




Contents lists available at ScienceDirect

Journal of Parallel and Distributed Computing

journal homepage: www.elsevier.com/locate/jpdc

Leveraging Multi-Instance GPUs through moldable task scheduling

Jorge Villarrubia^{ID,*}, Luis Costero^{ID}, Francisco D. Igual^{ID}, Katalin Olcoz

Universidad Complutense de Madrid, Departamento de Arquitectura de Computadores y Automática, Madrid, 28040, Spain

ARTICLE INFO

Keywords:

Multi-Instance GPU (MIG)
Moldable Resource Management
Task Scheduling

ABSTRACT

NVIDIA MIG (Multi-Instance GPU) allows partitioning a physical GPU into multiple logical instances with fully-isolated resources, which can be dynamically reconfigured. This work highlights the untapped potential of MIG through moldable task scheduling with dynamic reconfigurations. Specifically, we propose a makespan minimization problem for multi-task execution under MIG constraints. Our profiling shows that assuming monotonicity in task work with respect to resources is not viable, as is usual in multicore scheduling. Relying on a state-of-the-art proposal that does not require such an assumption, we present FAR, a 3-phase algorithm to solve the problem. Phase 1 of FAR builds on a classical task moldability method, phase 2 combines Longest Processing Time First and List Scheduling with a novel repartitioning tree heuristic tailored to MIG constraints, and phase 3 employs local search via task moves and swaps. FAR schedules tasks in batches offline, concatenating their schedules on the fly in an improved way that favors resource reuse. Excluding reconfiguration costs, the List Scheduling proof shows an approximation factor of $7/4$ on the NVIDIA A30 model. We adapt the technique to the particular constraints of an NVIDIA A100/H100 to obtain an approximation factor of 2. Including the reconfiguration cost, our real-world experiments reveal a makespan with respect to the optimum no worse than $1.22\times$ for a well-known suite of benchmarks, and $1.10\times$ for synthetic inputs inspired by real kernels. We obtain good experimental results for each batch of tasks, but also in the concatenation of batches, with large improvements over the state-of-the-art and proposals without GPU reconfiguration. Moreover, we show that the proposed heuristics allow a correct adaptation to tasks of very different characteristics. Beyond the specific algorithm, the paper demonstrates the research potential of the MIG technology and suggests useful metrics, workload characterizations and evaluation techniques for future work in this field.

1. Introduction

1.1. Context and motivation

Graphics Processing Units (GPUs) are nowadays established as the elemental computing device for accelerating massively parallel computations in fields such as Artificial Intelligence (AI) and High Performance Computing (HPC). Hand in hand with the emergence of much more resource-demanding tasks, manufacturers are offering increasingly powerful GPUs featuring tens of thousands of computing cores and a massive memory bandwidth. These resources, however, can be difficult to be fully utilized in applications that exhibit lower levels of parallelism or bandwidth necessities, resulting in resource underutilization. Co-scheduling independent applications is a natural way to regain full occupation of the available computing resources in a single GPU. Still, it poses important challenges to avoid or control resource contention and

to adapt the assigned per-application resources to attain fair scheduling policies.

Manufacturers offer different solutions to address the problem of resource isolation in multi-tenant scenarios on a single GPU. For example, the Multi-Processing Service (MPS) from NVIDIA [1] allows multiple processes to use the same GPU simultaneously by provisioning the stream multiprocessors (SMs), but sharing memory bandwidth. This is not a good solution for the cloud, as the quality of service a client receives (e.g. processing speed) cannot be guaranteed in advance since it depends on the resource usage of the tasks running simultaneously on the GPU. Some of the latest GPU models for data centers, including the NVIDIA Ampere, Hopper and Blackwell families, integrate the Multi-Instance GPU (MIG) technology,¹ which allows the GPU to be logically and physically split into disjoint instances that have their own fully isolated hardware resources, including memory bandwidth, cache and cores, among others.

* Corresponding author.

E-mail address: jorvil01@ucm.es (J. Villarrubia).¹ See <https://docs.nvidia.com/datacenter/tesla/mig-user-guide/index.html>.<https://doi.org/10.1016/j.jpdc.2025.105128>

Received 18 July 2024; Received in revised form 14 April 2025; Accepted 2 June 2025

Available online 6 June 2025

0743-7315/© 2025 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

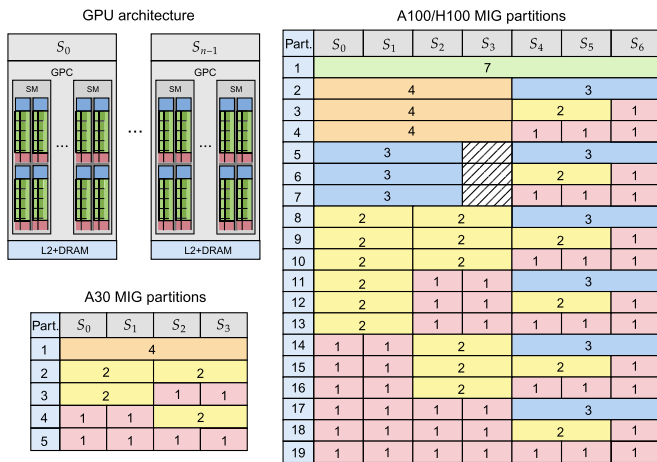


Fig. 1. Slice per instance distributions for the NVIDIA A30 and A100/H100.

For Cloud Service Providers (CSPs), MIG solves resource contention in multi-tenant scenarios. Typically deployed in a *static* and *task-oblivious* fashion (that is, the partitioning scheme does not vary frequently and it does not depend on the specific features of the task(s) to execute), the use of MIG ensures that one client cannot impact the work of others, providing robust Service Level Agreements (SLA) on a shared GPU without needing to deploy tasks across multiple GPUs, and at the same time improving the effective resource utilization. However, the reconfigurable nature of MIG-capable GPUs enables *dynamic* partitioning and scheduling schemes for tasks with different scalability properties. As of today, the joint exploitation of task scheduling and MIG reconfiguration is a topic hardly covered in the literature.

1.2. MIG concepts and restrictions

The minimum units in MIG are called *slices*, comprising one block of computational resources called GPCs (GPU Processing Clusters), along with a portion of L2 cache and DRAM. Several consecutive slices can be grouped in an *instance* on which to execute a task. Furthermore, due to physical limitations, it is not possible to form instances with any sequence of consecutive slices. Finally, a set of instances that divide the GPU is called a *partition*.

Fig. 1 shows all the possible partitions of the NVIDIA A30 (left), consisting of 4 slices with 6 GB of memory assigned to each one. Note that it is not possible to form instances with 3 slices, nor to create an instance with the two middle ones: $\{S_1, S_2\}$. The NVIDIA A100 and H100 extend this strategy to 7 slices, with 5 and 10 GB of memory in each slice respectively, presenting a total of 19 possible partitions, also shown in Fig. 1 (right). Note that partitions 5, 6 and 7 do not use S_3 . In these partitions, the memory allocated to S_3 can be assigned to the instance composed by $\{S_0, S_1, S_2\}$, imitating the same instance formed by $\{S_4, S_5, S_6\}$ (as S_6 comprises more memory than the rest of slices). Except for a slightly lower energy consumption, it never seems more interesting to use $\{S_0, S_1, S_2\}$ disabling S_3 than directly using $\{S_0, S_1, S_2, S_3\}$; however, since NVIDIA offers it, we will also consider that instance.

Finally, MIG allows to reconfigure partitions dynamically, so that even if a task assigned to an instance has not completed its execution, other free instances can be repartitioned/merged. This reconfiguration philosophy means that (1) instances can be organized hierarchically (they are joined or split to form other instances), and (2) valid partitions are precisely the possible combinations of disjoint instances. These two properties will guarantee that our proposal works correctly. To make the explanations more concrete and easier to follow, we will base them on the GPUs previously described (NVIDIA A30, A100 and H100), but they can be extrapolated to any future generations as long as they maintain these properties, which is expected. In fact, NVIDIA has just released




Fig. 2. Speedup of Rodinia's *Gaussian* kernel for instance sizes on A100.

two models of the Blackwell family (B100 and B200) with exactly the same MIG restrictions as the A100/H100 models: up to 7 slices with the same valid partitions. In addition to this, task migration between instances is not allowed, nor any kind of task preemption.

1.3. Motivation of MIG usage for task scheduling

While the amount of available resources in modern GPUs has increased drastically, not all applications can take advantage of these resources, as their degree of parallelism can be insufficient to use all of them, or the communication overhead can limit their speedup. In these situations the use of MIG can improve the resource usage of the GPU without impact in performance.

As an illustrative example, let us consider the *Gaussian* kernel within the Rodinia GPU benchmark [2], whose performance for an increasing number of slices within an A100 GPU is reported in Fig. 2. Each color bar represents a different input size, ranging from 0.32 s up to 20 s for 1 slice (blue and orange bars respectively). The results reveal that, for this application and both input sizes, it stops scaling beyond 3 slices. In this situation, co-scheduling independent applications in the remaining slices would increase the effective GPU occupation, without degrading individual performance thanks to the isolation capabilities of MIG.

In addition, in some scenarios, it is common to run the same kernel multiple times, so that all runs can be seen as a longer task for MIG utilization. If the task times are long, the use of MIG makes even more sense as reconfiguration times become negligible, and it is possible that during their executions many new pending jobs accumulate. This is specially interesting in datacenters or clusters where typically large tasks are executed. In these scenarios, where many new jobs accumulate in request queues while other long tasks are running, it would be possible to schedule them offline to optimize their joint execution through MIG. As of today, MIG is supported in job management systems like SLURM or HTCondor, and in workload orchestrators like Kubernetes; their scheduling policies, however, do not include coupled strategies for dynamic MIG reconfiguration and job scheduling.

1.4. Contributions

In this paper, we address the scheduling of independent tasks (applications) on a MIG-capable GPU, determining the instance and the time at which their execution starts. We focus our proposal towards a scenario in which we can form successive batches of pending tasks while executing other long tasks (e.g. taking a batch of tasks every periodic time interval), as is common in batch job management systems and orchestrators. For each batch, we consider an *offline* scheduling problem (its tasks are available from the start of batch execution), with the objective of minimizing the makespan (last completion time of a task in the batch). This scheduling is *moldable* , since the amount of resources that each task receives is not given, but decided by the scheduler upon task start, and it is restricted to always use feasible partitions according to the rules imposed by MIG. Between batches, we perform an *online* scheduling (batches are known on the fly), by concatenating their schedules in an *intelligent* way. Our proposal is suitable for any GPU where MIG still satisfies the properties mentioned in Section 1.2: (1) instances

can be hierarchically organized, and (2) valid partitions are precisely combinations of disjoint instances. Because of the dynamic and instance-independent reconfiguration philosophy, MIG is expected to preserve these conditions.

We assume that the execution time of a task on each instance size is known,² as tasks can be previously profiled to predict their execution time (for example, on an auxiliary node) [3,4], or we operate in a private environment with tasks that have known behaviors and are easily predictable. In fact, we could use the MISO predictor presented in [5], which precisely predicts execution times for MIG with different instance sizes using profiling-based machine learning.

The main contributions of the paper, together with the section of the manuscript in which they are covered, are:

- We formulate a scheduling problem for MIG that, to the best of our knowledge, is new. We also review similarities and differences with other problems addressed in the literature, showing experimentally that it is not acceptable to assume monotony in the work of the tasks, as most of the moldable scheduling proposals do, although it is reasonable to assume that their times are monotonic. This contribution is covered in Section 2.
- We propose FAR (*Family of Allocations and Repartitioning*), a three-phase algorithm for solving the offline problem for each batch of tasks. The first phase is a direct application of the molding method presented in [6]. The second phase combines Longest Processing Time First and List Scheduling [7] with a new repartitioning tree heuristic derived from MIG constraints. Finally, the third phase refines the result through task moves and swaps. The details of the algorithm are given in Section 3. A C++ implementation for MIG-capable GPUs is offered to the community in [8].
- We propose a strategy to concatenate FAR schedules for multiple task batches to favor resource utilization. This strategy is developed in Section 4.
- We applied the List Scheduling proof to derive an approximation factor of 7/4 for the A30 model, without considering reconfiguration costs. We adapted the proof to the casuistry of the A100/H100 models proving an approximation factor of 2 for them. This proof is developed in Section 5.
- We empirically test our proposal, verifying that it provides good solutions, with differences to the optimum of less than 10%, and does not require large task batches to achieve it. In particular, we show large improvements over other approaches: between 55% and 230% versus a scheduler without reconfiguration and between 19% and 114% versus the state-of-the-art in MIG scheduling [5]. We use real execution times from Rodinia, as well as synthetic times that mimic real task behaviors and allow adjusting different workloads parametrically. The experimental evaluation corresponds to Section 6.

2. Scheduling problem

2.1. Reconfiguration times

Reconfiguration times in MIG can greatly influence on the final execution time depending on their number and cost. The reconfiguration of a MIG partition is performed by the isolated destruction and creation of instances; for example, to transform partition 2-2 (partition number 2 for the NVIDIA A30 in Fig. 1) into partition 1-1-2 (partition number 4 for the NVIDIA A30 in Fig. 1) it is enough to destroy $\{S_0, S_1\}$, and create $\{S_0\}$ and $\{S_1\}$, without modifying $\{S_2, S_3\}$. Note that this way of

² The time only depends on the number of slices in the instance, since the slices are identical. The exception is S_6 in the A100, which offers more memory to instances of size 3, but we explained that they also get it in other cases.

Table 1

Average time (in seconds) for the creation/destruction of an instance of each possible size on each MIG-capable GPU.

Instance Size	A30		A100		H100	
	Create	Destroy	Create	Destroy	Create	Destroy
1	0.11	0.10	0.16	0.20	0.16	0.21
2	0.12	0.10	0.17	0.20	0.21	0.23
3	–	–	0.20	0.21	0.33	0.25
4	0.13	0.10	0.21	0.21	0.38	0.26
7	–	–	0.24	0.22	0.42	0.26

reconfiguring by creating/destroying instances independently supports the properties that our proposal will use: (1) hierarchical organization of instances, and (2) any feasible combination of instances is part of a valid partition.

It would be expected that creating/destroying multiple independent instances could be parallelized, but our experience is that the NVIDIA drivers currently sequentialize this process. On the three currently available MIG-capable GPUs (A30, A100 and H100) we have obtained virtually the same reconfiguration time between any pair of partitions by running all builds/destructions sequentially versus using different threads to perform in parallel those that are independent.

Therefore, the reconfiguration overhead is determined by the sum of the instance creation/destruction times which only depend on their size, as we have also experienced (e.g., it took practically the same time to build/destroy $\{S_0, S_1\}$ as $\{S_2, S_3\}$ because both have size 2). Table 1 reports the times of each operation on the MIG capable-GPUs for the different instance sizes. Deviations are not reported because in all cases they are of the order of milliseconds, i.e., none of the decimals shown change. We can see that the times are small in all cases, and therefore the reconfigurations should have little impact relative to the duration of the tasks unless they are very short. Reconfigurations were carried out using the NVML library,³ which does not introduce the overhead observed when using the `nvidia-smi` tool. The times in Table 1 are practically stable between different executions of the same operation, and we will use them for the problem specification and in our scheduling algorithm.

2.2. Specification and notation

Given a set of independent tasks to be executed on a MIG-capable GPU denoted by G (currently, $G \in \{A30, A100, H100\}$), we assume that tasks are grouped in batches periodically.⁴ For each batch, we pose the following problem:

Input:

A batch of tasks $\{T_1, \dots, T_n\}$ to be executed with MIG on G , and their execution times for its instance sizes:

$$t_i : C_G \rightarrow \mathbb{R}^+ \quad \forall i \in \{1, \dots, n\},$$

$$\text{where } C_G = \{|I| : I \in \text{MIG-Instances}_G\}.$$

In our case, $C_{A30} = \{1, 2, 4\}$ and $C_{A100} = C_{H100} = \{1, 2, 3, 4, 7\}$.⁵ The destruction/creation times of each size are also known:

$$t_{destroy}, t_{create} : C_G \rightarrow \mathbb{R}^+.$$

³ https://docs.nvidia.com/deploy/nvml-api/group_nvmlMultiInstanceGPU.html.

⁴ Long tasks favor a correct batching according to these strategies, but nothing prevents to use some online heuristic for cases with few accumulated tasks.

⁵ The number of slices and instance sizes C_G are constant for a particular GPU. However, we include them in the complexity analyses, emphasizing that they do not matter in asymptotic terms and currently have little impact on cost.

Goal:

Assign to each task T_i an instance $I_i \in \text{MIG-Instances}_G$ and a begin time $b_i \geq 0$, trying to minimize the makespan (last finish time f_i) without preemption:

$$\min_{(I_i, b_i)_{i=1}^n} \left(\max_{1 \leq i \leq n} f_i \right), \quad \text{where } f_i = b_i + t_i(|I_i|).$$

Constraints:

1. Each resource can only be in use by one task, i.e., two tasks with slices in common do not run at the same time:

$$[b_i, f_i) \cap [b_j, f_j) = \emptyset, \quad \text{if } I_i \cap I_j \neq \emptyset, \quad 1 \leq i \neq j \leq n.$$

2. Let $P_t = \{I_i : t \in [b_i, f_i)\}$ be the set of instances used at time t , and MIG-Part_G the set of valid partitions for G (those shown in Fig. 1). At any time, the running instances are part of a valid MIG partition:

$$\forall t \in \mathbb{R}^+ \exists P \in \text{MIG-Part}_G, \quad \text{with } P_t \subseteq P.$$

3. There is enough time for reconfiguration, i.e., between two instants $t < t'$ there is time to sequentially destroy the instances that have disappeared ($I \in P_t, I \notin P_{t'}$) and create the instances that have appeared ($I' \notin P_t, I' \in P_{t'}$):

$$\forall t < t' \in \mathbb{R}^+, \quad t + \sum_{I \in P_t, I \notin P_{t'}} t_{\text{destroy}}(|I|) + \sum_{I' \notin P_t, I' \in P_{t'}} t_{\text{create}}(|I'|) \leq t'.$$

The schedule of a batch is planned after that of the previous batch, but it is possible to overlap part of their executions while complying the constraints. Section 4 will show how to optimize this concatenation of batch schedules to maximize such overlap.

2.3. Related work

Parallel task scheduling has been extensively studied in the literature for different levels of flexibility [9,10]. In the *rigid* form, the amount of resources allocated to each task is predefined and given as input [11,12]. In the *modal* form, the scheduler statically determines the amount of resources for each task, which remains fixed during execution [13,14]. That is what happens for MIG. Finally, in the *malleable* form, the scheduler dynamically adjusts the number of resource allocated, varying them during task execution [15–17]. Additionally, the problem can be *offline* if a set of tasks is available from the beginning [18,19,13], or *online* if tasks are received and scheduled individually on the fly [20–22]. Although the online form is more flexible and suitable for real-time scenarios, it is poorly informed and its results are much worse than the offline equivalent [23]. In our scenario, where tasks are long enough to form batches, offline scheduling is preferable.

Several papers have explored scheduling problems involving setup times and costs, with theoretical contributions relevant to dynamic and even stochastic environments [24,25], but most are practical applications in concrete scenarios [26,27]. Although these approaches offer some flexibility in handling diverse environments, their applicability to the rigid partitioning constraints of MIG is not yet possible. Some authors consider also a bi-criteria objective problem, where another metric is tried to be optimized together with the makespan [28,29]. Finally, many papers deal with the problem under task precedence constraints given as a graph [30,17,31].

As discussed above, MIG presents an opportunity via *modal scheduling* as the number of slices for a task can be chosen at the beginning but cannot be changed until it is finished. In this case, it is usual to approach the problem in two phases: first, by allocating the number of resources that will be available for each task, and second, by applying an algorithm for the resulting problem [20]. Our proposal will work like this, and will treat the rigid problem as the Strip Packing

problem [32], since they are equivalent when the allocated resources are imposed to be consecutive (this happens with MIG). Due to its strong NP-hardness [33], this problem has been treated with bio-inspired algorithms [34,35], relaxed linear programming models [36,37], machine-learning approaches [38] and, mostly, using heuristics [39,40]. Among them, we adapt the list scheduling heuristic with LPT (Long Processing Time first) [7] for MIG, due to its simplicity and suitability for our problem.

Nevertheless, to the best of our knowledge, there is no version of the problem with constraints such as those imposed by MIG: the set of valid partitions is not the total set of partitions. Thus, using previous approaches to our problem could create partitions such as 2-4-1 in the A100/H100, when the valid one would be 4-2-1 (see Fig. 1). Although in such a case one may think of simply reordering the instances to make the partitioning valid, this breaks the inter-task placement intended by the scheduler and tends to produce very poor results. Moreover, none of the existing proposals allows to easily introduce MIG constraints to astutely schedule tasks.

In practice, most proposals focus on multicores and assume a task monotony hypothesis that, as we will see shortly, is not acceptable on MIG-capable GPUs. The problem is easier to solve under the monotony assumption since it becomes NP-hard in the weak sense, and polynomial algorithms emerge to arbitrarily approximate optimal scheduling [41]. In fact, for task execution times it is usual to assume a more restrictive scaling model, such as Roofline, the communication model or Amdahl's model, which in particular satisfy the monotony hypothesis [30]. In contrast, approaches for non-monotonic tasks are forced to relax the problem in another way to approximate arbitrarily well the optimal solution: by assuming a constant number of processors [42,43], or a polynomial number of processors with respect to the number of tasks [44]. Although these relaxations are adequate in our case (the number of slices is constant and, in fact, small), the cost hides excessively large multiplicative constants and degrees that limit their practical use. In this sense, it is preferable to use much more efficient heuristic proposals, even if they do not guarantee such a good approximation [45,6]. Among the few proposals in this line we highlight the *family of allotments* presented by Turek et al. [6], which is oriented to independent tasks and will be used for phase 1 of our scheduler.

Finally, there are not many proposals of scheduling with MIG, probably due to its novelty. The aim of [46] is to minimize the number of GPUs satisfying the QoS requirements of the tasks. Instead, MISO optimizer presented in [5] iteratively selects the partition that maximizes the sum of the speedups of tasks in FIFO order. Although we previously highlighted the time predictor of that paper, we believe that their scheduling optimizer has room for improvement. We will experimentally compare our scheduler against it in Section 6. Very recently, some proposals aiming at improving the use of MIG in multi-tenant scenarios have been published. Among them, ElasticBatch [47] employs a machine learning based pipeline to reduce the number of GPU instances and increase their utilization. Moreover, Jormungandr [48] improves the use of MIG for the particular case of deep neural networks (DNNs) with a containerized implementation using Kubernetes. Finally, MIGER [49] combines the use of MIG and MPS to improve the performance of offline task submission with QoS requirements.

2.4. Non-monotonic tasks

Most proposals for modal scheduling include a task monotony hypothesis consisting of the following two points:

1. A task takes the same or less time to execute when it has more resources:

$$t_i(s) \geq t_i(s'), \quad \forall 1 \leq i \leq n, \quad s < s'.$$

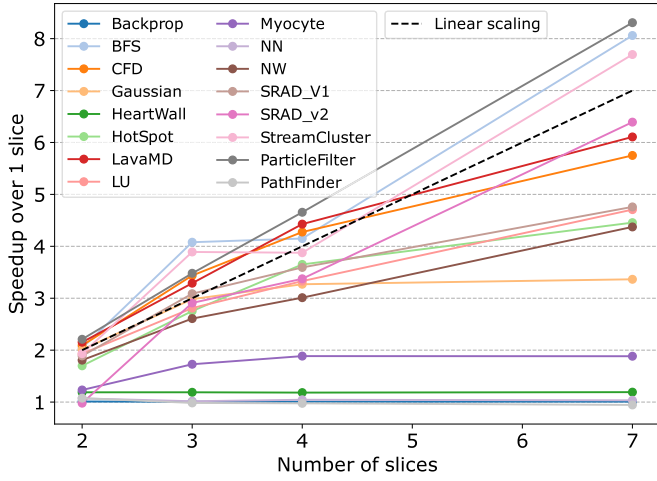


Fig. 3. MIG speedup of selected Rodinia benchmarks on a NVIDIA A100.

- The work of a task (product of the amount of resources by its execution time) does not decrease with more resources. This means that no task scales super-linearly:

$$s \cdot t_i(s) \leq s' \cdot t_i(s') \equiv \text{speedup}_i(s) / s \geq \text{speedup}_i(s') / s',$$

$$\forall 1 \leq i \leq n, s < s',$$

where $\text{speedup}_i(m) = t_i(1) / t_i(m)$.

The second point, which is of capital importance to take advantage of the hypothesis in algorithmic terms, assumes linear scaling as ideal and conceives sub-linear scaling by overhead in communications. However, sometimes it is possible to observe super-linear scaling due to memory hierarchy effects, the increasing available bandwidth or scheduling anomalies such as those described in [7].

In our experience with MIG, it is relatively common for applications to superscale. Fig. 3 reports the scaling of a subset of the Rodinia kernels on an A100 according to the number of slices, many of which widely exceed linear scaling. Of course, there are also applications that scale approximately linearly or clearly sub-linearly from a certain instance size. Superscaling is probably more prevalent in MIG than in multicore processors as the memory and the bandwidth of the instances are isolated rather than shared, which favors great improvements for memory-bound kernels (e.g. *BFS* and *StreamCluster* superscale up to 7 slices but hardly improve from 3 to 4, as both sizes offer the same memory bandwidth).

From what we have discussed, to solve the problem, we will not assume the second point of the monotony hypothesis. However, we will assume the first point since we have rarely observed tasks taking longer when using more slices, and in the few cases that have occurred, the increase has been less than 2%.

2.5. Instance isolation

As discussed, instance isolation is the distinguishing feature of MIG compared to other techniques for GPU task co-execution. Although NVIDIA claims that the isolation is total, an experimental verification is pertinent, as it is fundamental for the scheduling problem posed. In particular, it is necessary to verify that the execution of a task on an instance is independent of the tasks running simultaneously on other instances.

To achieve this, we run each of the 16 Rodinia kernels mentioned earlier on an instance of every possible size. First, the kernels are executed without any activity on the remaining slices, and then together with one or two randomly selected kernels. For example, we start by measuring the execution time of the *Backprop* kernel on a size-4 instance S_0, S_1, S_2, S_3 while keeping slices S_4, S_5 and S_6 idle. Next, we repeat

Table 2

Maximum percentage (%) difference in execution time with MIG when running a kernel alone vs. co-running with other tasks on NVIDIA A100.

Kernel	Instance size			
	1	2	3	4
Backprop	1.13	1.11	0.54	0.35
Gaussian	0.62	0.51	0.82	0.81
HotSpot	1.48	1.96	1.72	0.83
LavaMD	0.34	0.41	0.23	0.23

the measurement but this time running a randomly chosen kernel simultaneously on the instance $\{S_4, S_5, S_6\}$. Then, *Backprop* is executed again on the size-4 instance, but now with two additional random kernels running on instances of sizes 2 and 1, assigned to $\{S_4, S_5\}$ and $\{S_6\}$, respectively. Finally, we repeat this for two random kernels, each using instances $\{S_4\}$ and $\{S_5\}$.

From this experiment, we calculate the deviation in *Backprop*'s execution times compared to its standalone execution on a size-4 instance. This process is repeated for *Backprop* across instances of sizes 3, 2, and 1. Subsequently, the entire procedure is performed for the other 15 kernels. We limit co-execution to 1 or 2 additional kernels and select them randomly to avoid excessive combinatorial complexity while ensuring the scenarios are diverse and representative.

Table 2 shows the maximum percentage difference in execution time of some cores when running together with other applications compared to running them alone, for each instance size. For brevity, we have chosen 4 kernels instead of reporting all 16. Among them, we observe at most a 1.96% maximum difference in *HotSpot* runtime when using 2 slices, which is in fact also the maximum value if we would report all 16 kernels. In all cases we see very small time differences, which are probably more related to data caching between runs than to lack of isolation. These results demonstrate that the isolation provided by MIG is really effective.

3. FAR: family of allocations and repartitioning

To solve the aforementioned problem, we propose FAR, a three-phase algorithm for the MIG scheduling problem of a batch of tasks. The first phase addresses moldability by generating a family of allocations with the number of slices each task will use (Section 3.1). In the second phase, instance repartitioning is used to solve the rigid problem associated with each allocation, determining the start time and specific slices to execute the tasks (Section 3.2). Finally, the third phase refines the schedule by moving tasks from slices that reach the makespan to other slices (Section 3.3). An implementation of the algorithm is provided in [8], and Section 3.4 briefly explains how it performs FAR output execution on a MIG-capable GPU.

3.1. Phase 1: generating a family of allocations

To decide how many slices to assign to each task we adapt the method proposed by [6], which generates a family of candidate allocations, assuming the first point of monotony explained in Section 2.4, but not the second.

The idea is to minimize the *work* of the tasks, defined as the product of slices by the execution time, but not to make tasks too long by receiving few slices. As the work is the area that the task will occupy in the timetable, reducing it favors a smaller makespan. However, very long tasks can produce a large makespan by themselves. Therefore, the first allocation selects the minimum number of slices that minimizes the work of each task. Denoting with a_i^1 the number of slices assigned to the i -th task in the first allocation, we have that:

First allocation:

$$a_i^1 = \arg \min_{s \in C_G} (s \cdot t_i(s)), \quad 1 \leq i \leq n.$$

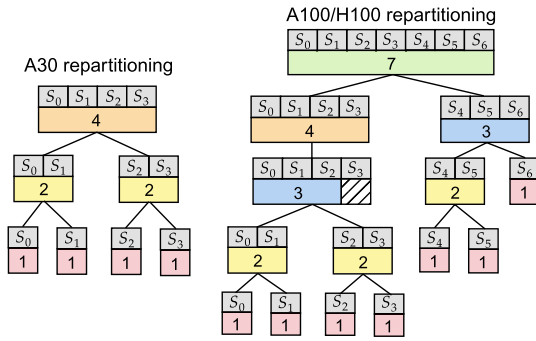


Fig. 4. Trees with instance repartitioning order for MIG-capable GPUs.

The rest are built from the previous one by increasing the slices of the longest task while possible. That is, all tasks in allocation a^{k+1} have the same assignment as in a^k , except for the longest one, which is increased to the next possible number of slices (if it cannot be increased, a^k is the last allocation in the family):

Allocation (k+1)-th from the k-th:

$$a_i^{k+1} = \begin{cases} \arg \min_{s \in C_G, s > a_i^k} (s \cdot t_i(s)) & \text{if } t_i(a_i^k) = \max_{1 \leq j \leq n} t_j(a_j^k), \\ a_i^k & \text{otherwise.} \end{cases}$$

The amount of allocations in the family is $\mathcal{O}(|C_G| \cdot n)$. As C_G is fixed for a GPU and has a small cardinal (3 for the A30 and 5 for the A100/H100), the number of allocations is de facto $\mathcal{O}(n)$.

Finally, the following theorem, proved in [6], bounds the error of some allocation of the family, with respect to the optimal solution of the moldable problem, using an algorithm A that solves the rigid problem. Thus, by applying A to all the allocations, and choosing the solution with the lowest makespan, we have a scheduler with guaranteed maximum error for the moldable problem. The requirement is that the makespan when solving the rigid problem of a fixed allocation is bounded by a combination of the total task area ($\sum_{i=1}^n s_i \cdot t_i(s_i)$) and the duration of the longest task (h_{max}). In successive phases, such an algorithm will be presented, and its makespan will be bounded in Section 5 according to that combination, using this theorem to derive an approximation factor for the moldable version.

Theorem 1. Let A be an algorithm that solves the rigid (non-moldable) scheduling problem, producing a makespan ω bounded as follows:

$$\omega \leq \max \left\{ \alpha_1 \sum_{i=1}^n s_i \cdot t_i(s_i) + \beta_1 h_{max}, \dots, \alpha_m \sum_{i=1}^n s_i \cdot t_i(s_i) + \beta_m h_{max} \right\},$$

where s_i is the number of slices given for the i -th task, $h_{max} = \max_i t_i(s_i)$, and α_i, β_i are coefficients. Then, there is at least one allocation a^k of the family for which using A yields a makespan ω^k bounded as follows:

$$\omega^k \leq \max \{ \alpha_1 \cdot \#slices_G + \beta_1, \dots, \alpha_m \cdot \#slices_G + \beta_m \} \cdot \omega^*,$$

where ω^* is the optimal makespan for the moldable problem.

3.2. Phase 2: instance scheduling by repartitioning

The second phase of the algorithm is in charge of scheduling each of the allocations obtained in the previous phase, returning the one with the lowest makespan.

For each allocation, the algorithm follows a list scheduling strategy by trying to place tasks in instances that are repartitioned according to the trees of Fig. 4. This method is outlined in Algorithm 1 (lines 7-16 for task placement and lines 17-23 for repartitioning). For example, for an A30 GPU, the algorithm first places consecutively the tasks assigned to its 4 slices; when there are no unscheduled tasks assigned to this size,

it repartitions the instance into two instances of size 2 that are treated independently to place tasks assigned to this size; finally, each of these instances can be repartitioned into two instances of size 1 when there are no unscheduled tasks assigned to 2 slices. The process is similar for an A100/H100, but before repartitioning the instance with the first 4 slices, we try to use it for tasks assigned to only 3, due to the particularity explained in the introduction. In addition, the next instance to be used is the first one to be released (line 6), which self-balances the duration of the slices and makes it preferable to leave short tasks for the end; therefore, the tasks of each instance size are selected in LPT (Longest Processing Time first) order [7] (lines 1-2).

The result is feasible, since the combinations of instances that could be generated (nodes from different branches at different levels of a tree) are precisely the valid MIG partitions shown in Fig. 1. The philosophy of MIG is that instances are independent and their dynamic repartitioning should not be conditioned by other instances. Thus, the existence of this tree and the feasibility of its combinations of disjoint instances should be fulfilled also in future GPUs, as we already explained in Section 1.2, so that the algorithm would be easily extrapolated and applicable.

Algorithm 1: Schedule an allocation by repartitioning.

Data: Functions with task times $t_{i=1}^n$ and creation/destruction times $t_{create}/t_{destroy}$, instance sizes C_G and on allocation a

Result: Instance tree of Fig. 4, where each node has a list with the indexes, begin times, and durations of the tasks running on it

- 1 Group the task times by the number of slices $s \in C_G$ given in a
 - 2 Order each group $s \in C_G$ by decreasing time (LPT)
 - 3 $reconfig_end := 0$
 - 4 Add the tree root instance R of Fig. 4 to a min heap ordered by end time with: $R.end := 0, R.tasks := []$
 - 5 **while** the heap is not empty **do**
 - 6 Pop from the heap the first instance I to end
 - ▷ Task placement
 - 7 **if** there are unscheduled tasks assigned by a to $|I|$ slices **then**
 - ▷ Give time for I creation
 - 8 **if** $I.tasks = []$ **then**
 - 9 $reconfig_end := \max(reconfig_end, I.end)$
 - 10 $reconfig_end += t_{create}(|I|)$
 - 11 $I.end := reconfig_end$
 - 12 Take the longest unscheduled task T_j using $|I|$ slices
 - ▷ Execute T_j right after in the same instance
 - 13 $task := (index := j, start := I.end, time := t_j(|I|))$
 - 14 Append $task$ to $I.tasks$
 - 15 $I.end += t_j(|I|)$
 - 16 Add I to the heap again
 - ▷ Repartitioning
 - 17 **else if** there are unscheduled tasks **then**
 - ▷ Give time to destroy I
 - 18 **if** $I.tasks$ is not $[]$ **then**
 - 19 $reconfig_end := \max(reconfig_end, I.end)$
 - 20 $reconfig_end += t_{destroy}(|I|)$
 - 21 **foreach** instance child C of I **do**
 - 22 $C.end := I.end$
 - 23 $C.tasks = []$
 - 24 Add C to the heap
-

In each repartitioning decision we make sure that the first task of a new instance is not executed until the parent instance is destroyed and the new one is created, so that the periods of creation/destruction of instances are sequenced (this is what currently happens with MIG as we saw in Section 1.2). In Algorithm 1 this is done through the *reconfig_end* variable (line 3), which is global to the repartitioning. Although the creation/destruction periods are not explicitly returned, they are implicit in the output tree: the nodes of the tree that have allocated tasks will have to be created with the time indicated by the t_{create} function (lines 8-11), and destroyed with the time indicated by $t_{destroy}$ if there is any descendant instance with allocated tasks (lines 18-20). A BFS traversal




Fig. 5. Example of task movement with a makespan improvement of more than $3/2 = 1.5$ by moving the T_j task to other slices.

of the output tree allows to properly extract this reconfiguration info together with the obtained schedule.

In the worst case, LPT sorting (line 2) involves a complexity $\mathcal{O}(n \log(n))$, which dominates over the task placement that requires $\mathcal{O}(n \cdot \log(\#slices_G))$ to insert an instance in a heap with as many nodes as the binary tree of Fig. 4 for each task (line 16).⁶ Also, the cost of repartitioning decisions is $\mathcal{O}(\#slices_G)$ due to the number of nodes in a binary tree as in Fig. 4.⁶ Thus, the complexity of Algorithm 1 is $\mathcal{O}(n \log(n) + \#slices_G)$, but since $\#slices_G$ is constant for a fixed GPU and small for current models, it is actually $\mathcal{O}(n \log(n))$. Since the allocation family can have up to $\mathcal{O}(n)$ elements, executing the method on each one takes $\mathcal{O}(n^2 \log(n))$ time in the worst case. This cost is adequate for batches with tens or few hundreds of tasks that could accumulate in a real scenario. It must be noted that we control the batch formation criteria, and we could limit it by the number of accumulated tasks. In fact, in Section 6 we will experimentally verify that with only a few tens of tasks the algorithm obtains very good results, and therefore the complexity is good enough.

The number of reconfigurations is limited by the number of non-leaf nodes in the tree, i.e., only 3 reconfigurations can occur in a schedule for A30, 7 reconfigurations for A100/H100, and at most $2 \cdot \#slices_G$ reconfigurations in the general case.⁷ Based on the reconfiguration times shown in Table 1, its overhead should not be significant in a scenario with long tasks, especially if there are many of them. We will analyze this experimentally in Section 6.

In addition, the method can be used seamlessly for multiple A30s and multiple A100/H100s; for that, there would be as many trees as GPUs, and initially, there is one node for the root of each tree to start repartitioning on. In such a case, the time complexity would also remain the same and the maximum number of reconfigurations would increase by a factor equal to the number of GPUs.

3.3. Phase 3: solution refinement

Task movement. To place a task, the previous phase considered the duration and location of those previously scheduled by using the first instance to be released, but not the duration of those that will be placed later on its slices. An example of its consequences is shown in Fig. 5, where task T_k uses a slice previously occupied by T_j , and it is long enough so that T_j should have been placed on other slices, even if they were not the first to be released. In our experience, this situation is relatively common and has a major impact in some cases (especially if there are few tasks in the batch), but can be easily mitigated by post-processing the solution to move some tasks. Fig. 5 shows an example where moving T_j reduces the makespan by a factor greater than 1.5.

Algorithm 2 outlines the movement of tasks in the solution to improve scheduling in a low complexity manner. Iteratively, the algorithm searches for *critical tasks*, i.e., using some slice that reaches the makespan, that can be moved to another instance without critical slices.

⁶ The height of the tree its $\leq \lceil \log(\#slices_G) \rceil$ since $\#slices_G$ is halved at each level, so the number of nodes its $\leq \sum_{l=0}^{\lceil \log(\#slices_G) \rceil} 2^l = 2^{\lceil \log(\#slices_G) \rceil + 1} - 1 = 2 \cdot 2^{\lceil \log(\#slices_G) \rceil} - 1 \leq 2 \cdot 2^{\log(\#slices_G) + 1} = 4 \cdot \#slices_G$.

⁷ Since non-leaf nodes are at a tree level lower than $\lceil \log(\#slices_G) \rceil$, their amount is $\leq \sum_{l=0}^{\lceil \log(\#slices_G) \rceil - 1} 2^l = 2^{\lceil \log(\#slices_G) \rceil} - 1 \leq 2 \cdot \#slices_G$.

Algorithm 2: Schedule refinement.

Data: Tree S of Fig. 4 with the schedule resulting from phase 2
Result: Tree S with the schedule refined

```

1 Create an empty queue  $Q$  for the open instance nodes
2  $stop := false$  ▷ Can be used to limit iterations
3  $\omega := makespan$  of  $S$ 
4 while not  $stop$  do
5   Push the leaf nodes of  $S$  with slices reaching  $\omega$  into  $Q$ 
6   while  $Q$  is not empty do
7     Pop the first instance node  $I$  of  $Q$ 
8     if  $I$  is the root node then
9        $stop := true$ 
10      break
11     ▷ Task movement
12     Get  $I^a \in S$  with  $|I^a| = |I|$ ,  $I^a \neq I$  for which the end time of its last
13     slice  $end(I^a) = \max_{s \in I^a}(s.end)$  is minimum
14     Bin search task  $T$  in  $I^a.tasks$  with  $T.time < \omega - end(I^a)$  which is closest
15     to  $(\omega - end(I^a))/2$ 
16     if task  $T$  has been found then
17       Pop  $T$  from  $I^a.tasks$ 
18       Insert  $T$  in  $I.tasks$  ordered by  $T.time$ 
19       Update end time of the slices  $s.end$  for  $s \in I \cup I^a$ 
20     else
21     ▷ Task swapping
22     Find  $T_k \in I.tasks, T_j \in I^a.tasks$  with:
23      $0 < (T_k.time - T_j.time) < (\omega - end(I^a))$  and  $(T_k.time - T_j.time)$ 
24     closest to  $(\omega - end(I^a))/2$ 
25     if  $T_k$  and  $T_j$  have been found then
26       Pop  $T_k$  from  $I.tasks$  and  $T_j$  from  $I^a.tasks$ 
27       Insert  $T_k$  neatly in  $I^a.tasks$  and  $T_j$  in  $I.tasks$ 
28       Update end time of the slices  $s.end$  for  $s \in I \cup I^a$ 
29     else if parent of  $I$  has not been added to  $Q$  then
30       Push the node instance parent of  $I$  into  $Q$ 
31   Update makespan  $\omega$  from new slice endings
32 Traverse tasks lists of  $S$ , updating task start times and reconfig. times

```

To find such tasks the algorithm follows a greedy strategy by traversing the critical task subtree in a reverse BFS path (from the leaves towards the root). To do so, it starts with the leaves corresponding to the slices that reach the makespan (line 5), and expands them to their parents in case it is not possible to move any of their tasks (lines 23-24). For an instance I of a *critical path* (i.e., in a sequence of critical tasks), it considers the alternative instance I^a with the same size (i.e., $I^a \neq I$ s.t. $|I^a| = |I|$) and minimum completion time for its slices (line 11). Then, among the tasks that are assigned to instance I , it looks for the most suitable one that does not reach the makespan when moving it to I^a (line 12), that is, $T.time < \omega - end(I^a)$ and $T.time$ closer to $(\omega - end(I^a))/2$, where ω denotes the makespan. The reason for considering half of the available time is that it will be divided between I and I^a during the movement, and ideally, this division should be as balanced as possible. As the tasks of each node are sorted by duration, this can be done efficiently with a binary search. If successful, the algorithm inserts T neatly into I^a (lines 13-16), and if not, it opens the parent instance of I unless it has been previously opened by another node (lines 23-24). The algorithm will have reduced the makespan of all the critical slices when all the nodes of an iteration are closed. Then, a new iteration will be started with the new critical slices (line 25 and again line 5). For each iteration the algorithm only recomputes the new makespan, so that the start times of the task that provide the detailed scheduling, and the reconfiguration times added/removed, are not computed until all the movements are done (line 26). **Task swapping.** It may also happen that a critical task cannot be moved to another instance without reaching the current makespan, but can be swapped for a shorter task to improve it. This is illustrated in Fig. 6 with an improvement factor of almost 1.25. Thus, if during an iteration of Algorithm 2 we cannot perform any movement from an instance I to its alternative I^a , we look for a task T_k of I and a task T_j of I^a to swap them (lines 18-22). The requirements to improve with swapping are that $0 < t_k(s) - t_j(s) < \omega - end(I^a)$, where $s = |I| = |I^a|$.




Fig. 6. Example of task swapping with a makespan improvement of almost $5/4 = 1.25$.

However, we again try to choose the tasks so that $t_k(s) - t_j(s)$ is as close as possible to $(\omega - \text{end}(I^a))/2$ to balance the time reduction in I with the increase in I^a in the best way according to the available margin.

Refinement stop and complexity. Obviously, the refinement process ends if the root node is opened in an iteration because the makespan of some slice could not be reduced (lines 8-10). As we will see in Section 6, it almost always ends that way after few moves/swaps, but it is still advisable to set a sufficiently strict stopping condition to avoid too long execution times in the general case. We propose to set a maximum number of iterations, or a minimum percentage of makespan improvement to avoid wasting time with too subtle refinements. In Section 6 we experimentally analyze the improvements produced by refinement in the schedule.

It is easy to check that each iteration can be implemented in $\mathcal{O}(n \cdot \#slices_G)$ time complexity. This is because, at most, $\#slices_G$ critical tasks are moved or swapped (as much as initially opened critical slices) and the swapping requires $\mathcal{O}(n)$ time because the lists $I.tasks$ and $I^a.tasks$ are sorted, so we can efficiently search the tasks T_j and T_k that comply with line 18 of Algorithm 2 by a double-pointer strategy. In addition, the movement operation only requires $\mathcal{O}(\log(n))$ time due to the task ordered insertion in its new instance. However, the actual cost will usually be much lower because not all tasks will be critical and not all will be allocated in a pair of instances.

3.4. Scheduling execution

Finally, to execute tasks on a MIG-capable GPU according to the schedule provided by the FAR algorithm, the nodes of the repartitioning tree (GPU instances) must be managed by different threads. As shown in Algorithm 3, the implementation described in [8] traverses the repartitioning tree generated by FAR in a descending manner. Each node with tasks calls to create its corresponding MIG instance on the GPU, sequentially executes the assigned tasks in that instance, and then destroys the instance. Additionally, every node –whether it has tasks or not– requests its descendants to perform the same actions in separate threads. This approach enables tasks from different instances to run concurrently. Algorithm 3 also reflects the calculation and reporting of the start and end time of each task, relative to the initial start of the batch execution. These times should be very similar to the ones used by FAR, due to the real isolation of the instances that we checked in Section 2.5 and the stability in the duration of reconfigurations and tasks that we also observed; in any case, this report allows verifying that the estimate of times simulated by FAR fits reality due to these hypotheses.

4. Multi-batch

As seen, FAR solves the moldable offline scheduling for a single batch of tasks, but in a real scenario we deal with several batches over time; this implies the necessity of appending consecutive solutions in an intelligent way to maximize the resource usage.

4.1. Concatenation issue

If we simply concatenate the scheduling produced by FAR, the first task of a new batch typically waits until the last task of the previous

Algorithm 3: Scheduling execution on GPU.

```

Data: Repartitioning tree  $S$  with the schedule produced by FAR
1 Measure the current time as  $init\_time$ 
2 Call  $execute\_tree(\text{root of } S, init\_time)$  ▷ Recursive execution
3 Function  $execute\_tree(\text{tree instance } I, init\_time)$ :
  ▷ Create, execute and destroy physically on GPU
  if  $I.tasks$  is not empty then
    4  $instance\_handler := GPU\_instance\_create(I)$ 
    5 foreach task  $T$  in  $I.tasks$  do
    6   Show current time difference with  $init\_time$  as  $T$  start
    7    $GPU\_execute(T, instance\_handler)$ 
    8   Show current time difference with  $init\_time$  as  $T$  end
    9    $GPU\_instance\_destroy(instance\_handler)$ 
  ▷ Repeat concurrently on children
  11 foreach instance child  $C$  of  $I$  do
  12   Run  $execute\_tree(C, init\_time)$  in a new thread
  13 Wait for all children threads to complete
    
```

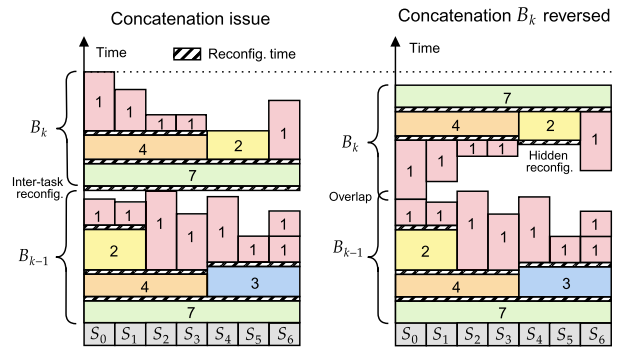


Fig. 7. Example of concatenation issue and improvement by reversing B_k .

batch finishes because FAR schedules initial tasks using many slices (commonly all of them), and it is very likely that one of them will be the last slice to be released by the previous batch. For example, as shown in Fig. 7 (left), a task that uses 7 slices in A100 will have been scheduled at the beginning of the scheduling of batch B_k and cannot start until the last one of batch B_{k-1} is finished since it will use its slices.

4.2. Reverse batch scheduling to overlap

To mitigate this, our proposal reverses the order of schedules for consecutive batches. Tasks using fewer slices from one batch can be followed by similar tasks from the next batch, enhancing batch overlap. The schedule depicted in Fig. 7 (right) shows the improvement when reversing the scheduling order of B_k relative to B_{k-1} . Note that this overlap between batches refers to running tasks of the current batch B_k in periods where its slices were idle in B_{k-1} ; it does not mean to run tasks of both batches on the same slices at the same time. Fig. 7 also shows that batch reversing can alleviate reconfigurations by hiding them in the overlap or eliminating them when the last instance of B_{k-1} coincide with the first instance of B_k , which is more likely.

The best feasible overlap between the inverted schedule of B_k and that of B_{k-1} can be obtained in time $\mathcal{O}(\#slices_G)$ (constant time) by determining which slice forces B_k to start later: just try to start B_k after B_{k-1} slice by slice and choose the option with the highest makespan (this is the best one that ensures feasibility on all slices).

4.3. Moving/swapping tasks to overlap

To increase the overlap between batches, we can move/swap tasks of the current batch B_k to better fit with the schedule of the previous batch B_{k-1} . We focus on task in B_k that starts immediately after the tasks in B_{k-1} , as these are critical for increasing inter-batch overlap.

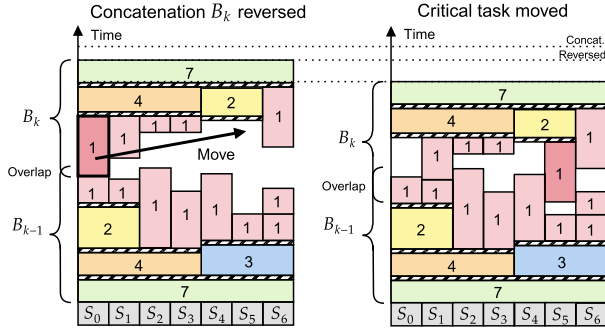


Fig. 8. Moving a critical task to improve the reversed batch concatenation.

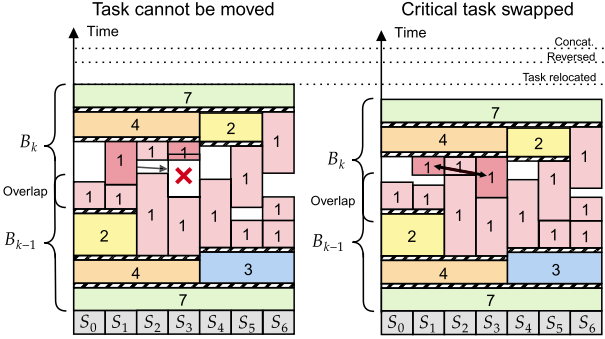


Fig. 9. Swapping a critical task which cannot be relocated to improve the reversed batch concatenation.

Fig. 8 illustrates how the inverted concatenation from the previous example is further improved by moving a task from B_k , and Fig. 9 demonstrates how further improvement is possible by swapping tasks in B_k , even when no critical tasks can be moved. The operations are carried out in the same manner as the refinement described in Section 3.3, with the difference that the idle time of the slices between batches now takes the role that the makespan margin ($makespan - end(I^a)$) played in the refinement. In addition, to perform the overlap increasing, the limiting slices must be recalculated after each iteration with complexity $\mathcal{O}(\#slices_G)$, but, as in the refinement, without having to update the start time of each task until the end of the iterations (see line 26 of Algorithm 2).

5. Approximation factor

Next, we bound the makespan resulting from phase 2 of FAR according to the form required by Theorem 1, to obtain an upper bound of its factor with respect to the optimum of the moldable problem by applying the theorem. Note that phase 3 for refinement can only improve the solution, so that factor is also an upper bound if it is used. Reconfiguration times are not considered in this factor, but they must be negligible in relation to the duration of the tasks in the target scenario of the algorithm.

As proved in [21], concatenating schedules with an approximation factor ρ has a *competitive ratio* 2ρ respect to the optimum, even if it is done without overlapping the executions. The *competitive ratio* is widely used in the literature to evaluate online schedulers [30], and is defined as the maximum possible ratio between the makespan achieved by the scheduler processing the tasks on the fly and the optimal offline makespan if the tasks were available from the beginning in an offline way.

5.1. Factor of NVIDIA A30

For modeling the A30 it is sufficient to apply the List Scheduling proof to deduce the approximation factor [6]. For the sake of self-

containment and since we will reuse the reasoning for the rest of the models, we briefly explain the particularized argument for this GPU.

The method can always use the released slices to immediately run other unscheduled tasks. This is because each node in its repartition tree has as ancestors all possible instance sizes larger than its own (Fig. 4). Therefore, all tasks matched to larger sizes have already been scheduled and the remaining ones fit in its slices. As exemplified in Fig. 10a, it means a schedule with the entire area occupied by tasks up to time l , when a critical task T_j (one reaching the makespan ω) starts. It follows that the area up to l , which is $4l$ because A30 has 4 slices always occupied up to that point, is at most the total task area removing the $h_j = t_j(s_j)$ portion of T_j : $4l \leq \sum_{i=1}^n s_i \cdot t_i(s_i) - h_j$. Since also $h_j \leq h_{\max} = \max_{i=1}^n t_i(s_i)$, it follows:

$$\omega = l + h_j \leq \frac{1}{4} \sum_{i=1}^n s_i \cdot t_i(s_i) - \frac{1}{4} h_j + h_j \leq \frac{1}{4} \sum_{i=1}^n s_i \cdot t_i(s_i) + \frac{3}{4} h_{\max}.$$

Now, we have the conditions to apply Theorem 1 deducing an approximation factor of $7/4 = 1.75$ respect to the optimal makespan of the moldable problem:

$$\omega \leq \left(\frac{1}{4} \cdot 4 + \frac{3}{4}\right) \omega^* = \frac{7}{4} \omega^*.$$

If we use FAR for two A30s instead of one, we have the same with 8 slices and can deduce that $\omega \leq 15/8 \omega^* = 1.875 \omega^*$ by the same argument. In general, if we use g A30s, $\omega \leq (8g - 1)/(4g) \omega^*$, which is bound by $2 \omega^* \forall g$.

Applying the result of [21] for one A30 we get a competitive ratio of $2 \cdot 7/4 = 7/2 = 3.5$ for the tasks of all batches received online with respect to their offline optimum. In general, for any number of A30s we will have a factor of $2 \cdot 2 = 4$ for all batches.

5.2. Factor of NVIDIA A100/H100

In this case, it is possible that the algorithm cannot place tasks on the slices of the first task to finish, not even partitioning it. For example, there may be unscheduled tasks requiring 4 slices and the first instance to finish is on slices $\{S_4, S_5, S_6\}$ (right branch of the tree), so no task can be placed and idle slices (a gap) are left. Following the same notation as in the previous section we will refer to it as idle period before time l , which is the start time of a critical task that reaches the makespan ω . This can only happen in tree branches where not all instance sizes are present, since a node can be reached without having scheduled all tasks for larger instances. Therefore, it can only happen in the slices $\{S_4, S_5, S_6\}$, and in S_3 when we assign to the first 4 slices a task of size 3 (see Fig. 4). For bounding we distinguish 3 cases:

1. $\{S_4, S_5, S_6\}$ are not idle before time l .

The only slice that can be idle before time l is S_3 . Therefore, the occupied area under l is at least $6l$. Just as we reasoned for the A30 above, we have that $6l \leq \sum_{i=1}^n s_i \cdot t_i(s_i) - h_j$, with $h_j = \omega - l \leq h_{\max}$, so:

$$\omega = l + h_j \leq \frac{1}{6} \sum_{i=1}^n s_i \cdot t_i(s_i) - \frac{1}{6} h_j + h_j \leq \frac{1}{6} \sum_{i=1}^n s_i \cdot t_i(s_i) + \frac{5}{6} h_{\max}.$$

2. S_4 or S_5 are idle before time l .

When one of these slices is left idle, all unscheduled tasks need instances of size 4 because their tree branches have all the other sizes: 7, 3, 2, 1. The makespan is given by a task running in slices S_0 to S_3 , or S_4 to S_6 , so: $\omega = \max(\omega_{S_0-S_3}, \omega_{S_4-S_6})$. In $S_0 - S_3$ all tasks use 7 or 4 slices because, as we have explained, the possible unscheduled tasks are for 4 slices (the 4-slice node in the $S_0 - S_3$ branch cannot have been partitioned). Therefore, the area under $\omega_{S_0-S_3}$ has no gaps and is bounded by all the task area: $4 \cdot \omega_{S_0-S_3} \leq \sum_{i=1}^n s_i \cdot t_i(s_i)$.

On the other hand, all slices except perhaps S_6 are always busy below the first time q when S_4 or S_5 are idle: as we argue, $S_0 - S_3$ is




Fig. 10. Cases to deduce the upper bound of the approximation factor.

filled with task using 7 or 4 slices with no gaps (Fig. 10b). We use the same argument as before with $\omega_{S_4-S_6} - q \leq h_j \leq h_{\max}$ for the 6 slices always used below q , so $\omega_{S_4-S_6} \leq 1/6 \sum_{i=1}^n s_i \cdot t_i(s_i) + 5/6 h_{\max}$. In short:

$$\begin{aligned} \omega &= \max(\omega_{S_0-S_3}, \omega_{S_4-S_6}) \\ &\leq \max\left(\frac{1}{4} \sum_{i=1}^n s_i \cdot t_i(s_i), \frac{1}{6} \sum_{i=1}^n s_i \cdot t_i(s_i) + \frac{5}{6} h_{\max}\right). \end{aligned}$$

3. S_6 is idle before time l , but S_4 and S_5 not.

Unscheduled tasks when S_6 is idle may need instances of sizes 2 or 4. If there are only tasks for size 4, we can use the argument of case 2 obtaining the same bound. If there are tasks for size 2, S_3 may also be idle by passing through the node using only 3 slices of $S_0 - S_3$. However, there are no tasks for size 1, since there are none to place in S_6 and the parents of them in the tree have size 2, so there are tasks to place in 2. Therefore, the time h_j of the critical task uses at least 2 slices and its area is at least $2h_j$ instead of h_j (Fig. 10c). As there are always 5 slices occupied up to l and exceeds an area of at least $2h_j$:

$$\omega \leq \frac{1}{5} \sum_{i=1}^n s_i \cdot t_i(s_i) - \frac{2}{5} h_j + h_j \leq \frac{1}{5} \sum_{i=1}^n s_i \cdot t_i(s_i) + \frac{3}{5} h_{\max}.$$

Joining results, we have the following bound in the form of Theorem 1 and we can apply Theorem 1 to deduce a 2-approximation for the moldable problem:

$$\begin{aligned} \omega &\leq \max\left(\frac{1}{6} \sum_i s_i \cdot t_i(s_i) + \frac{5}{6} h_{\max}, \frac{1}{4} \sum_i s_i \cdot t_i(s_i), \frac{1}{5} \sum_i s_i \cdot t_i(s_i) + \frac{3}{5} h_{\max}\right) \\ &\implies \omega \leq \max\left(\frac{1}{6} \cdot 7 + \frac{5}{6}, \frac{1}{4} \cdot 7, \frac{1}{5} \cdot 7 + \frac{3}{5}\right) \omega^* = 2\omega^*. \end{aligned}$$

Applying the result of [21] we get a competitive ratio of $2 \cdot 2 = 4$ for the tasks of all batches received online with respect to their offline optimum.

6. Experimental evaluation

For the experimental evaluation, we will employ the execution times obtained when using MIG on an A100 GPU for the Rodinia suite [2]; however, due to the limited number of examples in the benchmark, and for the sake of generality, we also generated synthetic times mimicking the behaviors observed in those tests with some randomness. The logic of the task generator and the specific scenarios evaluated are described in Section 6.3. All times, both measured and synthetically generated, are in seconds, to keep them in agreement with the reconfiguration times presented in Table 1, which uses the FAR algorithm.

6.1. Experimental objectives

Specifically, we develop a thorough evaluation of FAR with a number of goals in mind, that are dissected throughout the section as follows:

1. Quantify the gap between simulated schedules generated by FAR (using its estimated times) and the real execution of these schedules on the GPU. Specifically, the goal of Section 6.2 is to verify if the simulation accurately reflects real-world performance, ensuring that it can be reliably used to evaluate the algorithm. This would allow it to be tested with synthetic workloads generated on demand, covering a wide range of scenarios.
2. An estimation of the error of FAR with respect to the optimum for realistic workloads, comparing it with the theoretical factor deduced in Section 5, is given in Section 6.4. In addition, we analyze its error according to the number of scheduled tasks, checking if it is necessary to schedule large task batches for good results.
3. A comparison of FAR versus alternative proposals such as fixed MIG partitioning without reconfiguration and the only MIG scheduler found in the literature [5] is given in Section 6.5.
4. An analysis of the contribution of refinement techniques to improve scheduling and its cost according to the number of iterations performed in this process is given in Section 6.6.
5. Finally, an evaluation of the extension of FAR for multi-batch, together with an analysis of the contribution of the improved batch combination is presented in Section 6.7.

The results of FAR shown include the reconfiguration time, which is limited and should not have a significant impact in a scenario with long tasks, as explained in Section 3.2. In particular, experiments will allow us to confirm this assumption, especially when comparing against a fixed partitioning without that overhead.

6.2. Simulated execution

To evaluate the proposal in a wide range of scenarios, the task execution is simulated, but supported by real data. It should be noted that this simulation should reproduce such a simple process with high fidelity due to two validation keys: first, the verification of sequential reconfiguration behavior with highly stable real execution times used by FAR, as detailed in Section 2.1; and second, the confirmation of true instance isolation, as discussed in Section 2.5. Under these assumptions, the transition from FAR's output to the actual execution schedule is straightforward, simple and deterministic by the scheme explained in Section 3.4. In any case, it is convenient to corroborate the accuracy of this simulation by comparing with the real GPU scheduler provided [8].

Table 3 presents the end times for the simulated and real executions of 9 Rodinia kernels scheduled in NVIDIA A30 as a batch using FAR. The real execution times are derived from running the tasks on the GPU with the software we provide to the community [8]. The end time refers to the total time from the start of the execution of all tasks until the completion of each task, which means that it represents partial durations of the entire execution, not the durations of the individual tasks. Along with the end times, the table shows the percentage deviation between the simulated and actual executions, providing a clear indication of their

Table 3

Comparison of simulated vs. real end execution times, in a batch of 9 Rodinia kernels scheduled by FAR on NVIDIA A30.

Kernel	End time (s)		Deviation (%)
	Simulation	Real execution	
PathFinder	21.09	20.64	-2.18
LavaMD	21.72	21.62	-0.46
HotSpot	22.31	22.25	-0.26
Gaussian	22.45	22.92	2.05
NW	22.79	23.26	2.02
Huffman	23.01	23.54	2.25
HeartWall	23.05	22.98	-0.30
ParticleFilter	23.23	23.38	0.64
LU	28.86	28.24	-2.20

similarity. The critical task determining the makespan is highlighted in bold. For the sake of reproducibility, the scheduling software includes these 9 kernels as a pre-configured test, and the documentation clearly explains how to run it.

As can be seen in Table 3, the deviations are minimal, with real execution times deviating at most 2% from the simulated counterparts. This justifies that simulations based on the scheduler's timing parameters closely reflect real executions. This approach allows us to experiment with diverse on-demand workloads without the need to locate applications that exhibit these characteristics.

6.3. Synthetic times generation

As seen in Fig. 3, tasks initially scale linearly with the slice count and become sub-linear beyond a certain size. Also, some memory-bound tasks exhibit super-linear scaling initially. Our generator receives the number of tasks n , the percentage of tasks p_s that scale well up to each instance size s (near- or super-linear), and the percentage p_{sup} that start with super-linear speedup from 1 to 2 slices since they are memory-bound. Also, we consider the transition of a memory-bound task to compute-bound with 0.3 of probability per slice.

For example, the times modeling for the A30 would be:

- $n_1 = \lfloor n \cdot p_1 \rfloor$ tasks only exploit one slice,⁸ i.e., they have sub-linear speedup for all slice increments.
- $n_2 = \lfloor n \cdot p_2 \rfloor$ tasks exploit two slices.⁸ Of them, $\lfloor p_{sup} \cdot n_2 \rfloor$ improve super-linearly from 1 to 2 slices (memory-bound), and $\lfloor (1 - p_{sup}) \cdot n_2 \rfloor$ improve near-linearly (compute-bound). Both scale sub-linearly from 2 to 4 slices.
- $n_4 = \lfloor n \cdot p_4 \rfloor$ tasks exploit all slices.⁸ On the one hand, $\lfloor p_{sup} \cdot n_4 \rfloor$ improve super-linearly from 1 to 2 slices, maintaining a super-linear speedup from 2 to 4 slices with a probability of $(1 - 0.3)^2$ (still memory-bound despite the increase of 2 slices), but moving to sub-linear speedup with a probability of $1 - (1 - 0.3)^2$ (becomes compute-bound). On the other hand, $\lfloor (1 - p_{sup}) \cdot n_4 \rfloor$ tasks improve near-linearly from 1 to 2 and from 2 to 4 slices.

We generate the time $t_i(1)$ of each task on 1 slice with a uniform distribution $U(t_{min}, t_{max})$, where t_{min} and t_{max} are configurable parameters. Then, we generate the time $t_i(s+1)$ from $t_i(s)$ according to the speedup type as follows:

$$t_i(s+1) = (s+r)/(s+1) \cdot t_i(s), \text{ with } r \text{ taken random as}$$

⁸ Applying percentages may result in non-integer values, so we floor the result. To create exactly n tasks we iteratively add one to the size farthest from the exact: $n_j := n_j + 1$ for $j = \arg \max_{i \in C_G} (n \cdot p_i - n_i)$, while $\sum_{i \in C_G} n_i < n$.

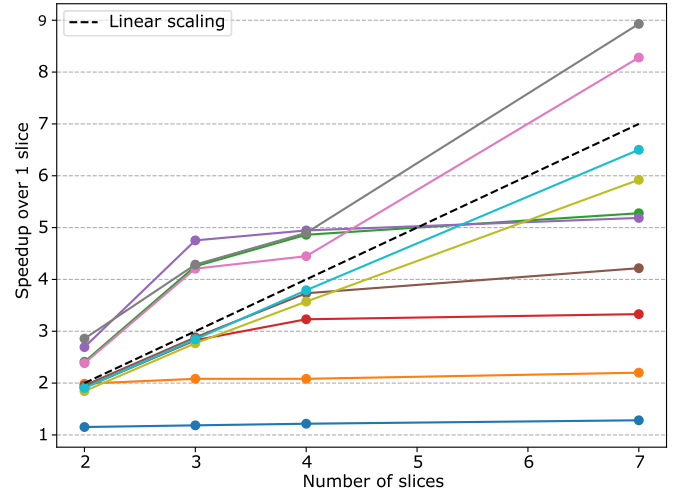


Fig. 11. Synthetic speedups for $n = 10$ tasks in A100/H100 using $p_1 = p_2 = 10\%$, $p_3 = 20\%$, $p_4 = p_7 = 30\%$ and $p_{sup} = 50\%$.

$$r \in \begin{cases} N(-0.25, 0.25)_{[-0.5, 0]} & \text{for super-linear speedup,} \\ N(0.1, 0.1)_{[0, 0.2]} & \text{for near-linear speedup,} \\ N(0.75, 0.25)_{[0.5, 1]} & \text{for sub-linear speedup,} \end{cases}$$

where $N(\mu, \sigma)_{[a, b]}$ is a normal distribution of mean μ and standard deviation σ clipped to $[a, b]$.

With $r = 0$ the speedup would be linear and the higher r the worse the improvement; therefore, we generate $r < 0$ for super-linear, $r \approx 0$ for near-linear, and $r > 0$ for sub-linear speedup. We use normal distributions since application behaviors follow patterns and speedup factors that are unlikely to be uniformly distributed. Since the distribution is not bounded, we clip the values to guarantee the desired speedup type (e.g., no $r > 0$ in a super-linear speedup case). In particular, we guarantee that $r \leq 1$, so that $t_i(s+1) \leq t_i(s)$, fulfilling property 1 of monotony (Section 2.4). Times for $s \notin C_G$ are useful for generation, but they are discarded afterwards as they are not input for the algorithm. Fig. 11 shows an example of generated speedups, which resemble those shown with Rodinia in Fig. 3.

Armed with the generator, we will employ synthetic loads with different combinations of parameters, covering several representative and diverse scenarios. In this sense, we denote the parameter configurations we will use for A100 as follows:

- **POORSCALING:** The percentages of tasks p_s that scale well up to s slices is set to $(p_1, p_2, p_3, p_4, p_7) = (50, 50, 0, 0, 0)\%$, i.e., half of the tasks can only exploit 1 slice well, the other half only scale well to 2 slices, and no task scales to more slices.
- **MIXEDSCALING:** The scaling percentages are set to $(p_1, p_2, p_3, p_4, p_7) = (20, 20, 20, 20, 20)\%$, i.e., the same 20% of tasks can get to exploit well to every possible instance size.
- **GOODSCALING:** The task scaling percentages are set to $(p_1, p_2, p_3, p_4, p_7) = (0, 0, 0, 50, 50)\%$, i.e., half of the tasks that can exploit up to 4 slices well and the other half that can exploit all 7 GPU slices.
- **WIDETIMES:** The range of task duration times is set to $[t_{min}, t_{max}] = [1, 100]$, i.e., a wide range between 1 and 100 seconds that conceives short tasks along with long duration tasks.
- **NARROWTIMES:** The range of task duration times is set to $[t_{min}, t_{max}] = [90, 100]$, i.e., all tasks are long with a narrow range of times between 90 and 100 seconds.

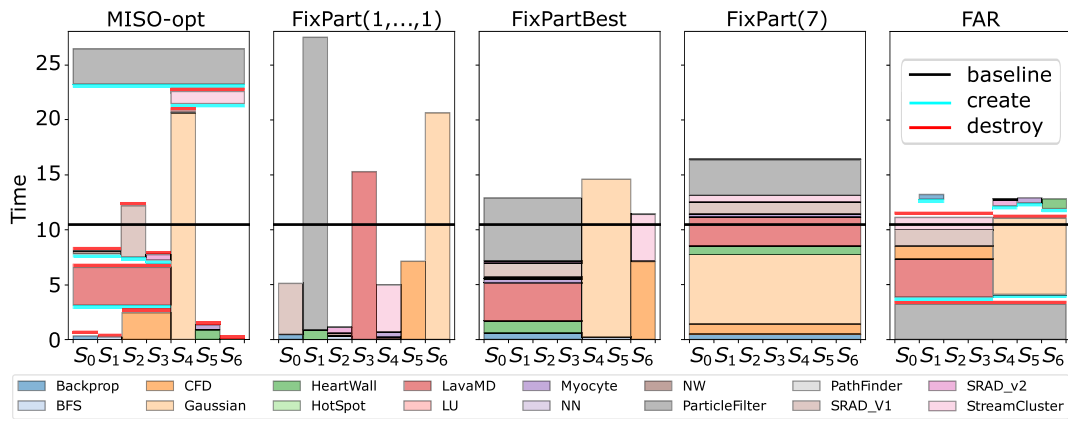


Fig. 12. Scheduling of the Rodinia kernels for NVIDIA A100 with different approaches.

Table 4
Average of ρ in 1000 execution of FAR with WIDETIMES for the A100.

Task configuration	Number of tasks (n)					
	10	15	20	25	30	35
POORSCALING	1.23	1.08	1.04	1.03	1.02	1.02
MIXEDSCALING	1.20	1.08	1.04	1.03	1.02	1.02
GOODSCALING	1.21	1.07	1.05	1.03	1.02	1.01

For the different metrics that we will be presenting, the mean values of 1000 executions are reported; the deviations are not included for brevity, since they are relatively small in all cases and are not relevant for the discussion.

6.4. Error with respect to the optimum

Since the optimal solution cannot be computed in a reasonable time (as it is an NP-hard problem), to evaluate FAR we use the lower bound of the optimal makespan that arises from assigning minimum area to each task and distributing the total area evenly among the slices:

$$baseline = \sum_{i=1}^n \min_{s \in C_G} (s \cdot t_i(s)) / \#slices_G \leq \omega^*.$$

Results are presented hereafter as

$$\rho = \omega / baseline,$$

which overestimates the optimal makespan error: $\omega \leq \rho \cdot \omega^*$.

6.4.1. Comparison using Rodinia workloads

FAR obtains $\rho = 1.22$, i.e., at most 22% error with respect to the optimal solution, using as input the times of the 16 Rodinia kernels profiled in the A100 (those shown in Fig. 3). This is a first indication that our proposal obtains good results even with reconfiguration overhead (*baseline* does not include it but FAR makespan does); in fact, this value is much better than the theoretical worst case (approximation factor of 2 for the A100, i.e., 100% of error). In this case, the execution time of FAR has been between 1 and 2 milliseconds in several runs, confirming the efficiency of the algorithm. It should be noted that the tests have been performed on a personal machine that is not excessively powerful (AMD Ryzen 5 4600H processor).

6.4.2. Comparison using synthetic workloads

For the sake of generality, Table 4 shows the ρ value with synthetic inputs for A100, according to the number of tasks n , for the three task configurations according to their scalability described above. For the sake of brevity, only results for tasks of duration WIDETIMES are reported, observing similar results for NARROWTIMES times.

We observe negligible errors when using $n \geq 15$ tasks: below 10% of error with 15 tasks, which progressively decreases to 2% with 30-35 tasks. With just 10 tasks we get errors of around 20% due to imbalances in the slice utilization of the last tasks to be placed, which has more impact with a reduced number of tasks. However, these errors are probably much smaller in reality since ρ overestimates more with few tasks because the optimal schedule is further away from the baseline (the baseline is the perfect distribution of the task area among the slices). Also, the error for $n = 15$ tasks, which is a similar amount to the Rodinia kernels used previously, is quite smaller because the execution times are longer and hence the reconfiguration overhead less significant. In all cases the errors are well below the approximation factor of 2 that was proved for the A100. Similar trends are observed for A30, but with much smaller errors for each n : $\rho \sim 1.01$ for $n = 10$.

Additionally, for the same n no significant differences are observed when varying task scaling capabilities: approximately the same error is made when scheduling a batch of tasks that can only leverage 1 and 2 slices (POORSCALING), a batch of tasks that can leverage up to 4 and 7 slices (GOODSCALING), and a very diverse batch of leverage capabilities evenly distributed among the slices (MIXEDSCALING).

The average times to compute the schedule with the configurations in Table 4 are between 1 and 2 milliseconds, which is a negligible cost. The numbers of tasks reported in the table are those that can probably be accumulated to run FAR. In any case, we have tested the time cost of the algorithm with $n = 100$, $n = 500$ and $n = 1000$ tasks, obtaining average execution times of 12.32 ms, 532.21 ms and 1620.82 ms respectively (1000 executions with MIXEDSCALING configuration). This is still affordable, as with hundreds of tasks we are even below the time of a GPU reconfiguration (reported in Table 1). Although a thousand tasks is probably excessive for a real scenario, it would still be equivalent to about 3 or 4 reconfigurations of the A100, which is adequate considering that we have posed an application scenario where the tasks last considerably longer than the reconfiguration time (otherwise the dynamism of MIG does not make much sense and it is preferable not to reconfigure).

As a summary of conclusions:

- FAR schedules close to the optimum (*baseline*) and its error is actually well below the approximation factor.
- The reconfiguration time is not relevant with long duration tasks.
- It is not necessary to schedule a large number of tasks per batch to obtain good results with long duration tasks.
- The algorithm adapts well to different (synthetic) workloads in terms of the GPU usage capabilities of the tasks.
- The cost of FAR is affordable even for several hundred tasks.

6.5. Evaluation of FAR versus other approaches

In the following, we compare FAR against two alternative solutions:

- **MISO-OPT**: This is the solution proposed in [5]. In this proposal, tasks are scheduled in FIFO order using the partition that maximizes the sum of speedups: the next partition will be $P = \{T_0, \dots, T_{|P|-1}\} \in \text{MIG-Part}_G$ maximizing $\sum_{i=0}^{|P|-1} \text{speedup}_{k+i}(|T_i|)$, where T_k is the next unscheduled task in FIFO order.
- **FIXPART(P)**: Given a partition P , all tasks are scheduled in FIFO order on the first free instance. This represents a static configuration of the GPU done before the start of the execution. Note that, for the A100 and H100 GPUs, there are 19 different **FIXPART(P)** strategies (5 for the A30). We denote by **FIXPARTBEST** the variant using the partition that provides the lowest makespan among all possible options.

6.5.1. Comparison using Rodinia workloads

Fig. 12 illustrates the scheduling comparison between those proposals and FAR for the 16 Rodinia kernels. It highlights a major issue with **MISO-OPT**, whose partitioning decisions are independent, resulting in longer idle periods as instances wait for preceding tasks to complete. In fact, as a good sum of speeds is intended without considering the individual duration of the tasks, it is common that very long tasks run in parallel with very short ones and high idle periods arise between them (for example, *Gaussian* in Fig. 12). This is solved in phase 1 of FAR by using more slices to reduce the duration of the longest task in the successive candidate allocations.

Also note that a fixed partition with 7 instances of size 1 generates a problem for long tasks that far exceed the ideal makespan for scheduling all tasks (*ParticleFilter*, *LavaMD* and *Gaussian* in Fig. 12). Meanwhile, with 1 fixed instance of size 7, i.e., the equivalent of using the entire GPU for each task, there is also considerable room for improvement due to the poor scaling of some tasks to so many slices. Even with the best possible fixed partition (**FIXPARTBEST**), there is room for improvement that FAR exploits, which lies in the inflexibility to deal with tasks with different scaling capabilities without changing partitions. In fact, we find that these problems outweigh the sequential reconfiguration avoided by **FIXPART**, even though the execution of these Rodinia kernels are not too long. Note that FAR provides the least makespan schedule, even with idle periods due to sequential reconfiguration: the cyan/red rectangles of instance creation/destruction do not overlap leaving visible gaps.

Let σ_A be the ratio between the makespan of an approach A and the makespan of FAR:

$$\sigma_A = \omega_A / \omega_{\text{FAR}}.$$

For the Rodinia kernels shown in Fig. 12, we obtained $\sigma_{\text{MISO-OPT}} = 2.10$, $\sigma_{\text{FIXPART}(1, \dots, 1)} = 2.18$, $\sigma_{\text{FIXPARTBEST}} = 1.16$ and $\sigma_{\text{FIXPART}(7)} = 1.26$. These results show an important advantage over these schedulers, consolidating the idea that task co-scheduling is worthwhile (**FIXPART(7)** does not do it), betting on flexible dynamic reconfiguration despite the overhead due to partition changes (**FIXPART** does not do it), and conditioning partitions to task execution times trying not to leave idle slices (**MISO-OPT** does not do it).

In terms of execution costs, there is not much to report compared to the other proposals. The scheduling calculation times for all of them are around the millisecond for the Rodinia subset, as we obtained for FAR in Section 6.4.1.

6.5.2. Comparison using synthetic workloads

Table 5 shows the σ_A ratios for synthetic inputs based on 15 tasks over six workload types.

Note that FAR improves more over **MISO-OPT** when the range of times is wide (**WIDETIMES**) versus more homogeneous task duration

(**NARROWTIMES**). This is because the coexistence of long and short tasks accentuates the idle gaps of that approach, but FAR can handle with it.

A great improvement over **FIXPART(1, ..., 1)** is observed when tasks are able to exploit up to 4 and 7 slices (**GOODSCALING**), and analogously versus **FIXPART(7)** when they barely exploit 1 and 2 slices (**POORSCALING**), which are obviously the worst cases for such proposals. However, FAR also outperforms them by far for the opposite configurations, despite those proposals offer their best performance: around 30% of improvement versus **FIXPART(1, ..., 1)** with tasks that barely scale (**POORSCALING**), and also 30% better versus **FIXPART(7)** with tasks that scale well (**GOODSCALING**). In contrast, **FIXPARTBEST** comes close to FAR when the tasks exploit almost the entire GPU (**GOODSCALING**), but is far below in the rest of the cases, which primarily motivates the use of MIG. Finally, FAR improves substantially to all approaches when the GPU resource exploitation is mixed (**MIXEDSCALING**), which is probably the most plausible case in real scenarios.

For the size of $n = 15$ tasks reported in Table 4, the cost of the different proposals have been very similar: in the case of FAR each planning has taken on average 19.43 ms to compute, while the rest of the proposals move from 9.41 ms for **FIXPART(1, ..., 1)** to 13.45 ms for **MISO-OPT**. These differences widen slightly if we increase the number of tasks but, as we already explained in Section 6.4.2, the costs of FAR are still quite affordable, especially if we consider the significant improvement obtained for the schedule. For example, for $n = 1000$ tasks, FAR took 1620.82 ms to compute the schedule, **MISO-OPT** took 1110.65 ms, **FIXPART(1, ..., 1)** took 935.65 ms, **FIXPARTBEST** took 1035.26 ms and **FIXPART(7)** took 933.70 ms (average of 1000 runs with the **MIXEDSCALING** configuration).

In summary, FAR improves significantly on both proposals by properly using GPU reconfiguration in very diverse scenarios. For example, as the last column of Table 5 shows, if the tasks do not scale perfectly, **FIXPART(7)** produces a makespan between 2 and 3 times larger than FAR. FAR even improves significantly over the proposals without GPU reconfiguration in ideal cases such as **FIXPARTBEST**, despite the reconfiguration costs that such proposals do not need. Only **FIXPARTBEST** comes close to FAR with tasks that can exploit almost all GPU resources, but it is an unremarkable case because the interest in using MIG arises from the inability of tasks to fully exploit the available resources. The cost of FAR is slightly higher than the other proposals, but the scheduling improvements make up for it. In fact, FAR is still very affordable even with many tasks.

6.6. Refinement contribution

Another aspect to assess is the improvement provided by refinement (phase 3 of FAR) in relation to its cost. It is intended to confirm that the movement/swapping of tasks raised for this phase has impact on the result without running too much operations.

Table 6 reports the percentage of improvement

$$p_{\text{ref.}} = (\omega_{\text{no ref.}} / \omega_{\text{ref.}} - 1) \cdot 100,$$

between the makespan without and with refinement for some of the workloads used before, along with the number of moves and swaps made until reaching an iteration where the method cannot improve.

It can be observed in all cases a rather moderate number of moves and task swaps: at most around 1 move or 4 task swaps on average per refinement. As introduced in Section 3.3, the worst cost for a movement is $\mathcal{O}(\log(n))$ and for a swap is $\mathcal{O}(n)$, so it can be intuited that the executions are not costly. In fact, the moderate costs for FAR that we have reported in the previous sections, even for a large number of tasks, include the refinement phase. Nevertheless, significant improvements over scheduling without refinement are obtained, especially with $n = 20$ tasks, where the percentage $p_{\text{ref.}}$ shows more than 10% of improvement, but also with $n = 30$ tasks, where $p_{\text{ref.}}$ drops slightly but also stays around that improvement value. For its part, with only $n = 10$ tasks the refinement practically does not reduce the makespan, which may also explain

Table 5
Average of σ_A in 1000 executions with $n = 15$ for A100.

Task configuration	Approach			
	MISO-OPT	FIXPART(1,...,1)	FIXPARTBEST	FIXPART(7)
POORSCALING, NARROWTIMES	1.19	1.25	1.24	3.29
POORSCALING, WIDETIMES	1.55	1.29	1.22	3.39
MIXEDSCALING, NARROWTIMES	1.62	1.39	1.13	2.17
MIXEDSCALING, WIDETIMES	2.03	1.47	1.09	2.16
GOODSCALING, NARROWTIMES	1.83	1.61	1.00	1.31
GOODSCALING, WIDETIMES	2.14	1.78	1.01	1.28

Table 6
Average of p_{ref} (%), along with the average number of movements and swaps operations performed in 1000 executions of refinement for A100.

Task configuration	Number of tasks (n)								
	10			20			30		
	p_{ref}	moves	swaps	p_{ref}	moves	swaps	p_{ref}	moves	swaps
POORSCALING, NARROWTIMES	0.31	0.00	0.02	13.15	0.00	3.25	11.45	0.00	4.26
POORSCALING, WIDETIMES	0.28	0.02	0.01	14.98	1.26	0.13	8.76	1.74	0.15
MIXEDSCALING, NARROWTIMES	0.76	0.02	0.06	13.87	0.00	3.23	9.04	0.00	4.31
MIXEDSCALING, WIDETIMES	3.21	0.17	0.06	11.45	1.14	0.09	9.01	1.84	0.15
GOODSCALING, NARROWTIMES	0.78	0.01	0.01	13.44	0.00	3.01	7.54	0.00	3.64
GOODSCALING, WIDETIMES	1.34	0.05	0.02	12.56	1.13	0.08	9.32	1.93	0.21

the slightly higher errors we obtained for that number of tasks in Table 4 with respect to the ideal baseline. The conclusion is that refinement significantly and efficiently improves task scheduling for not very large batches, but if the number of tasks is too small it probably does not contribute anything. No intermediate or larger n values have been included for brevity, but the trends are progressively maintained and stabilize at around $n = 30$.

The different ranges for task duration in Table 6 (NARROWTIMES and WIDETIMES) highlight the importance of considering both refinement operations: movement and swapping. It is observed that for a small range of long tasks (NARROWTIMES) the operation that allows improvement is exclusively task swapping, but for a wide range including both long tasks and some short task (WIDETIMES) the improvement comes primarily from movements. This is logical because with all long tasks it is very difficult to have idle time until the makespan of an instance to relocate a task of such duration (movement operation), but it will be likely to make a trade-off by gaining the difference in duration between the swapped tasks. When the duration of the tasks is heterogeneous, there are large idle gaps up to the makespan with respect to the duration of other tasks, which can be moved. On the other hand, the different scaling capabilities do not show great influence on the performance of the refinement, and simply corroborate their usefulness for diverse workloads.

6.7. Multi-batch evaluation

Finally, we evaluate the performance of our proposal for many task batches. We first focus on the contribution of the batch concatenation step presented in Section 4, and then on evaluating the algorithm as a whole for multi-batch task loads.

6.7.1. Improved concatenation contribution

To evaluate the contribution of concatenation with overlap we take as a reference the trivial concatenation consisting of starting the scheduling of a batch just after the last task of the previous one. The improvement over it is disaggregated according to the contribution of batch reversal, and task move/swap operations performed after.

Table 7 reports the improvement percentage

$$p_{rev} = (\omega_{trivial\ concat.} / \omega_{rev.\ concat.} - 1) \cdot 100,$$

between the makespan of the trivial concatenation and that obtained after task reversal, together with the percentage

$$p_{move/swap} = (\omega_{trivial\ concat.} / \omega_{move/swap.\ concat.} - 1) \cdot 100,$$

between trivial concatenation makespan and the reversed one with also task moving/swapping.

For each workload, a concatenation of 1001 different schedules produced by FAR has been executed so that each one represents a batch. Therefore, we have 500 concatenations of a batch B_k with a batch B_{k+1} reverse and 500 concatenations where B_{k+1} is not reverse as the one that is reversed is B_k . The mean improvement percentages are reported aggregated since no great difference was observed between the results of each type: practically the same when B_k is the reversed one as when it is B_{k+1} . Moreover, in a real scenario, the number of tasks in each batch would be different, but for dimensionality issues and due to the analytical purposes of the experiment, all concatenated batches in each workload have the same number of tasks.

Table 7 clearly shows that enhanced concatenation is especially useful versus trivial concatenation when batches have few tasks: $p_{move/swap}$ around 15% improvement with $n = 10$ tasks, which decays to just 4% with $n = 20$ tasks and disappears with $n = 30$ tasks. This is logical given that: (1) the scheduling of each batch is closer to the ideal (perfectly balanced block) as its number of tasks increases and, therefore, a more basic concatenation has little room for improvement in such a case; (2) the relative weight in the total time of a move/swap of unbalanced tasks at the end of the schedule is lower when there are more tasks. In any case, the results show good performance for batches where FAR has more error, i.e., those with very few tasks, and augur that the concatenation technique can mitigate it in the combination of schedules. In addition, we observe that the move/swap operations are much more important for improvement than the reversion of the batch to concatenate: p_{rev} , 2-5% of improvement with $n = 10$, but $p_{move/swap}$ shows more than 10% extra improvement than only reversing. However, as we saw in Section 4, the reversion has practically no computational cost, and moving/swapping has the same as we discussed in the refinement ($\mathcal{O}(\log(n))$ the move and $\mathcal{O}(n)$ the swap). The rest of the task characteristics do not seem to have a clear influence on the results, and simply reaffirm the usefulness of the proposed concatenation for different scenarios.

Table 7
Average of $p_{rev.}$ (%) and $p_{move/swap}$ (%) in the 1000 concatenations of 1001 FAR schedules for A100.

Task configuration	Number of tasks (n)					
	10		20		30	
	$p_{rev.}$	$p_{move/swap}$	$p_{rev.}$	$p_{move/swap}$	$p_{rev.}$	$p_{move/swap}$
POORSCALING, NARROWTIMES	2.32	16.19	2.64	4.60	1.03	1.45
POORSCALING, WIDETIMES	3.23	14.47	3.82	4.10	0.98	1.01
MIXEDSCALING, NARROWTIMES	5.43	16.22	4.53	4.82	0.86	1.01
MIXEDSCALING, WIDETIMES	4.87	14.30	2.41	4.26	0.44	0.46
GOODSCALING, NARROWTIMES	3.65	15.43	1.05	3.86	0.51	1.01
GOODSCALING, WIDETIMES	5.54	13.12	2.03	3.76	0.23	0.30

Table 8
Average number of movements or swaps in the 1000 concatenations of 1001 FAR schedules for A100.

Task configuration	Number of tasks (n)					
	10		20		30	
	Moves	Swaps	Moves	Swaps	Moves	Swaps
POORSCALING, NARROWTIMES	0.48	9.72	0.04	1.32	0.01	0.82
POORSCALING, WIDETIMES	3.13	1.21	0.10	0.05	0.03	0.00
MIXEDSCALING, NARROWTIMES	0.30	9.52	0.02	1.12	0.00	0.62
MIXEDSCALING, WIDETIMES	3.51	1.12	0.13	0.06	0.05	0.01
GOODSCALING, NARROWTIMES	0.72	9.62	0.03	1.36	0.01	0.94
GOODSCALING, WIDETIMES	3.42	1.04	0.15	0.06	0.05	0.01

We next analyze the cost and the particular utility of task movement or swapping, by the number of operations of each type that have been performed with each workload, as shown in Table 8. Again, it is the number of operations until an iteration where none can be performed to improve.

As for the refinement, Table 8 shows that concatenating batches of homogeneous times (NARROWTIMES) can only be improved in practice by task swaps, while concatenating batches of more varied duration (WIDETIMES) improves mainly by task movement. We also notice more operations on average until no further improvement is possible compared to the refinement phase (see Table 6), but it still does not represent an excessive cost. The average time to compute each concatenation in this experiment was less than half a millisecond for $n = 30$ tasks. Increasing the concatenation to batches of $n = 1000$ tasks, an average time of 23.76 ms was obtained to compute each concatenation. In any case, we reiterate the option of limiting the number of operations when their improvement is low or too many have already been performed.

6.7.2. Multi-batch performance

Finally, to gauge the overall error in multi-batch executions, we report a factor based on the competitive ratio, i.e., the maximum possible ratio between the makespan achieved by the online scheduler and the optimal offline makespan. Again, since it is not possible to compute the optimum in a reasonable time, we use the lower bound consisting in assuming that the minimum possible area of the tasks is evenly distributed among the slices⁹:

$$baseline_{multi\text{-}batch} = \sum_{B \in \text{Batches}} \sum_{T_i \in B} \min_{s \in C_G} (s \cdot t_i^B(s)) / \#slices_G.$$

That is, we report the improvement percentage:

$$p_{multi\text{-}batch} = (\omega_{multi\text{-}batch} / baseline_{multi\text{-}batch} - 1) \cdot 100.$$

Table 9 reports $p_{multi\text{-}batch}$ for the same workloads used to obtain the ρ ratios for a single batch in Table 4. We use tens of thousands of

Table 9
Value of $p_{multi\text{-}batch}$ (%) with WIDETIMES and 1001 batches for A100.

Task configuration	Number of tasks (n)					
	10	15	20	25	30	35
POORSCALING	84.42	94.12	95.21	95.21	92.32	93.21
MIXEDSCALING	89.56	94.00	93.01	95.01	90.21	92.03
GOODSCALING	82.67	94.53	94.46	94.32	92.32	91.54

tasks corresponding to 1001 batches to concatenate. Again, we report only the results for WIDETIMES, but the result of this experiment for NARROWTIMES is similar.

By simply concatenating the batches we could expect up to twice the approximation ratio for single batch [21], which would mean a percentage of error between 102-116% for $n \geq 15$ (see Table 4). However, the multi-batch results in Table 9 are 10-20% below those values. For $n = 10$, the results are up to 50% below twice the single-batch error, as $baseline$ is not very tight for so few tasks. In contrast, $baseline_{multi\text{-}batch}$ does not exhibit this issue as it is calculated with the tasks from all batches. On the other hand, for larger numbers of tasks, the least tightly fitting may be $baseline_{multi\text{-}batch}$; the simplification made by both lower bounds comes from assuming that all tasks can be perfectly balanced with minimum area and no idle periods, but this incurs a larger error in multi-batch since it is assumed to be offline, but in reality tasks are received online. In any case, these results suggest good performance for multi-batch, although it is already justified by the tighter results reported for single-batch in Section 6.4.2, and the improvement analysis with respect to trivial concatenation in Section 6.7.1.

The scheduling cost is basically the sum of the times we have reported for FAR in Section 6.4.2, and for the improved concatenation in Section 6.7.1. As we have already discussed in those sections, they are good enough to support the efficiency of the proposal.

7. Conclusions

In this paper we have motivated the opportunities of MIG by proposing a moldable scheduling problem to leverage this technology for task co-scheduling. Beyond presenting a complete scheduling proposal, accompanied by a useful C++ implementation on currently MIG-capable

⁹ For brevity, in the single-batch scheduling context, the batch has been omitted from the notation of execution times. In the multi-batch context we use $t_i^B(s)$ to denote the execution time of the i -th task of batch B using s slices.

GPUs [8], this work highlights a promising line of research that can be greatly enriched by other scheduling approaches in the future. In addition, it suggests many metrics, workload characterizations, and evaluation techniques that may be useful for such future work.

In this work, we have considered a scenario with long running tasks in which batches of tasks are accumulated during the executions. For scheduling a batch we have presented FAR, an efficient three-stage algorithm that combines several heuristics with an approximation factor of 1.75 for NVIDIA A30 GPU and 2 for the A100/H100 GPUs. In addition, we propose a method to concatenate the FAR batch solutions, trying to overlap their executions to improve their combined makespan.

Experimentally we have observed good solutions for real kernels profiled on the A100, with an error of less than 22% of the optimum for a batch. In addition, we have obtained errors of less than 10% on synthetic inputs that mimic real behaviors and cover a wide variety of workloads. These experiments have allowed us to verify that the overhead of MIG reconfiguration is not excessive if the scheduled tasks exhibit long durations, and in that case it is not necessary to schedule many tasks per batch. We observed a significant improvement over the only previous state-of-the-art approach (MISO-OPT) and over a simple proposal which does not use dynamic MIG reconfiguration (FIXPART).

The third phase of FAR significantly refines the solution for some batches, with improvements of around 15% of the makespan for batches of 20 tasks. These improvements come from moving tasks when execution times are heterogeneous, and from swapping tasks when execution times are homogeneous. In addition, the proposed concatenation with overlap significantly improves multi-batch scheduling for batches with few tasks, where FAR is a bit worse, achieving quite good overall performance and similar to larger batches.

For the future, we plan to leverage MIG in scenarios where tasks do not accumulate, for example because they are much shorter, and need to be scheduled *online*. We also consider extending the study for tasks with precedence constraints, deadlines, or different priorities. In some of these cases, GPU reconfiguration times may be more relevant and one of the challenge is to make the scheduler aware of them to avoid unhelpful reconfigurations.

CRedit authorship contribution statement

Jorge Villarrubia: Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Luis Costero:** Writing – review & editing, Supervision, Formal analysis, Conceptualization. **Francisco D. Igual:** Writing – review & editing, Supervision, Resources, Investigation, Conceptualization. **Katzalin Olcoz:** Writing – review & editing, Supervision, Investigation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Jorge Villarrubia, Luis Costero, Francisco D. Igual, Katzalin Olcoz reports financial support was provided by Spain Ministry of Science and Innovation (Grant PID2021-126576NB-I00).

Acknowledgments

This work is funded by Grant PID2021-126576NB-I00 funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe”. We thank the HPC&A group at Universitat Jaume I de Castellón for granting us access to A100 and H100 GPUs for profiling purposes.

Data availability

No data was used for the research described in the article.


References

- [1] NVIDIA, Multi-process service, <https://docs.nvidia.com/deploy/mps/>.
- [2] S. Che, M. Boyer, J. Meng, D. Tarjan, J.W. Sheaffer, S.-H. Lee, K. Skadron, Rodinia: a benchmark suite for heterogeneous computing, in: 2009 IEEE International Symposium on Workload Characterization (IISWC), 2009, pp. 44–54.
- [3] M. Amarís, R. Camargo, M. Dyab, A. Goldman, D. Trystram, A comparison of gpu execution time prediction using machine learning and analytical modeling, in: 2016 IEEE 15th International Symposium on Network Computing and Applications (NCA), 2016, pp. 326–333.
- [4] B. Johnston, G. Falzon, J. Milthorpe, Opencl performance prediction using architecture-independent features, in: 2018 International Conference on High Performance Computing & Simulation (HPCS), 2018, pp. 561–569.
- [5] B. Li, T. Patel, S. Samsi, V. Gadepally, D. Tiwari, MISO: exploiting multi-instance GPU capability on multi-tenant GPU clusters, in: Proceedings of the 13th Symposium on Cloud Computing, 2022, pp. 173–189.
- [6] J. Turek, J.L. Wolf, P.S. Yu, Approximate algorithms scheduling parallelizable tasks, in: Proceedings of the Fourth Annual ACM Symposium on Parallel Algorithms and Architectures, 1992, pp. 323–332.
- [7] R.L. Graham, Bounds on multiprocessing timing anomalies, SIAM J. Appl. Math. 17 (2) (1969) 416–429, <https://doi.org/10.1137/0117039>.
- [8] J. Villarrubia, Code repository: FAR MIG scheduler, https://github.com/artecsgroup/FAR_MIG_scheduler, 2025.
- [9] M. Drozdowski, Scheduling multiprocessor tasks — an overview, Eur. J. Oper. Res. 94 (2) (1996) 215–230, [https://doi.org/10.1016/0377-2217\(96\)00123-3](https://doi.org/10.1016/0377-2217(96)00123-3).
- [10] D.G. Feitelson, L. Rudolph, U. Schwiegelshohn, K.C. Sevcik, P. Wong, Theory and practice in parallel job scheduling, in: Job Scheduling Strategies for Parallel Processing, Springer Berlin Heidelberg, Berlin, Heidelberg, 1997, pp. 1–34.
- [11] E. Bampis, M. Caramia, J. Fiala, A.V. Fishkin, A. Iovanello, Scheduling of independent dedicated multiprocessor tasks, in: Algorithms and Computation, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 391–402.
- [12] K. Jansen, A $(3/2+\epsilon)$ approximation algorithm for scheduling moldable and non-moldable parallel tasks, in: Twenty-Fourth SPPA, 2012, pp. 224–235.
- [13] G. Mounie, C. Rapine, D. Trystram, A $3/2$ -approximation algorithm for scheduling independent monotonic malleable tasks, SIAM J. Comput. 37 (2) (2007) 401–412, <https://doi.org/10.1137/S0097539701385995>.
- [14] X. Wu, P. Loiseau, Efficient approximation algorithms for scheduling moldable tasks, Eur. J. Oper. Res. 310 (1) (2023) 71–83, <https://doi.org/10.1016/j.ejor.2023.02.044>.
- [15] J. Blazewicz, M. Kovalyov, M. Machowiak, D. Trystram, J. Weglarz, Preemptible malleable task scheduling problem, IEEE Trans. Comput. 55 (4) (2006) 486–490, <https://doi.org/10.1109/TC.2006.58>.
- [16] L. Costero, F.D. Igual, K. Olcoz, Dynamic power budget redistribution under a power cap on multi-application environments, Sustain. Comput. Inf. Syst. 38 (2023) 100865, <https://doi.org/10.1016/j.suscom.2023.100865>.
- [17] L. Marchal, B. Simon, O. Sinnen, F. Vivien, Malleable task-graph scheduling with a practical speed-up model, IEEE Trans. Parallel Distrib. Syst. 29 (6) (2018) 1357–1370, <https://doi.org/10.1109/TPDS.2018.2793886>.
- [18] T. Abdelzaher, K. Shin, Combined task and message scheduling in distributed real-time systems, IEEE Trans. Parallel Distrib. Syst. 10 (11) (1999) 1179–1191, <https://doi.org/10.1109/71.809575>.
- [19] J. Blazewicz, M. Machowiak, G. Mounié, D. Trystram, Approximation algorithms for scheduling independent malleable tasks, in: Euro-Par 2001 Parallel Processing, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 191–197.
- [20] L. Perotin, S. Kandaswamy, H. Sun, P. Raghavan, Multi-resource scheduling of moldable workflows, J. Parallel Distrib. Comput. 184 (2024) 104792, <https://doi.org/10.1016/j.jpdc.2023.104792>.
- [21] D.B. Shmoys, J. Wein, D.P. Williamson, Scheduling parallel machines online, SIAM J. Comput. 24 (6) (1995) 1313–1331, <https://doi.org/10.1137/S0097539793248317>.
- [22] D. Ye, D.Z. Chen, G. Zhang, Online scheduling of moldable parallel tasks, J. Sched. 21 (6) (2018) 647–654, <https://doi.org/10.1007/s10951-018-0556-2>.
- [23] R. Guerreiro, A.S. Santos, A. Tereso, Online scheduling: a survey, in: 2023 18th Iberian Conference on Information Systems and Technologies (CISTI), 2023, pp. 1–6.
- [24] A. Allahverdi, A survey of scheduling problems with uncertain interval/bounded processing/setup times, J. Proj. Manag. 7 (4) (2022) 255–264, <https://doi.org/10.5267/j.jpjpm.2022.3.003>.
- [25] A. Allahverdi, The third comprehensive survey on scheduling problems with setup times/costs, Eur. J. Oper. Res. 246 (2) (2015) 345–378, <https://doi.org/10.1016/j.ejor.2015.04.004>.
- [26] L.R. Abreu, J.O. Cunha, B.A. Prata, J.M. Framinan, A genetic algorithm for scheduling open shops with sequence-dependent setup times, Comput. Oper. Res. 113 (2020) 104793, <https://doi.org/10.1016/j.cor.2019.104793>.
- [27] M.D. Toksari, G. Toğa, Single batch processing machine scheduling with sequence-dependent setup times and multi-material parts in additive manufacturing, CIRP J. Manuf. Sci. Technol. 37 (2022) 302–311, <https://doi.org/10.1016/j.cirpj.2022.02.007>.
- [28] P.-F. Dutot, L. Eyraud, G. Mounié, D. Trystram, Bi-criteria algorithm for scheduling jobs on cluster platforms, in: Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, 2004, pp. 125–132.

- [29] R. Sudarsan, C.J. Ribbens, Combining performance and priority for scheduling resizable parallel applications, *J. Parallel Distrib. Comput.* 87 (2016) 55–66, <https://doi.org/10.1016/j.jpdc.2015.09.007>.
- [30] A. Benoit, L. Perotin, Y. Robert, H. Sun, Online scheduling of moldable task graphs under common speedup models, in: Proceedings of the 51st International Conference on Parallel Processing, ICPP '22, Association for Computing Machinery, New York, NY, USA, 2023.
- [31] L. Perotin, H. Sun, P. Raghavan, Multi-resource list scheduling of moldable parallel jobs under precedence constraints, in: Proceedings of the 50th International Conference on Parallel Processing, 2021, pp. 1–10.
- [32] B.S. Baker, E.G. Coffman Jr., R.L. Rivest, Orthogonal packings in two dimensions, *SIAM J. Comput.* 9 (4) (1980) 846–855, <https://doi.org/10.1137/0209064>.
- [33] J. Du, J.Y.-T. Leung, Complexity of scheduling parallel task systems, *SIAM J. Discrete Math.* 2 (4) (1989) 473–487, <https://doi.org/10.1137/0402042>.
- [34] A. Bortfeldt, A genetic algorithm for the two-dimensional strip packing problem with rectangular pieces, *Eur. J. Oper. Res.* 172 (3) (2006) 814–837, <https://doi.org/10.1016/j.ejor.2004.11.016>.
- [35] G. Gómez-Villouta, J.-P. Hamiez, J.-K. Hao, Tabu search with consistent neighbourhood for strip packing, in: Trends in Applied Intelligent Systems, Springer, Berlin Heidelberg, 2010, pp. 1–10.
- [36] C.A. Floudas, X. Lin, Mixed integer linear programming in process scheduling: modeling, algorithms, and applications, *Ann. Oper. Res.* 139 (2005) 131–162, <https://doi.org/10.1007/s10479-005-3446-x>.
- [37] K.M. Tarplee, R. Friese, A.A. Maciejewski, H.J. Siegel, Scalable linear programming based resource allocation for makespan minimization in heterogeneous computing systems, *J. Parallel Distrib. Comput.* 84 (2015) 76–86, <https://doi.org/10.1016/j.jpdc.2015.07.002>.
- [38] J. Fang, Y. Rao, X. Zhao, B. Du, A hybrid reinforcement learning algorithm for 2d irregular packing problems, *Mathematics* 11 (2) (2023), <https://doi.org/10.3390/math11020327>.
- [39] R. Harren, R. van Stee, Improved absolute approximation ratios for two-dimensional packing problems, in: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 2009, pp. 177–189.
- [40] A. Steinberg, A strip-packing algorithm with absolute performance bound 2, *SIAM J. Comput.* 26 (2) (1997) 401–409, <https://doi.org/10.1137/S0097539793255801>.
- [41] K. Jansen, F. Land, Scheduling monotone moldable jobs in linear time, in: 2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS), IEEE, 2018, pp. 172–181.
- [42] A.K. Amoura, E. Bampis, C. Kenyon, Y. Manoussakis, Scheduling independent multi-processor tasks, *Algorithmica* 32 (2002) 247–261, <https://doi.org/10.1007/s00453-001-0076-9>.
- [43] K. Jansen, Scheduling malleable parallel tasks: an asymptotic fully polynomial-time approximation scheme, in: 10th Annual ESA, 2002, pp. 562–573.
- [44] K. Jansen, R. Thöle, Approximation algorithms for scheduling parallel jobs, *SIAM J. Comput.* 39 (8) (2010) 3571–3615, <https://doi.org/10.1137/080736491>.
- [45] S. Hunold, One step toward bridging the gap between theory and practice in moldable task scheduling with precedence constraints, *Concurr. Comput., Pract. Exp.* 27 (4) (2015) 1010–1026, <https://doi.org/10.1002/cpe.3372>.
- [46] C. Tan, Z. Li, J. Zhang, Y. Cao, S. Qi, Z. Liu, Y. Zhu, C. Guo, Serving DNN models with multi-instance GPUs: a case of the reconfigurable machine scheduling problem, <https://doi.org/10.48550/arXiv.2109.11067>, 2021.
- [47] J. Qi, W. Xiao, M. Li, C. Yang, Y. Li, W. Lin, H. Yang, Z. Luan, D. Qian, Elasticbatch: a learning-augmented elastic scheduling system for batch inference on mig, *IEEE*

Trans. Parallel Distrib. Syst. 35 (10) (2024) 1708–1720, <https://doi.org/10.1109/TPDS.2024.3431189>.

- [48] X. Wei, Z. Li, C. Tan, Optimizing gpu sharing for container-based dnn serving with multi-instance gpus, in: Proceedings of the 17th ACM International Systems and Storage Conference, SYSTOR '24, Association for Computing Machinery, New York, NY, USA, 2024, pp. 68–82.
- [49] B. Zhang, S. Li, Z. Li, Miger: integrating multi-instance gpu and multi-process service for deep learning clusters, in: Proceedings of the 53rd International Conference on Parallel Processing, ICPP '24, Association for Computing Machinery, New York, NY, USA, 2024, pp. 504–513.



Jorge Villarrubia received a BSc Degree in Computer Science and a BSc Degree in Mathematics in 2021 from the Complutense University of Madrid (UCM). In 2023 he obtained a MSc Degree in Computer Science at UCM. Currently he is pursuing a Ph.D. in Computer Engineering at UCM. Since 2023, he is an assistant professor with the Department of Computer Architecture and System Engineering. His main research areas include resource management and high performance computing.



Luis Costero received the PhD degree in Computer Sciences from the Complutense University of Madrid (UCM) in 2021. In 2021, he joined the Embedded Systems Laboratory (ESL) at EPFL as postdoctoral researcher. Since 2022, he is an assistant professor with the Department of Computer Architecture and System Engineering, Complutense University. His main research areas involve high performance computing, heterogeneous architectures, power consumption and real-time resource management.



Francisco D. Igual obtained a bachelor degree in Computer Engineering from University Jaume I de Castellón (Spain) in 2006, and a Ph.D. degree in Computer Science from the same University in 2011. In 2011, he moved as a postdoctoral researcher to The University of Texas at Austin, and in 2012 he joined the Universidad Complutense de Madrid where he is currently Associate Professor. His research interests include high-performance and energy-aware computing, dense linear algebra library development and optimization, and runtime task scheduling on massively heterogeneous architectures.



Katzalin Olcoz received the PhD degree in physics from the Complutense University (UCM) of Madrid, in 1997. She is an associate professor with the Department of Computer Architecture and System Engineering, Complutense University. From 2012 to 2016 she served as head of the Department of Computer Architecture and System Engineering of the same university. She was a visiting professor at EPFL (Lausanne, Switzerland) from April to June, 2018. Her research interests include high-performance and energy-aware computing, resource management, and energy efficiency and virtualization. Within the computer architecture group of the Complutense University, she has been involved in several projects in the field of computer architecture and design automation from high-level specifications.