

Exploring the Relation Between Two Levels of Scheduling Using a Novel Simulation Approach

Ahmed Eleliemy, Ali Mohammed, and Florina M. Ciorba

Department of Mathematics and Computer Science

University of Basel

Basel, Switzerland

Email: {ahmed.eleliemy, ali.mohammed, florina.ciorba}@unibas.ch

Abstract—Modern high performance computing (HPC) systems exhibit a rapid growth in size, both “horizontally” in the number of nodes, as well as “vertically” in the number of cores per node. As such, they offer additional levels of hardware parallelism. Each level requires and employs algorithms for appropriately scheduling the computational work at the respective level. The present work explores the relation between two scheduling levels: batch and application. To understand and explore this relation, a novel simulation approach is presented that bridges two existing simulators from the two scheduling levels. A novel two-level simulator that implements the proposed approach is introduced. The two-level simulator is used to simulate all combinations of three batch scheduling and four application scheduling algorithms from the literature. These combinations are considered for allocating resources and executing the parallel jobs from a workload of a production HPC system. The results of the scheduling experiments reveal the strong relation between decisions taken at the two scheduling levels and their mutual influence. Complementing the simulations, the two-level simulator produces abstract parallel execution traces, which can visually be examined and illustrate the execution of different jobs and, for each job, the execution of its tasks at node and core levels, respectively.

Keywords—High performance computing; Batch level scheduling; Application level scheduling; Two-level scheduling; GridSim; Alea; SimGrid; SimDag; OTF2; Vampir.

I. INTRODUCTION

Maximizing cluster utilization, minimizing job execution time and job waiting time are key optimization metrics for successful management and operation of modern high performance computing (HPC) clusters [1]. Scheduling of different jobs at the cluster level affects the cluster utilization and the job waiting time. Scheduling of different tasks within an application affects its execution time. Scheduling is known to be an NP-Complete [2] problem. Therefore, in the literature, a large number of heuristics have been proposed for job and application scheduling. Each of these heuristics has its own advantages in terms of the achieved performance. The state-of-the-art job and application scheduling algorithms are no longer sufficient to exploit *individually* all available parallelism of modern and future HPC systems. Modern (and future) HPC systems (will) exhibit massive parallelism at different hardware levels. Each of these hardware parallelism levels uses a corresponding scheduling technique that manages and schedules its computational load. For instance, batch, appli-

cation, and thread level scheduling techniques exist at the cluster, node, and core levels, respectively. It is important to examine the relation between the different scheduling levels to take full advantage of the different hardware and software parallelism levels of modern HPC systems. For instance, making the individual schedulers at the batch and application levels aware of each other’s decisions enables them to work *in concert* for an optimized execution of applications and improved utilization of the underlying resources. The focus of this work is on exploiting multi-level parallelism through scheduling [3].

Simulation-based approaches are widely used in the literature for studying various aspects of single level scheduling. Yet, to the best of our knowledge, no simulators exist that support the study of two levels of scheduling. The main contributions of this work are: (1) Introduces a novel generic simulation approach by bridging two existing simulators from batch and application scheduling levels; (2) Develops and evaluates a new two-level scheduling simulator based on GridSim [4] and SimGrid [5]; and (3) The novel two-level simulator generates standard open trace format (OTF2) [6] traces that allow visualization of the execution of jobs and tasks on allocated resources.

In this work, combinations of three BLS algorithms, First Come First Serve (FCFS), Earliest Deadline First (EDF), and Shortest Job First (SJF) [7] and four ALS algorithms, static chunking (STATIC), self scheduling (SS) [8], guided self scheduling (GSS) [9], and factoring (FAC) [10], are performed to explore the relation between BLS and ALS during operation. These combinations are used to schedule two real workloads from HPC production systems. Certain customizations are applied to these batch workloads to envision the missing information that is necessary to simulate ALS. The results which are presented in this paper cover only one workload while the complete set of the results is available online¹.

This paper is structured as follows: In Section II, the most relevant work in the literature is discussed. In Section III-A, a generic simulation approach for connecting two simulators of different scheduling levels is introduced. Based on this approach, in Section III-B, a new two-level simulator for

¹<https://drive.switch.ch/index.php/s/tAIBHGWiWW1vyhZ>

BLS and ALS is proposed. In Section IV, the results of executing the workload of a production HPC system with several combinations of BLS and ALS algorithms using the proposed two-level simulator are discussed. The information needed to reproduce this work is presented in Section V. Finally, in Section VI, the conclusions of this work and its potential future work are outlined.

II. RELATED WORK

Implementing, comparing, verifying, and validating a scheduling solution for HPC production systems involves numerous technical details and steps. Therefore, simulation approaches have been widely used to examine the performance of different scheduling algorithms. Consequently, the crux of this section is to highlight the most relevant and influential simulators used for simulating scheduling of problems at the levels considered herein: batch and application scheduling, respectively.

Simulation of BLS algorithms: One popular simulation toolkit is GridSim [4]. It facilitates simulation of grids, clusters, and single processing elements. It offers support for a broad range of heterogeneous resources including shared and distributed memory architectures. GridSim is built on top of a reliable discrete event simulation library called SimJava [11]. The GridSim toolkit is fully implemented in Java which promotes its portability and extensibility.

Simulation of ALS algorithms: SimGrid [5] is a widely used simulation toolkit for ALS. SimGrid supports the development of parallel and distributed applications in heterogeneous/homogeneous parallel and distributed environments. Recent releases of SimGrid have three different interfaces: MetaSimGrid (MSG), SimDag (SD), and Simulated MPI (SMPI) for simulating applications as a group of concurrent processes, applications as directed acyclic graphs (DAGs), and for running unmodified applications written using the message passing interface (MPI), respectively.

Due to the reliability and the active support community of the two toolkits, GridSim and SimGrid, they are used to build a large number of notable simulators, such as Alea [12] and Batsim [13]. Alea and Batsim are the most relevant simulators to the current work, and are used for the BLS and ALS simulations, respectively.

Alea [12] [14] is a well-known simulator, developed on top of the GridSim toolkit. It extends GridSim and improves certain of its limitations. For instance, it supports reading of jobs files written in the standard workload format (SWF) [15] and reading of platform files instead of defining the simulated platform within the code. Alea provides a group of data structures that enhance job and resource modeling. It also implements a set of scheduling algorithms and gives the opportunity to integrate other scheduling algorithms. It can be considered a suitable candidate for BLS simulations. In the present work, simulations at the BLS level are performed with Alea.

Batsim [13] is one of the most recent SimGrid-based simulators. It is based on the separation of concerns between system simulation and scheduling algorithms using two main components: *batsim main* and *batsim decision*. The main component is responsible for simulating the computational resources and it uses the SimGrid simulation toolkit underneath. The decision component is responsible for the scheduling decisions at the resources management level and it can be implemented in any programming language. Batsim depends on a Unix socket layer to allow communication between its two components. An approach similar to the Batsim communication approach that depends on Unix sockets is used in this present work.

The GridSim and SimGrid toolkits are preferably used (not restricted) to support batch and application level scheduling, respectively. Certain research efforts are described below that studied extensions of one of these two simulation toolkits to support the simulation of other scheduling levels.

ALS simulation based on a BLS simulation toolkit: In [16], the authors prove the ability of extending Alea to support ALS algorithms. The authors extended Alea to support ALS algorithms, such as STATIC, SS, GSS, and FAC. The work in [16] carried over all the Alea's advantages to the ALS domain, such as application tasks being expressed in the SWF format and the effect of system failures being examined with different ALS techniques. However, the simulator provided by [16] supports ALS algorithms in such a way that it can no longer support BLS algorithms.

BLS simulation based on an ALS simulation toolkit: Simbatch [17] is a SimGrid-based simulator. Simbatch uses the MSG interface of SimGrid to support simulations and development of BLS algorithms. Simbatch's uniqueness comes from the fact that it swaps the focus of SimGrid from the ALS perspective to the BLS perspective.

All aforementioned simulators and simulation toolkits are designed to support *single level* scheduling simulations, such as at the BLS or ALS. However, to explore the relation between multiple levels of scheduling, simulators are needed that can combine the required methods, tools, and techniques from the single levels. In [18], the concept of combining resource allocation (RA) with dynamic loop scheduling (DLS) was first proposed, under the name of "CDS", a combined dual-stage RA and DLS approach. CDS is a two-stage framework that maximizes the probability that applications complete by common deadline under certain levels of variation in the resources availability to compute and variation in system input. The RA techniques used in [18] were initially simplistic and, subsequently, more sophisticated in [19].

As a preliminary step for the work in the present paper, the original Alea simulator [12] has been *redesigned* and *reimplemented* to support ALS algorithms in addition to BLS algorithms, in [3]. Moreover, in [3] a new simulator based on the SimGrid-SD interface [5] was designed to support BLS algorithms in addition to ALS algorithms. These two simulators showed similar results in terms of total execution time for the simulated batches and applications at BLS and ALS, respectively. In the case of large batch workloads, the extended

Alea simulator showed improved performance in terms of total simulation time, while the SimGrid-SD-based simulator showed improved performance in the case of applications that contain large numbers of tasks. The simulators presented in [3] support only single level scheduling: BLS or ALS.

Attempting to simultaneously simulate the two levels of scheduling using the simulators in [3] revealed certain technical challenges. In particular, both simulators are based on discrete events. Each maintains an individual simulation clock, updated according to the events occurring at the scheduling level they simulate. Extending any of the two simulators to support multiple simulation clocks involves numerous changes that may affect the functionality of simulation toolkit used to build that simulator. Another challenge is related to the initialization functions `GridSim.init` and `SD_init` of `GridSim` and `SimGrid-SD`, respectively. These functions were designed to be called one time at the beginning of the simulation. Thus, multiple calls would cause simulation errors during execution. In certain cases, the `SimGrid-SD`-based simulator would require multiple calls to `SD_init` to reinitialize the simulation environment, when the simulator is launched for a different application.

In the present work, a novel simulation approach is proposed to develop two-level simulators. The proposed approach overcomes the challenges encountered in [3] by bridging scheduling simulators from two levels, in such a way that each remains responsible for simulating a specific scheduling algorithm at a certain level of hardware and software parallelism. Bridging simulators according to the proposed approach aims to minimize the changes in the single-level simulation toolkits. Minimal changes in the simulation toolkits constitutes an advantage for obtaining the support of the user community and to naturally maintain compatibility with the new toolkit releases.

III. THE PROPOSED TWO-LEVEL SIMULATION APPROACH

As discussed in Section II, numerous simulators exist for scheduling in parallel and distributed systems. Each simulator has its own capabilities and limitations. It becomes necessary to define a set of objectives for the proposed two-level simulation approach, to crystallize its capabilities and limitations.

Objectives of the proposed approach

- (i) Preserve the current level of user involvement, to avoid that users learn new APIs or new simulation toolkits to perform their simulations.
- (ii) The scheduling algorithms in the literature, at either of the BLS or ALS levels, can be easily ‘plugged’ into the new simulator.
- (iii) The approach exploits parallel computing systems with shared and/or distributed memory to reduce the overall simulation run-time.

A. Bridging simulators via a connection layer

As shown in [3], different simulators for parallel and distributed systems support simulation of scheduling algorithms at different levels of scheduling, i.e., BLS and ALS. However,

certain simulators have a strong potential to support simulation of scheduling algorithms at certain levels of hardware parallelism. For instance, at the grid or cluster level, `GridSim` has the capability to implement and simulate BLS algorithms, while `SimGrid` has advantages to implement and simulate ALS algorithms. The two-level scheduling idea proposed in this work is, therefore, based on simultaneously executing two simulators where each simulates a certain level of scheduling, both simulators feeding each other with their scheduling decisions when needed throughout execution. Fig. 1 illustrates an example in which the BLS simulator simulates a batch of jobs and requires as input three important parameters: the *set and characteristics of the batch jobs*, the *set and characteristics of the cluster resources*, and the *BLS algorithm*. The BLS simulator decides which cluster resources should be allocated to execute a certain job from the batch at a certain time. This decision is fed into the ALS simulator, instantiated for the particular job, with the specifications of the three parameters: *tasks of particular job*, *description of allocated resources*, and the *ALS algorithm*. The arrows in green color denote the connection layer between the two levels of scheduling. The illustrative example in Fig. 1 is extended in Fig. 2 (described later in Section III-B), in which the BLS simulator is `Gridsim-Alea` and the ALS simulators are instances of `SimGrid-SD`.

B. The two-level simulator

The proposed approach meets objectives (i) and (ii) stated earlier in this section, due to the fact that users are free to select any two simulators for the two levels of scheduling and the scheduling algorithms are easily ‘plugged’ in. This approach helps users to leverage their expertise on the use of certain simulation toolkits. The proposed approach depends on simultaneously running several simulation instances as separate processes. Moreover, these instances can simultaneously

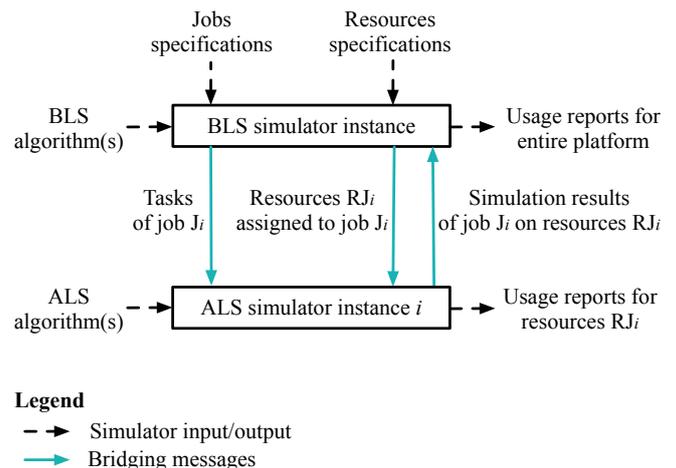


Fig. 1. Bridging between simulator instances for allocating resources R_{J_i} to job J_i using a certain BLS algorithm and executing J_i on R_{J_i} according to a given ALS algorithm.

TABLE I. Notation

Symbol	Description
N	Number of batch jobs
M	Number of cluster resources
J	Set of batch jobs $\{J_i \mid 0 \leq i < N\}$
R	Set of cluster resources $\{R_j \mid 0 \leq j < M\}$
RJ_i	Set of resources allocated to job J_i $RJ_i \subseteq R, RJ_i \neq \emptyset, 0 \leq i < N$
AT	Set of jobs arrival times $\{AT_i \mid 0 \leq i < N\}$
FT	Set of jobs finishing times $\{FT_i \mid 0 \leq i < N\}$
ST	Set of jobs starting times $\{ST_i \mid 0 \leq i < N\}$
LJ_i	Length of job J_i (in GFLOP), where $0 \leq i < N$
TJ_i	Set of all tasks belonging to job J_i , where $0 \leq i < N$
LT_k	Length of task T_k (in GFLOP) of job J_i , where $0 \leq k < TJ_k $ and $0 \leq i < N$
Υ	Task variation factor $0 \leq \Upsilon < 1$
$T_{\text{par}}^{\text{ideal}}$	Time to complete all jobs of a certain workload, where each job has a number of equal size tasks, $\max(FT) - \min(AT) \mid \Upsilon = 0, \forall J_i \in J$
$T_{\text{par}}^{\Upsilon}$	Time to complete all jobs of a certain workload, where the sizes of tasks within each job vary according to Υ $\max(FT) - \min(AT) \mid 0 < \Upsilon < 1, \forall J_i \in J$

run on shared and/or distributed memory systems. Thus, the proposed approach also meets objective (iii) stated earlier. All the results obtained within this work are based on running a new proposed two-level simulator on a multi-core processor with shared memory, described below.

Based on the approach discussed in Section III-A, a new two-level simulator is designed and implemented by connecting and integrating two different simulators. The GridSim-based simulator, Alea [12], is used to simulate BLS algorithms, while the SimGrid-SD-based simulator [3] simulates ALS algorithms. Although both simulation toolkits are well-known and have an active support community, connecting them has not yet been attempted and poses certain implementation challenges.

The first challenge is interfacing two different programming models: object-oriented and structured programming used for developing the GridSim-Alea (in Java) and the SimGrid-SD (in C) simulators, respectively. The second challenge is synchronizing the independent simulation clocks of the simulators instances. Both simulation toolkits are based on discrete events, and each keeps its own discrete simulation clock that is only advanced when an internal event occurs. The third challenge is merging the output results generated by the multiple instances of the two simulators to enable a proper informative presentation.

To address these challenges, a *connection layer* (as a part of the proposed two-level simulator) was designed to provide the following functions: (i) Manage simulator instances, (ii) Synchronize the clocks of the simulator instances, and

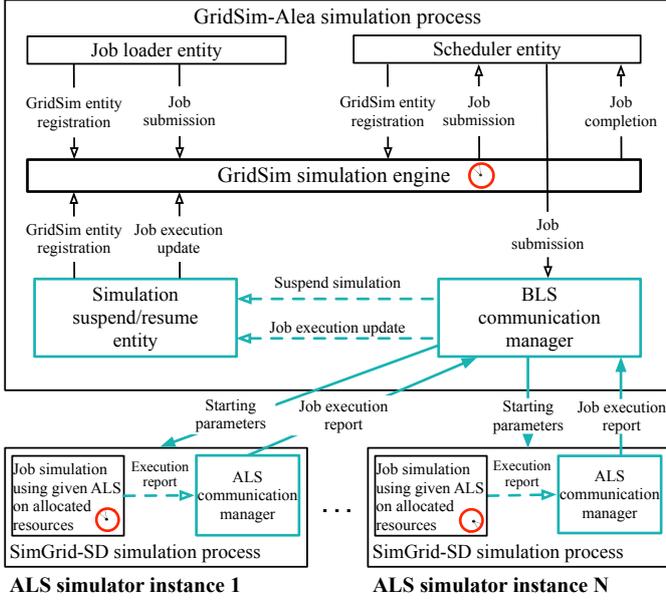
(iii) Exchange necessary information regarding jobs, tasks, and other execution parameters.

To illustrate the use of the connection layer, consider the following scheduling scenario: A batch J consists of four jobs $\{J_0, J_1, J_2, J_3\}$. Each job consists of three tasks. In each job, the sum of the lengths of the first two tasks is equal to the length of the third task, i.e., $LT_1 + LT_2 = LT_3$. A cluster R consists of five homogeneous resources $\{R_0, R_1, R_2, R_3, R_4\}$. The set of resources required by job J_i is denoted RJ_i , $0 \leq i < 4$. The following resource assignments are requested: $RJ_0 = \{R_0, R_1\}$, $RJ_1 = \{R_2, R_3\}$, $RJ_2 = \{R_2, R_4\}$, and $RJ_3 = \{R_0, R_4\}$. The arrival time of job J_i is AT_i , $0 \leq i < 4$, where $AT_0 = AT_1 = 0$ and $AT_2 < AT_3$. The finishing time of job J_i is FT_i , $0 \leq i < 4$, and $FT_0 = FT_1 > AT_3 > AT_2$. FCFS and STATIC are used as BLS and ALS, respectively.

Since $AT_0 = AT_1 = 0$, the connection layer **manages** the BLS and ALS simulator instances by *starting* two separate instances of the SimGrid-SD-based simulator to simulate the execution of jobs J_0 and J_1 on RJ_0 and RJ_1 , respectively, using STATIC. Given that $FT_1 > AT_2$ and $RJ_1 \cap RJ_2 = \{R_2\}$, the connection layer *holds* the simulation of J_2 until the SimGrid-SD-based simulation instance for J_1 reports its completion. Since $AT_3 > AT_2$, J_2 starts before J_3 , and, thus, the connection layer holds the simulation of J_3 until the SimGrid-SD-based simulation instances for J_0 and J_2 report their completion, given that $RJ_3 \cap RJ_0 = \{R_0\}$ and $RJ_3 \cap RJ_2 = \{R_4\}$. Therefore, the time at which simulation of J_3 begins depends on the times at which the simulation of J_0 and J_2 completes. The finishing times of J_0 and J_2 are dominated by the scheduling decisions of the ALS algorithms. Recall that for jobs J_0 and J_2 , the sum of the lengths of the first two tasks equals the length of the third task. Due to using STATIC as ALS and having homogeneous resources, load imbalance will arise in executing the three tasks of J_0 and J_2 on the sets of resources RJ_0 and RJ_2 that are assigned to J_0 and J_2 . As a consequence, the BLS scheduler, FCFS, needs to delay the beginning of the execution of J_3 . The influence between BLS and ALS becomes visible via the fact that STATIC as ALS not only affects the individual performance of J_0 and J_2 , but also the performance achieved by FCFS as BLS for scheduling the other jobs in the batch. In this scenario, if the FCFS algorithm passed certain information to the STATIC algorithm to prioritize the release of resources, the STATIC algorithm would assign the smallest chunk of tasks to the resources needed to be released for other jobs, such as R_0 and R_4 .

To support this type of scenarios, the connection layer **synchronizes** the running simulators using two strategies: *simulation suspend/resume* and *event injection*, as illustrated in Fig. 2. A *simulation suspend/resume entity* registered in GridSim-Alea, is used to *suspend* and *resume* the BLS simulation. It performs a busy loop that ends *if and only if* all running instances of the SimGrid-SD-based simulator report their completion and results. Due to the fact that the suspend entity is a registered GridSim entity, its busy loop can pause

BLS simulator instance



Legend

- \rightarrow Internal GridSim events
- \rightarrow External messages of the connection layer
- \rightarrow Internal synchronization events of the connection layer
- Connection layer entities
- Simulation clock within a simulation instance

Fig. 2. Two-level scheduling consisting of a single BLS and several ALS simulation instances. The connection layer synchronizes the independent simulation clocks of the GridSim-Alea-based and SimGrid-SD-based simulators.

the simulation clock of the GridSim-Alea-based simulator until the busy loop ends.

The internal synchronization events in Fig. 2 created by the BLS communication manager are used to update the simulation suspend/resume entity. Thus, the suspend/resume entity can incrementally *inject* the execution reports of the running SimGrid-SD simulation instances into the GridSim engine (see Fig. 2) and end its busy loop when there are no more running SimGrid-SD simulation instances. The simulation suspend/resume entity injects the execution reports as *GridSim events*. Therefore, the GridSim engine is able to use them to advance its simulation clock. Fig. 2 depicts the independent simulation clocks of the GridSim-Alea-based and SimGrid-SD-based simulators and their synchronization by connection layer. The connection layer uses socket-based communication and application arguments to **exchange** the information between the GridSim-Alea-based simulator and the SimGrid-SD-based simulator instances. The connection layer launches SimGrid-SD-based simulator instances as independent application processes, and passes certain parameters as application arguments to each launched process. For example, it sends the port number on which it expects to receive the ALS simulation results.

IV. EXPERIMENTS AND RESULTS

To explore the relation between BLS and ALS algorithms in HPC systems and its impact on the batch and application execution, this work considers real workloads from production HPC systems. These workloads are selected from the *parallel workload archive*² (PWA) [20]. The PWA is a public archive containing detailed information on 38 workloads of large scale parallel machines from around the world, collected between 1993 and 2015. The workloads are provided in the standard workload format (SWF). A workload may exist in two versions: cleansed and raw. The cleansed version is a minor modification of the raw version to guarantee self-consistency. The founders of PWA recommend to use the cleansed version, when available.

In this work, the cleansed version of the CEA CURIE is used and referred to as W_1 . The W_1 workload contains 312,000 jobs and spans 20 months from February 2011 to October 2012. In all experiments reported herein, a simulated platform that consists of four hosts is used. Each of the hosts has a processor that contains 64 cores. The maximum performance of one host is 3 tera floating point operations per second. A fully connected network topology is used to connect the four hosts. The network model used herein is an InfiniBand model with a link bandwidth and latency of 50 gigabits per second and 500 nanoseconds, respectively.

A. Workload customization

The workloads from the PWA only keep information relevant for batch level scheduling, such as *job ID*, *submission time*, *wait time*, *allocated resources* and *user ID*. For the purpose of this work, additional details regarding the application level are important and needed, such as the characteristics of the application each job represents, the number of parallel tasks within the application, the resource usage by each task, etc. These details are necessary for performing application level scheduling in the present work. Since this information is not present in the PWA workloads, certain assumptions are made about the applications:

- (i) All jobs in the workload are computationally-intensive. Consequently, all communication or I/O tasks that may exist in the original jobs are not considered in this work. This assumption is not a limitation of the proposed approach. It is simply used to convert the existing workloads to one of the possible cases where jobs are computationally-intensive, as such jobs are among the main incentives for using HPC systems. Other forms of jobs, such as communication-intensive, I/O-intensive and combinations thereof will be considered in future work.
- (ii) Although the number of tasks and length of each task are application-dependent, this work considers the case of *ideal parallelism*: all available hardware parallelism is exploited, execution is perfectly load balanced, communication is virtually instantaneous, and the resources allocated to tasks are identical. This work considers the case of ideal parallelism as an important baseline case.

²<http://www.cs.huji.ac.il/labs/parallel/workload/>

In addition, other cases are generated and examined by introducing variation at the task length level using the task variation factor Υ . By considering job J_i and its allocated set of resources RJ_i , the elements of the set TJ_i of tasks of job J_i can be randomly generated according to a probability distribution with mean $\mu = \frac{LJ_i}{|RJ_i|}$ and standard deviation $\sigma = \mu \times \Upsilon$.

Workloads from the PWA store the execution time and the number of resources requested by each job. The PWA also states certain detailed about the hardware platform where the workloads were obtained from. Based on his information, an estimate about each job length (in giga floating point operations) can be obtained by multiplying the minimum performance of request resources (in giga floating point operations per second) to the job execution time (in seconds). The job length LJ_i is deduced by accumulating the length of the generated tasks, until LJ_i become greater than or equal to the estimate.

Many researchers modeled the arrival, finishing, execution times, and number and type of requested resources of different jobs in the context of HPC. Few efforts exist that can be used to model (with certain adaptations) the number of tasks and the task length within HPC applications. In the present work, for simplicity, the task lengths are generated according to normal distribution with the aforementioned μ and σ parameters. Future work will consider other models to generate the task lengths within jobs of a given workload.

B. Analysis of the relation between BLS and ALS algorithms

To explore the relation between BLS and ALS algorithms, experiments were performed following two different strategies. The first strategy supports a *coarse-grain analysis* of the relation between BLS and ALS, examining the effect of changing the ALS algorithms on the BLS performance, measured as the makespan of the entire batch workload. The second strategy supports a *fine-grain analysis* of the relation between BLS and ALS that examines the effect of changing the ALS algorithm in one job on the starting time of its successor job(s) in the batch.

In this work, combinations of three BLS algorithms FCFS, EDF, and SJF [7] and four ALS algorithms STATIC, SS [8], GSS [9], and FAC [10] are considered. To perform the experiments for the coarse-grain analysis, jobs of the most intensive 24 hours in terms of job arrival time have been selected from W_1 . The most intensive 24-hour interval of W_1 is referred to as W_1^{24} . A Java-based tool was developed as part of this work and used to extract the jobs of the most intensive period of 24 hours from a given workload.

As discussed in Section IV-A, the task length variation factor Υ is used to vary the length of tasks within a certain job to represent more realistic jobs. Fig. 3 shows the parallel execution time of W_1^{24} when executing the twelve BLS-ALS combinations in two different cases. In the first case, $\Upsilon = 0$ and corresponds to T_{par}^{ideal} . In the second case, $\Upsilon \in \{0.1, 0.15, 0.2, 0.25\}$ and leads to different values of T_{par}^Υ denoted as $T_{par}^{0.1}$, $T_{par}^{0.15}$, $T_{par}^{0.2}$, and $T_{par}^{0.25}$, respectively. In the case of $\Upsilon = 0$, the results represent the best case

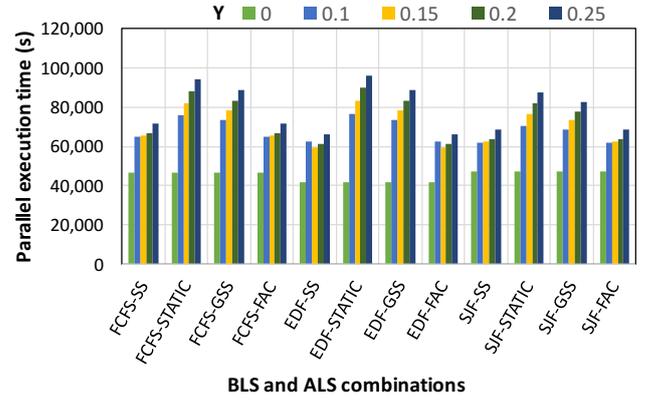


Fig. 3. Effect of changing the task variation factor Υ from 0 to 0.25 on the parallel execution time for the twelve combinations of selected BLS and ALS algorithms for the jobs within W_1^{24}

scenario in which all submitted jobs are perfectly optimized for their allocated resources. Although such a scenario is highly desirable at the cluster operation and at the user levels, it is difficult to be encountered in practice. The results of Fig. 3 show the effect of increasing Υ from 0 to 0.25 for the twelve BLS-ALS combinations. One can infer that increasing Υ leads to an increase in the total parallel execution time regardless of the BLS-ALS combination that has been used. The increase in the parallel execution time of the workload is not constant for different BLS-ALS combinations. Certain BLS-ALS combinations showed the ability to better absorb the effect of increasing Υ than others. Further analysis is needed to understand this behavior and pinpoint its root-cause(s).

The results in Fig. 4 show the ratio between the parallel execution time T_{par}^Υ and the ideal parallel execution time T_{par}^{ideal} , respectively. The ratio $\frac{T_{par}^\Upsilon}{T_{par}^{ideal}}$ can be used to characterize the sensitivity of the system performance to given BLS-ALS combinations.

To perform a fine-grain analysis, the connection layer between GridSim-Alea-based and SimGrid-SD-based simulators was extended with an additional task: to collect all text-based traces generated from the SimGrid-SD-based simulator and to combine them into a single text-based trace file. The main

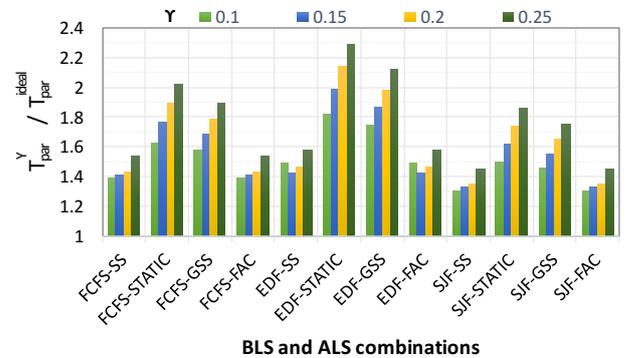


Fig. 4. The ratio between T_{par}^Υ and T_{par}^{ideal} for the twelve combinations of selected BLS and ALS algorithms for the jobs within W_1^{24}

challenge associated with this task is that each instance of the SimGrid-SD-based simulator does not have the global view of the entire batch workload simulation. For instance, to simulate jobs J_1 and J_2 on the sets of resources JR_1 and JR_2 at times t_1 and t_2 , respectively, the connection layer runs two instances of the SimGrid-SD-based simulator. Each SimGrid instance, however, simulates its corresponding job as J_i on the set of resource JR_i at time t_x .

In this work, a tool was used to convert the collected text-based traces to binary traces in the OTF2 [6] format. Using OTF2 traces with the Vampir [21] trace visualizer, we are able to visualize for the first time, to the best of our knowledge, the cluster utilization from the node to the core level *and* from batch level to application level scheduling, as shown in Fig. 5. A snapshot captured from Vampir is included in Fig. 5 and shows the execution of three out of 1,700 running jobs, namely J_{1542} , J_{1543} , and J_{1544} , from the W_1^{24} workload. The execution was performed with a combination between FCFS as BLS and GSS as ALS. The tasks of the three jobs utilize host 1. The other three hosts are also utilized by the three jobs. Due to the limited space, the execution of the three jobs on hosts 0, 1, and 2 are collapsed and not shown. Fig. 5 illustrates a case of severe load imbalance of certain jobs, its effects on the starting times of subsequent jobs in the batch, and, consequently, the effects on the entire system performance and utilization.

Scalability is an interesting aspect of the proposed two-level simulation approach, in terms of increasing number of jobs and, consequently, in terms of increasing the number of

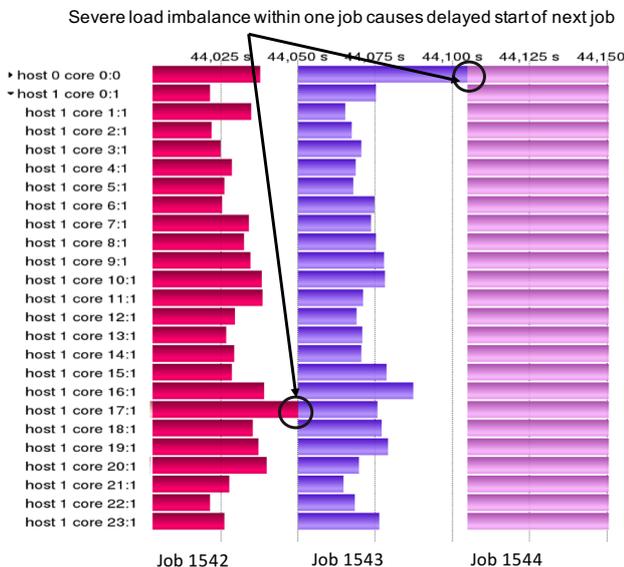


Fig. 5. Snapshot of the Vampir visualization tool showing the generated OTF2 trace of the proposed two-level scheduling simulator. Tasks of the same job are represented using horizontal bars of the same color, while the white space between the job bars represents the idle state of the allocated cores. The scheduling algorithms shown herein are FCFS and GSS at BLS and ALS, respectively. Jobs are obtained from workload W_1^{24} . This snapshot spans the execution time between 44,017 and 44,150 out of 88,441 seconds total execution time.

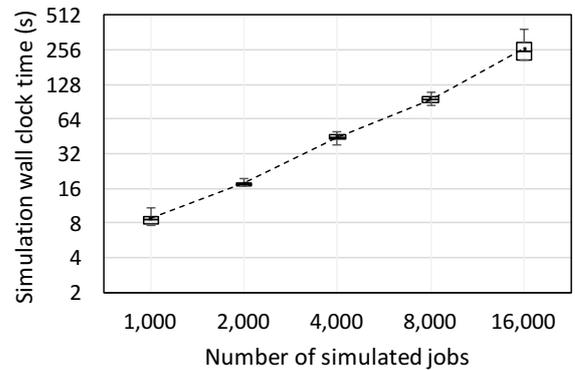


Fig. 6. The simulation wall clock time of the two-level simulator on an increasing number of simulated jobs from W_1 .

simultaneous ALS instances. An initial scalability assessment of the two-level simulator is presented in Fig. 6. In these experiments, the simulation wall clock times are reported for executing an increasing number of jobs (from 1,000 to 16,000) from the W_1 workload.

The simulation wall clock is defined as the total time required by the two-level simulator to simulate the execution of all jobs of a given workload and their tasks. The experiments were conducted with the least performing BLS-ALS combination, i.e., FCFS-STATIC, with $\Upsilon = 0.25$, selected from the results in Fig. 3. The results reveal a linear relation between the increase in the number of simulated jobs and simulation wall clock time consumed by the two-level simulator.

V. REPRODUCIBILITY OF THIS WORK

To ensure reproducibility of this work, apart from the information in Section IV about the workloads and the simulated platform considered in this work, the code of the proposed two-level simulator is developed under LGPL license, and is publicly available³. Under the same LGPL license, the developed Java-based tool used for extracting the jobs corresponding to the most intensive time period is also available upon request. Table II summarizes the software and the hardware specifications of the platform on which the proposed two-level simulator has been compiled and executed.

VI. CONCLUSION AND FUTURE WORK

With the growing complexity of modern and future HPC systems, parallelism becomes more massive and available

TABLE II. Characteristics of the platform used for conducting experiments

Software	
Operating system	OS X 10.11.5
Required libraries for build	GridSim v.5Alea v.4 SimGrid v.3.13
Compilers	for SimGrid, clang v.7.3.0 for GridSim, javac v.1.8.0.91
Hardware	
Processor model	Intel Core i7
Processor frequency	2.5 GHz
RAM (DDR4)	16 GB

³<https://c4science.ch/diffusion/3130/>

at additional hardware levels. As a consequence, efficient exploitation and scheduling at these levels of parallelism is required. It is, therefore, important and necessary to explore the relation between different levels of scheduling to enhance the performance and utilization of modern HPC systems as a whole and not only at individual scheduling levels. This work can be considered as an important first step in this direction. The proposed two-level simulation approach, connects simulators from two scheduling levels (BLS and ALS) and showed its validity to explore the relation between BLS and ALS. Based on the proposed approach, a novel two-level simulator was proposed and successfully used to simulate combinations of three BLS and four ALS well-known algorithms from the literature. The choice of ALS not only affects the performance of the applications for which it was employed, but also the performance of the chosen BLS for the other jobs in the batch.

The proposed two-level simulator also generates execution traces in binary format. This enables the visual analysis of the job execution at the batch level and task execution within each job at the application level. Such an analysis can help identify severe load imbalance and execution “hotspots”. This type of insight is an important advantage for the development of future multi-level scheduling algorithms. The first evaluations of the proposed two-level simulator indicate its usefulness and scalable performance with the number of simulated jobs. It was able to simulate 16,000 jobs and their tasks at BLS and ALS levels in less than 261 seconds.

Further work is needed and planned to deepen the understanding of the relation between different levels of scheduling in modern large-scale HPC systems. It is important to understand the root-cause behind certain BLS-ALS combinations being able to better absorb the effect of variable job and task lengths than others.

ACKNOWLEDGMENT

This work is in part supported by the Swiss National Science Foundation in the context of the Multi-level Scheduling in Large Scale High Performance Computers (MLS) grant, number 169123. The authors acknowledge the contribution of Mr. Aram Yesildeniz to the code that converts text-based SimGrid output to OTF2 traces.

REFERENCES

- [1] D. H. Ahn, J. Garlick, M. Grondona, D. Lipari, B. Springmeyer, and M. Schulz, “Flux: A Next-Generation Resource Management Framework for Large HPC Centers,” in *Proceedings of the 43rd International Conference on Parallel Processing Workshops (ICPPW)*, pp. 9–17, September 2014, Minneapolis, USA.
- [2] J. Ullman, “NP-complete scheduling problems,” *Journal of Computer and System Sciences*, vol. 10, no. 3, pp. 384–393, 1975.
- [3] A. Eleliemy, A. Mohammed, and F. M. Ciorba, “Simulating Batch and Application Level Scheduling Using GridSim and SimGrid,” Poster at the 28th International Conference for High Performance Computing, Networking, Storage, and Analysis (SC16), November 2016, Salt Lake City, USA.
- [4] R. Buyya and M. Murshed, “GridSim: A toolkit for the modeling and simulation of distributed resource management and scheduling for grid computing,” *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1175–1220, 2002.
- [5] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, “Versatile, scalable, and accurate simulation of distributed applications and platforms,” *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2899–2917, 2014.
- [6] D. Eschweiler, M. Wagner, M. Geimer, A. Knüpfer, W. E. Nagel, and F. Wolf, “Open Trace Format 2: The Next Generation of Scalable Trace Formats and Support Libraries,” in *Proceedings of the International Conference on Parallel Computing (ParCo)*, pp. 481–490, April 2011, Ghent, Belgium.
- [7] Z. R. M. Azmi, K. A. Bakar, M. S. Shamsir, W. N. W. Manan, and A. H. Abdullah, “Scheduling Grid Jobs Using Priority Rule Algorithms and Gap Filling Techniques,” *International Journal of Advanced Science and Technology*, vol. 37, pp. 61–76, 2011.
- [8] P. Tang and P.-C. Yew, “Processor Self-Scheduling for Multiple-Nested Parallel Loops,” in *Proceedings of the International Conference of Parallel Processing (ICPP)*, pp. 528–535, January 1986, Urbana, USA.
- [9] C. D. Polychronopoulos and D. J. Kuck, “Guided self-scheduling: A practical scheduling scheme for parallel supercomputers,” *IEEE Transactions on Computers*, vol. 100, no. 12, pp. 1425–1439, 1987.
- [10] S. F. Hummel, E. Schonberg, and L. E. Flynn, “Factoring: A method for scheduling parallel loops,” *Communications of the ACM*, vol. 35, no. 8, pp. 90–101, 1992.
- [11] F. Howell and R. McNab, “A Discrete Event Simulation Library for Java,” in *Proceedings of the 1st International Conference on Web-based Modeling and Simulation*, p. 6, January 1998, California, USA.
- [12] D. Klusáček and H. Rudová, “Alea 2: Job scheduling simulator,” in *Proceedings of the 3rd International Conference on Simulation Tools and Techniques (ICST)*, p. 10, March 2010, Malaga, Spain.
- [13] P.-F. Dutot, M. Mercier, M. Poquet, and O. Richard, “Batsim: A Realistic Language-Independent Resources and Jobs Management Systems Simulator,” in *Proceedings of the 20th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP)*, p. 20, May 2016, Chicago, USA.
- [14] D. Klusáček, L. Matyska, and H. Rudová, “Alea-Grid scheduling simulation environment,” in *Proceedings of the 7th International Conference on Parallel Processing and Applied Mathematics (PPAM)*, p. 10, September 2007, Gdansk, Poland.
- [15] S. J. Chapin, W. Cirne, D. G. Feitelson, J. P. Jones, S. T. Leutenegger, U. Schwiegelshohn, W. Smith, and D. Talby, “Benchmarks and standards for the evaluation of parallel job schedulers,” in *Proceedings of the 5th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP)*, pp. 67–90, April 1999, San Juan, Puerto Rico.
- [16] S. Srivastava, I. Banicescu, F. M. Ciorba, and W. E. Nagel, “Enhancing the Functionality of a GridSim-based Scheduler for Effective Use with Large-Scale Scientific Applications,” in *Proceedings of the 10th International Symposium on Parallel and Distributed Computing*, pp. 86–93, July 2011, Cluj-Napoca, Romania.
- [17] Y. Caniou and J.-S. Gay, “Simbatch: An API for simulating and predicting the performance of parallel resources managed by batch systems,” in *Proceedings of the European Conference on Parallel Processing (EuroPar)*, pp. 223–234, August 2008, Canary Island, Spain.
- [18] F. M. Ciorba, T. Hansen, S. Srivastava, I. Banicescu, A. A. Maciejewski, and H. J. Siegel, “A combined dual-stage framework for robust scheduling of scientific applications in heterogeneous environments with uncertain availability,” in *Proceedings of the 21st International Heterogeneity in Computing Workshop (HCW) of the 26th Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW)*, pp. 193–207, May 2012, Shanghai, China.
- [19] T. Hansen, F. M. Ciorba, A. A. Maciejewski, H. J. Siegel, S. Srivastava, and I. Banicescu, “Heuristics for Robust Allocation of Resources to Parallel Applications with Uncertain Execution Times in Heterogeneous Systems with Uncertain Availability,” in *Proceedings of the 19th International Conference of Parallel and Distributed Computing (ICPDC) of the 25th World Congress on Engineering (WCE)*, p. 6, July 2014, London, UK.
- [20] D. G. Feitelson, D. Tsafir, and D. Krakov, “Experience with using the Parallel Workloads Archive,” *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2967–2982, 2014.
- [21] A. Knüpfer, H. Brunst, J. Doleschal, M. Jurenz, M. Lieber, H. Mickler, M. S. Müller, and W. E. Nagel, “The Vampir performance analysis toolset,” in *Proceedings of the 2nd International Workshop on Parallel Tools for High Performance Computing*, pp. 139–155, July 2008, Stuttgart, Germany.