithm for task scheduling on heterogeneous computing systems using multiple priority queues," *Information Sciences*, vol. 270, pp. 255–287, 2014.

[35] K. R. Shetti, S. A. Fahmy, and T. Bretschneider, "Optimization of the heft algorithm for a cpu-gpu environment," in *2013 International conference on parallel and distributed computing, applications and technologies*. IEEE, 2013, pp. 212–218.

[36] H. J. Choi, D. O. Son, S. G. Kang, J. M. Kim, H.-H. Lee, and C. H. Kim, "An efficient scheduling scheme using estimated execution time for heterogeneous computing systems," *The Journal of Supercomputing*, vol. 65, pp. 886–902, 2013.

[37] M. A. Khan, "Scheduling for heterogeneous systems using constrained critical paths," *Parallel Computing*, vol. 38, no. 4-5, pp. 175–193, 2012.

[38] O. Beaumont, L.-C. Canon, L. Eyraud-Dubois, G. Lucarelli, L. Marchal, C. Mommessin, B. Simon, and D. Trystram, "Scheduling on two types of resources: a survey," *ACM Computing Surveys (CSUR)*, vol. 53, no. 3, pp. 1–36, 2020.

[39] A. K. Maurya and A. K. Tripathi, "On benchmarking task scheduling algorithms for heterogeneous computing systems," *The Journal of Supercomputing*, vol. 74, no. 7, pp. 3039–3070, 2018.

[40] B. Bramas, "Optimization and parallelization of the boundary element method for the wave equation in time domain," Ph.D. dissertation, Bordeaux, 2016.

[41] W. Fornaciari, G. Agosta, D. Atienza, C. Brandolese, L. Cammoun, L. Cremona, A. Cilardo, A. Farres, J. Flich, C. Hernandez *et al.*, "Reliable power and time-constraints-aware predictive management of heterogeneous exascale systems," in *Proceedings of the 18th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation*, 2018, pp. 187–194.

[42] G. Agosta, W. Fornaciari, D. Atienza, R. Canal, A. Cilardo, J. F. Cardo, C. H. Luz, M. Kulczewski, G. Massari, R. T. Gavilá *et al.*, "The recipe approach to challenges in deeply heterogeneous high performance systems," *Microprocessors and Microsystems*, vol. 77, p. 103185, 2020.

[43] N. Grenèche, T. Menouer, C. Cérin, and O. Richard, "A methodology to scale containerized hpc infrastructures in the cloud," in *European Conference on Parallel Processing*. Springer, 2022, pp. 203–217.

[44] C. Cérin, N. Grenèche, and T. Menouer, "Executing traditional hpc application code in cloud with containerized job schedulers," in *High Performance Computing in Clouds: Moving HPC Applications to a Scalable and Cost-Effective Environment*. Springer, 2023, pp. 75–97.

[45] I. Dagal, K. Tanriöven, A. Nayir, and B. Akın, "Adaptive stochastic gradient descent (sgd) for erratic datasets," *Future Generation Computer Systems*, vol. 166, p. 107682, 2025.

[46] S. Santra, J.-W. Hsieh, and C.-F. Lin, "Gradient descent effects on differential neural architecture search: A survey," *IEEE Access*, vol. 9, pp. 89 602–89 618, 2021.

[47] A. S. Berahas, L. Cao, K. Choromanski, and K. Scheinberg, "Linear interpolation gives better gradients than gaussian smoothing in derivative-free optimization," *arXiv preprint arXiv:1905.13043*, 2019.

[48] S. C. Chapra, *Numerical methods for engineers*. Mcgraw-hill, 2010.

[49] H. Arabnejad and J. G. Barbosa, "List scheduling algorithm for heterogeneous systems by an optimistic cost table," *IEEE transactions on parallel and distributed systems*, vol. 25, no. 3, pp. 682–694, 2013.

APPENDIX

Algorithm 1: Ideal makespan Algorithm

```
function ideal_makespan(graph_info G, double coef_par_eff)
    nb_workers = G.#GPUS + G.#CPUS;
    if (G.#GPUS == 0.0) || (G.#CPUS == 0.0) then
        if (G.GPUS_PR == 0.0) then
            m = G.GPU_W / G.#GPUS;
        if (G.GPUS_PR == 0.0) then
            m = G.CPU_W / G.#CPUS;
        return m × coef_par_eff (nb_workers−1);
    Compute tgpu_minM and tcpu_minM using equations 5 and 6;
    if (G.GPU_W == G.GPU_PW) || (G.CPU_W==G.CPU_PW) then
        return MAX(tgpu_minM, tcpu_minM);
    gpu_rem_wl = compute the remaining GPU workload;
    cpu_rem_wl = compute the remaining CPU workload;
    if (gpu_rem_wl == 0.0) || (cpu_rem_wl == 0.0) then
        return MAX(tgpu_minM, tcpu_minM) + gpu_rem_wl / G.#GPUS
          + cpu_rem_wl / G.#CPUS;
    Compute ideal_makespan using equation 7;
    return ideal_makespan;
```

Algorithm 2: DSR-GD (Dedicated plus Shared Resource with Gradient Descent)

```
//-->Dedicated workers per app;
rem_gpus = nb_gpus;
rem_cpus = nb_cpus;
foreach graph G in graphs do
    gpus_dedicated = floor(G.GPU_W / SUM(GPU_W)) × nb_gpus;
    cpus_dedicated = floor(G.CPU_W / SUM(CPU_W)) × nb_cpus;
    while (gpus_dedicated) do
        Assign the (rem_gpus)^−th GPU to graph G;
        gpus_dedicated − −;
        rem_gpus − −;
    while (cpus_dedicated) do
        Assign the (rem_cpus)^−th CPU to graph G;
        cpus_dedicated − −;
        rem_cpus − −;
if (Stdd_speedup > SPEEDUP_STDD_LIMIT) || (Stdd_idealm >
  MAKESPAN_STDD_LIMIT)) then
    bartering();
//-->Shared workers between apps Using Gradient Descent;
Configure GD axis (X, Y) using rem_cpus and rem_gpus;
foreach axis in X, Y, Z do
    while (No convergence) do
        Use the best indexes from previous axes;
        Fix the value of the following axis;
        Search the best index in the axis using Gradient Descent with
          Min-Max ideal_makespan as objective function;
        keep track of the best solution;
```

Algorithm 3: MinMaxWL (Min-Max Workload balancing)

```
//-->Ensure each graph has at least one worker;
//----->Graphs with pure workload;
foreach graph G in graphs do
    if (G.GPU_PW != 0.0) then
        if (Remaining GPUs) then
            Assign one GPU to G;
        else
            Share one GPU with the under-loaded pure gpu graph, with
              backpropagation;
    if (G.CPU_PW != 0.0) then
        if (Remaining CPUs) then
            Assign one CPU to G;
        else
            Share one CPU with the under-loaded pure cpu graph, with
              backpropagation;
//----->Graphs without pure workload;
foreach graph G in graphs do
    if (G.CPU_PW == 0.0 && G.CPU_PW == 0.0) then
        if (G.CPU_W > G.GPU_W) then
            // GPU is faster;
            if (Remaining GPUs) then
                Assign one GPU to G;
            else
                Share one GPU with the under-loaded graph, with
                  backpropagation;
        else
            // CPU is faster;
            if (Remaining CPUs) then
                Assign one CPU to G;
            else
                Share one CPU with the under-loaded graph, with
                  backpropagation;
//-->Dist. Remaining Workers: Load-balancing using Min-Max;
while (Remaining GPUs Workers) do
    Assign one GPU to graph leading to Min-Max ideal_makespan;
while (Remaining CPUs Workers) do
    Assign one CPU to graph leading to Min-Max ideal_makespan;
```