

Scheduling Distributed Applications: the SimGrid Simulation Framework

Arnaud Legrand

Loris Marchal

Henri Casanova

Laboratoire de l'Informatique du Parallélisme
École Normale Supérieure de Lyon

Dept. of Computer Science and Engineering
San Diego Supercomputer Center
University of California at San Diego

Abstract— Since the advent of distributed computer systems an active field of research has been the investigation of *scheduling* strategies for parallel applications. The common approach is to employ scheduling heuristics that approximate an optimal schedule. Unfortunately, it is not possible to obtain analytical results to compare and select appropriate heuristics for a given scheduling problem. One possibility is to conduct large numbers of back-to-back experiments on real platforms. While this is possible on tightly-coupled platforms, it is infeasible on modern distributed platforms such as the *Computational Grid* as it is labor-intensive and does not enable repeatable results. The solution is to resort to *simulations*. Simulations not only lead to repeatable results but also make it possible to explore wide ranges of platform and application scenarios.

In this paper we present the SimGrid framework which enables the simulation of distributed applications in distributed computing environments for the specific purpose of developing and evaluating scheduling algorithms. This paper focuses on SimGrid v2, which greatly improves on the first version of the software. Enhancements include improved and more realistic network models, as well as means to import platform simulation models from platform monitoring tools. SimGrid v2 also enables the simulation of distributed scheduling agents, which has become critical for current scheduling research in the area of Grid computing. After describing and validating these features, we present a case study by which we demonstrate the usefulness of SimGrid for conducting scheduling research.

I. INTRODUCTION

Since the advent of distributed computer systems, the question of *scheduling* parallel applications has generated a very large body of work. The process of scheduling consists in assigning the tasks of an application to available resources, both in time and space, with as objective the minimization of a metric (e.g. execution time, throughput, resource utilization, monetary cost). Most scheduling problems are NP-hard [46]. Even when an optimal solution to a scheduling problem can be found in polynomial time, it is often the case that small modifications of the underlying assumptions (e.g. addition of non-zero network latencies) ren-

der the problem NP-hard. As a result, a common approach is to use *heuristics* to approximate an optimal schedule. Classes of heuristics have been developed by several researchers for different scheduling problems (e.g. scheduling task graphs [27], scheduling independent tasks [6]). The advantage of these heuristics is that they have low complexity and can therefore be used in practice. However, it is generally not possible to quantify the efficacy of heuristics analytically. Consequently, scheduling heuristics must be compared by performing many experiments in many scenarios.

One possible approach is to perform experiments with real applications on real resources. However, modern computing platforms are increasingly distributed and often span multiple administrative domains (e.g. Computational Grids [22, 23]). Therefore, resource availability fluctuations makes it impossible to conduct repeatable experiments. While it is possible to run back-to-back runs of short applications in certain cases, experiments for relatively long running applications are problematic. Another problem with using real platforms is that the number of platform configurations that can be explored is limited. As a result of these difficulties with real experimentations, most researchers have resorted to discrete-event simulation (which we will refer to as “simulation”).

Simulation has been used extensively as a way to evaluate and compare scheduling strategies as simulation experiments are configurable, repeatable, and generally fast. In spite of these advantages, we claim that there are two main limitations to the simulation methodology used for scheduling research. First, there is no simulation standard in the scheduling research community. Typically, researchers build “throw-away” simulators using a variety of programming languages and tools. This makes it very difficult for other researchers to reproduce results presented in papers, and therefore to compare results among papers. This pre-

vents scientific advances and is in contrast to other research communities in which simulation standards have been adopted (e.g. networking, computer architecture). This lack of standard simulation procedure and software was somewhat justifiable when the simulation models in use were simplistic. However, traditional models and assumptions about computing platforms are no longer valid for modern platforms, which leads us to the second limitation of simulations used in the scheduling literature. The simplistic network models used in most scheduling literature do not hold for modern computing platforms in which compute resources are connected over complex network topologies with complex link contention behaviors. The assumption that the behavior of the compute platform is perfectly predictable also needs to be revisited as modern platforms exhibit dynamic resource availabilities. Realistic simulation must make it possible to simulate such fluctuations as they cause uncertainty about the platform knowledge that a scheduler can obtain.

Consequently, there is a need for a simulation framework designed for conducting research on distributed application scheduling. To be useful this framework must meet the following objectives. It must provide *high usability* so that users can focus on scheduling research and not on simulation issues. It must make it possible to run *fast simulations* because large numbers of simulation experiments must be conducted to evaluate and compare scheduling algorithms. This framework must make it possible to build *configurable*, *tunable*, and *extensible* simulations. In particular, it should be possible to go from “traditional” simulation models to more realistic ones without having to rethink the implementation of the simulation. Finally, given the scale of modern applications and platforms, the simulation framework must be *scalable* and sustain simulations with tens of thousands of resources and application tasks. We claim that no available simulation framework currently matches these requirements.

To address these issues, we have designed the SIMGRID simulation framework. This paper focuses on SIMGRID v2, which greatly improves on the first version of the software [9]. Enhancements include improved and more realistic network models, means to import platform simulation models from platform monitoring tools, and improved APIs. Most importantly, a new software layer makes it easier to build complex simulation by providing higher-level abstractions. In particular, it provides capabilities for simulating distributed scheduling agents. This is becoming critical in scheduling research as platforms and applications

increase in scale and a single scheduler is a single point of failure and reduces scalability. SIMGRID has already been used successfully for a number of research works [10, 41, 18, 43, 3, 4, 47] and promises to be of great value to the scheduling research community.

This paper is organized as follows. In Section II we review related work on simulation tools. Section III describes version 2 of the SIMGRID software. Section IV identifies issues and solutions for platform modeling with SIMGRID and Section V presents a case-study. Section VI concludes the paper and points at future work.

II. BACKGROUND AND RELATED WORK

A large number of software tools have been developed for users to build and run simulations in many application domains. In this section we review those tools that could be employed to conduct research in the area of distributed application scheduling.

Several software libraries and environments provide ways to construct generic discrete-event simulations [38, 45, 34, 20, 24, 25]. Such generic-purpose tools could be used for simulating parallel applications on distributed platforms. An impediment is that the Application Programming Interfaces (APIs) provided by these tools are not tailored to the investigation of scheduling algorithms. The APIs are low-level, because generic, and a user who wishes to do research on scheduling algorithms is faced with the task of implementing many higher-level abstractions. Each researcher would probably build different such abstractions, thereby not solving the problem of simulation standardization identified in Section I. Besides, researchers in the area of scheduling seldom use these low-level simulation packages in current practice.

An area in which simulation is used extensively is networking research. In order to understand how large telecommunication networks scale and behave in various situations, several network simulators have been developed. Example of such simulators include NS [35, 2], DaSSF [29] or OMNeT++ [36]. These simulators focus on precise simulation of packets traveling on the network, rather than on the network behavior as it is experienced by an application. The main intent is to help identifying limitations of network protocols and help developing improvements. Although it is possible to perform application-level simulation with certain of these tools (e.g. with DaSSF [29] or OMNeT++ [36]), it is labor-intensive and a number of capabilities need to be implemented from scratch. Furthermore, due to their highly detailed simulation

models, most network simulators induce long simulation times (e.g. they implement the TCP stack). This issue is exacerbated when one wishes to simulate “background” network traffic generated by other users/applications as one may then need to simulate thousands of connections at the packet-level. This is prohibitive for conducting the large numbers of simulation experiments mandated by scheduling research.

Another limitation of network simulators for our purpose is that they do not simulate the entire platform. Even though the network is a fundamental component of a distributed computing platform, other components are required for investigating scheduling algorithms: compute resources and application processes. A number of projects address the simulation of a complete distributed computing platforms for studying parallel applications. We review here three representative projects: LAPSE, MicroGrid, and Albatross.

LAPSE (Large Application Parallel Simulation Environment) [15] is a tool designed to simulate parallel applications. The main goal is to make it possible to test the scalability of parallel application while running them only on a few processors. A focus of the project is precise network event simulation. The application code is instrumented so as to account for the number of instructions executed between two consecutive communication. LAPSE was originally targeted to the Intel Paragon architecture, but allows the simulation of other platforms. One important limitation is that the network model cannot capture contention and cannot easily model background network traffic.

Microgrid [42], a more recent project, targets the simulation of Grid environments. It follows somewhat the same philosophy as LAPSE in that it emulates large, complex, virtual platforms on small, physical platforms. The software virtualizes every resource of a Grid platform (memory, CPU, and network). Just as in LAPSE, this virtualization is achieved by trapping every relevant library call. In Microgrid’s first version, the network was simulated through a modified version of the NS [2] simulator but the latest version uses DaSSF [29]. At the moment, Microgrid does not allow the simulation of variable, background network traffic.

The Albatross [26] project focuses on platforms that consist of multiple clusters or MPPs that are connected by wide-area networks (WANs). In the context of this project, work based on the Panda library [12] addresses the simulation of these platforms. Unlike LAPSE and Microgrid, the network is not simulated but emulated by slowing down LAN links to model WAN links. This

makes it straightforward to dynamically modify network behavior during a simulation. A limitation is that heterogeneity is modeled only for WAN links.

In these three projects, portions of the application are effectively executed on an emulated architecture. The rationale is that emulation is realistic because it can capture more detailed and idiosyncratic behaviors. However, a consequence is that the ratio of simulated time to simulation time can be very large. This is prohibitive for comparing scheduling algorithms in tens of thousands of scenarios.

We conclude that no existing simulation framework satisfies the requirements that we have identified in Section I. In the rest of this paper we describe our work on SIMGRID which addresses these requirements. Note that the GridSim [7] project is very related to our work in scope and intent. GridSim is based on SimJava [25] and implements a number of abstractions to enable simulation of Grid environments. At the time this article is being written, we were unable to successfully install and test the GridSim software distribution, hence preventing us from comparing it with our work.

III. SIMGRID v2.0

A. History

The first version of SIMGRID was a discrete-event simulation toolkit. It provided a set of core abstractions and functionalities that can be used to easily build simulators for specific application domains and/or computing environment topologies. This allows the simulation of arbitrary performance fluctuations such as the ones observable for real resources due to background load. However, this first version lacked a number of abstractions (e.g. routing, scheduling agents). With SIMGRID v2 we have added a new software layer to provide high-level abstractions and the software thus provides two interfaces:

SG : The original low-level toolkit presented in [9], by which the simulation is done in terms of explicitly scheduling tasks on some resources;

MSG : A simulator built using SG. This layer implements realistic simulations based on the foundational SG and is more application-oriented. Simulations are built in terms of communicating agents.

In all that follows, we focus on the MSG interface as we anticipate it will become the standard interface for the vast majority of SIMGRID users. Also, “SIMGRID” denotes version 2 of the software throughout.

B. Fundamental Concepts

SIMGRID implements the following core abstractions:

Agent – An agent is an entity which makes scheduling decisions. An agent is defined by a code, private data, and the location at which it executes.

Location – A location (or host) is the place in the simulated topology at which an agent runs. Thus it is defined by a computational resource, a number of mailboxes that enable communication with other agents, and private data that can be only accessed by agents at the same location.

Task – A task is an activity of the simulated application and for now can be a computation and/or a data transfer. A task is defined by an amount of computing, a data size, and private data.

Path – The low-level, original SG layer in the software did not provide any abstraction for message routing among locations. This made the task of the user arduous when simulating complex platforms. SIMGRID now provides a routing abstraction so that the user (and the scheduler) can rely of a logical topology of the platform rather than targeting the physical topology directly. A path is an agglomeration of communication resources representing a set of physical network links. Locations are then interconnected through paths. The simulated application cannot access links directly (in the same way as a real application does not choose which routers its packets go through).

Channel – Communication between agents is embedded in the channel abstraction. A channel embodies the concept of communication ports opened by agents at locations.

With these abstractions, scheduling algorithms with SIMGRID should always be described in terms of agents that run at locations and interact by sending, receiving, and processing simulated application tasks. Agent does not have direct access to paths but can send a task to another location using a channel. In fact, a location may have many mailboxes and a channel is then simply a mailbox number. So sending a task to a location using a channel amounts to transferring the task on a particular path, depending on the emitter location and on the destination, and to put it in a particular mailbox. We claim that these abstractions are sufficient to conduct scheduling research while isolating the user from details of the simulation that are not directly relevant to scheduling algorithms.

C. Building a Simulation with SIMGRID

A SIMGRID program always follows the following steps:

1. Definition of the code of each agent (i.e. modeling the application);

2. Creation of resources (i.e. modeling of the physical platform). It consist in defining hosts, links, and a routing table that specifies paths;
3. Creation and allocation of agents to locations (i.e. the deployment of the application);
4. Simulation.

The SIMGRID software distribution provides documentation as well as a number of tutorial examples that illustrate these four steps.

IV. PLATFORM MODELING WITH SIMGRID

SIMGRID provides several mechanisms for constructing simulated computing platforms. In this section we review the basic resource models, describe how realistic platform topologies can be generated, and describe how SIMGRID simulates bandwidth-sharing.

A. Basic Resource Models

Low-level resource objects in SIMGRID are *hosts* and *links*. A host is described by its computational speed relative to that of a reference host, and by its CPU availability (a value between 0 and 100%). A link is described by a latency and a bandwidth. By default, SIMGRID uses the traditional models of task execution time equal to the computational cost divided by the computational speed, and of data transfer time equal to the latency plus the data size divided by the bandwidth. Note that the units of these values are not specified. Instead the user must ensure coherence between units of data size and bandwidth, and of task computational cost and computational speed. All values characterizing SIMGRID resources can be either fixed, or varying according to a trace (a series of time-stamped values).

An important question is that of resource sharing among tasks of the simulation (including computational tasks and data transfer tasks). SIMGRID implements three different sharing modes: (i) First In First Out (FIFO); (ii) First Ready First Out (FRFO); (iii) Shared. In the FIFO mode tasks on a resource execute in the order they were scheduled on that resource. In the FRFO mode, tasks that are ready for execution execute first. If there is a tie, then the task that was scheduled first executes first. In the Shared mode all ready tasks execute concurrently on a resource and SIMGRID allows the user to implement fair sharing or proportional sharing with task priorities. These different modes provide some flexibility and make it possible to implement simulations with traditional assumptions from the scheduling literature.

The strength of SIMGRID is that the user can easily

use the API to tune the behavior of the simulated environment. For instance, moving from resources with constant performance characteristics to resources with characteristics that vary according to traces can be achieved by just modifying a few parameters of the resource creation API functions.

The current trend in high-performance computing is to move to Grid platforms in which resources are interconnected over complex wide-area networks. For applications to benefit from such platforms it is necessary that novel scheduling approaches be designed that can accommodate complex networks. As a result, it is necessary to provide simulation tools that can capture the behavior of these networks and which can in turn enable new scheduling research. The default resource models in SIMGRID described above do not fully enable the transition from traditional parallel computing scheduling research to Grid computing scheduling research.

We identify two challenges that must be addressed to enable this transition. First, it must be possible for users to construct large simulated platforms that are representative of existing platforms. Second, it must be possible to simulate the complex network contention behaviors of applications executing on these platforms. In the next two sections we describe how SIMGRID addresses these two challenges and makes it easy for users to build realistic simulation models for studying scheduling on Grid platforms.

B. Modeling Grid Topologies

Most simulations reported in the scheduling literature have used either simple topology models (e.g. a ring or a star), representations of a real topology (e.g. the Arpanet), or randomly-generated flat topologies using a variation of Waxman’s edge-probability function. Recently, complex randomly-generated hierarchical models have been used to better approximate the Internet’s hierarchical structure. Many studies on Internet topologies are available [16, 8, 19, 33]. Using snapshots of the Internet, simple power-laws of the Internet topology have been derived, leading to several public-domain generators of “representative” Internet topologies [1, 32, 44].

While these generators constitute an excellent starting point for building a realistic simulation model, they lack several critical elements for our purpose. Typically, there is no information concerning the traffic (e.g. the available bandwidth throughout time) and no information on the characteristics of network links. Common techniques involve the use of simulated sources of

traffic using random laws which are seldom validated against real traffic. Finally, several parameters have to be hand-tuned to obtain realistic platforms. In summary, the “art” of building a realistic platform model requires both significant experience and good intuition.

The alternative is, rather than simulating traffic sources, to collect traffic traces from a real environment. These traces can then be replayed for conducting repeatable simulation experiments, which is needed for the investigation of scheduling algorithms. To enable this, SIMGRID makes it possible for users to automatically *import* platform descriptions obtained with Effective Network View (ENV) [40]. ENV is a tool that discovers a variety of characteristics of a Grid testbed including the network and its effective topology. The main advantage of ENV is that it only uses user-level observations of network performance to create an effective profile of network configuration and thus, it makes it possible to get information about layer 2 and 3 routers directly. The obtained topology can be used to deploy NWS sensors and record real-time traces which are directly usable in SIMGRID. The combination of ENV and NWS makes it possible to instantiate platform models which represent realistic platforms both in terms of topology and in terms of traffic, as demonstrated in [28]. These models can then be used as is, or as basis for generating realistic randomized collections of models for conducting large numbers of simulations.

C. Bandwidth Sharing Models in SIMGRID

Modeling realistic topologies is a critical step forward to enable scheduling research for emerging computing platforms such as the Grid. An important question is then that of simulating contention among data transfers on these topologies. The default resource models in SIMGRID as described in Section IV-A do not enable realistic simulation of network traffic on topologies in which paths between hosts use multiple network links. Indeed, the basic models implement a simple store-and-forward scheme, without any notion of packets or network pipelining. Furthermore, if the Shared resource sharing mode is used for network links, it is not clear how to assign priorities among competing transfers (it is well-known that bandwidth-sharing among TCP flows is not fair [13, 21, 31]).

A possibility is to implement a notion of packets in the simulation in an attempt to be closer to the actual implementation of network protocols. This functionality is available in SIMGRID and allows pipelining of communications on multiple links and among multiple data transfers. The user can trade-off speed versus ac-

curacy by tuning the packet size (large packet sizes are cheaper to simulate). While this approach has some benefits, it also has a number of limitations. First, the actual network protocols are not modeled. Second, unless the packet size is large, the memory footprint and computation time of the simulation can be very large for large platforms and applications. This is problematic for conducting scheduling research as large numbers of simulation experiments are necessary to obtain statistically significant results.

An alternate approach is to investigate macroscopic models of bandwidth-sharing and to implement these models directly in the simulation. If such models can be analytically derived then they can be implemented extremely efficiently and lead to short simulation times even for very large-scale platforms. A number of recent efforts in network research have tackled the problem of finding models for bandwidth-sharing [30, 5]. The typical approach is to make an analogy with fluids in pipes and to ignore the granularity of packets in data transfers, or *flows*. An important question is: "What is the bandwidth-sharing behavior of TCP?" and a number of authors propose theoretical models [13, 21, 31]. Building on these results, SIMGRID implements an efficient algorithm for simulating TCP flows competing over multi-path routes. We have validated this model with the Network Simulator (NS) [35]. Bandwidth-sharing is implemented "under the cover" and is completely transparent to the SIMGRID user. Due to space constraints we refer the reader to [11] for details on the model and its validation.

SIMGRID makes it possible to define two types of links for the purpose of bandwidth-sharing: links on which the bandwidth is shared among flows of the simulated applications, and links on which it is not. On these latter links it is assumed that all competing flows of the simulation application achieve the same, constant data-transfer rate. The rationale is that, over a congested backbone, the contention due to interaction between flows of the simulated application is negligible compared to that of the many (thousands) flows going through the backbone. SIMGRID computes bandwidth-sharing for topologies that contain both types of links. This makes it possible to build Grid models as sites interconnected by backbone links while each site has its own internal topology (see [11] for a discussion). We conjecture that such a model is simple enough to be instantiated easily while it captures the relevant network behaviors necessary for conducting meaningful Grid computing scheduling research. We are validating this conjecture in current work and will report on

results in upcoming papers.

A key strength of SIMGRID is that its simple API makes it possible to use a wide variety of platform models interchangeably without impacting the implementation of the scheduling algorithms under investigation. It is possible for every user to pick and choose which model is appropriate for his/her research and to move to more complex models when warranted. The conjectural Grid model that we briefly described above is just one of the many possible models that can be implemented with SIMGRID. The question of which Grid model is "right" is not answered here, but SIMGRID provides an ideal research platform for making progress towards an answer.

V. CASE STUDY

To better illustrate the capabilities and ease-of-use of SIMGRID we present a small case study. The target application is a high-throughput master-worker application such as the ones being deployed in many global computing systems [39, 14, 37, 17]. A *master* dispatches input data for computational tasks to p *workers* one at a time. The question we wish to answer is: "Given n independent, identical tasks that must be distributed to workers, how should the master fulfill requests for work from the workers?"

This is a classical scheduling problem and a straightforward strategy is to employ a greedy algorithm by which the master allocates tasks on a first-come first-serve basis. We denote this strategy as *Work Queue*. Another approach is to maintain a "pool" of pending requests, $p/2$ such requests in our example, and to dispatch work according to priorities. Several criteria for prioritization are possible and in this case study we implemented the following three:

- **Computing power:** tasks are sent with higher priority to workers with faster computing speed;
- **Bandwidth:** tasks are sent with higher priority to workers with faster bandwidth;
- **Throughput:** if c_i (resp. w_i) denotes the time needed by the i -th worker to process (resp. receive data for) a task, then the ratio $w_i/(c_i + w_i)$ represent the throughput for that worker. Tasks are sent with higher priority to workers with higher throughput.

Using SIMGRID, it took under two hours to set up the entire simulation framework from scratch with the Work Queue scheduling algorithm and for a platform imported from ENV (shown in Figure 1). Twenty additional minutes were needed to add simulation of bandwidth and processing power measurements for the workers and to use these measurements for implement-

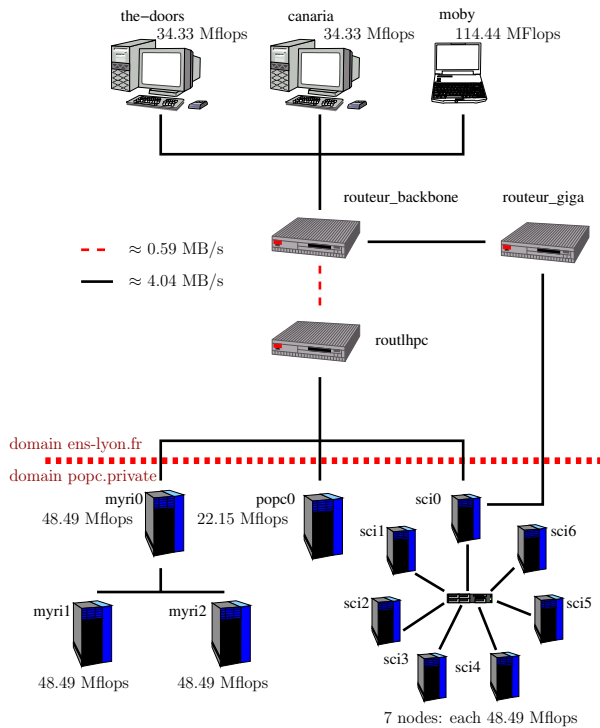


Fig. 1. Realistic platform imported from ENV.

ing the three aforementioned priority-based heuristics. In these simulations, the request data size was set to 100KB and the request computational cost (on the master) was set to 0.1 MFlops. The task data size was set to 10MB and the task computational cost was set to 1 000 MFlops. Network and processor speeds are shown in Figure 1.

A number of experiments can be conducted with the simulator and Figure 2 plots the performance of the four scheduling strategies on the simulated platform in terms of overall execution time versus task number (i.e. the time at which the x^{th} task completes). These simulations ran in less than four minutes on a PIII-733. The somewhat surprising result is that it is most efficient to prioritize workers according to bandwidth only. In other terms, the master should delegate work as fast as possible. Note that this type of result has been obtained in previous work [3] and that this case study simply confirms this phenomenon on a more realistic platforms.

This case study and other SIMGRID examples were used for educational purposes in a course on Parallel Algorithms and Architectures at the École Normale Supérieure de Lyon during the Fall 2002 semester.

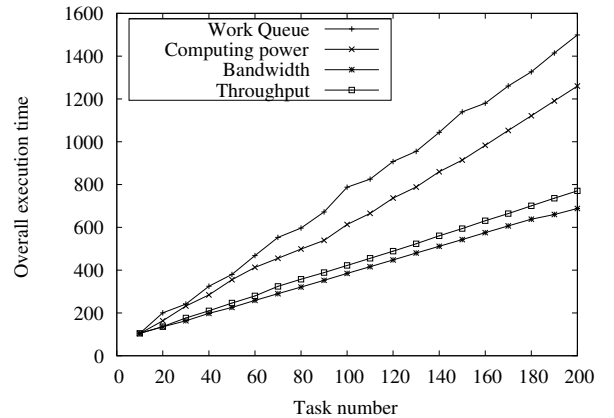


Fig. 2. Comparison of various strategies of independent task allocation.

VI. CONCLUSION

In this paper we have recognized two important limitations of simulations used for scheduling research: (i) there is no simulation standard; (ii) simulation methodologies are not appropriate for modern distributed computing platforms such as the Grid. We have identified a need for a simulation framework that makes it convenient for users to build scalable, configurable, extensible, and fast simulations for investigating novel scheduling techniques for these platforms. In this context we have presented SIMGRID v2, which addresses these issues. We have described the main concepts in SIMGRID and provided a detailed discussion of platform modeling issues. We have also demonstrated the usefulness of SIMGRID with a case study.

SIMGRID has already been used for a number of successful research endeavors [10, 41, 18, 43, 3, 4, 47] in our groups, and the SIMGRID user community is currently undergoing a dramatic expansion. SIMGRID is used for educational purposes in a course on Parallel Algorithms and Architectures at the École Normale Supérieure de Lyon. Future directions include mechanisms to import Brite [32] topologies and associated them link characteristics in an ENV [40] fashion for simulation in SIMGRID. We are also planning to create a repository where SIMGRID users can download pre-generated topologies that can be used directly in their simulations. The SIMGRID software has the potential for being of great value to the scheduling research community at large and these future directions will further expand its capabilities. All details, publications, and software for SIMGRID are available at: <http://grail.sdsc.edu/simgrid>.

REFERENCES

- [1] GT-ITM: Georgia Tech Internetwork Topology Models. <http://www.cc.gatech.edu/projects/gtitm/>.
- [2] BAJAJ, S., BRESLAU, L., ESTRIN, D., FALL, K., FLOYD, S., HALDAR, P., HANDLEY, M., HELMY, A., HEIDEMANN, J., HUANG, P., KUMAR, S., MCCANNE, S., REJAIE, R., SHARMA, P., VARADHAN, K., XU, Y., YU, H., AND ZAPPALA, D. Improving simulation for network research. Tech. Rep. 99-702, University of Southern California, 1999.
- [3] BEAUMONT, O., CARTER, L., FERRANTE, J., LEGRAND, A., AND ROBERT, Y. Bandwidth-centric allocation of independent tasks on heterogeneous platforms. In *International Parallel and Distributed Processing Symposium IPDPS'2002* (2002), IEEE Computer Society Press. Extended version available as LIP Research Report 2001-25.
- [4] BEAUMONT, O., LEGRAND, A., AND ROBERT, Y. Optimal algorithms for scheduling divisible workloads on heterogeneous systems. Tech. Rep. 2002-36, LIP, Oct. 2002.
- [5] BONALD, T., AND MASSOULIÉ, L. Impact of fairness on internet performance. In *SIGMETRICS/Performance* (2001), pp. 82–91.
- [6] BRAUN, T., SIEGEL, H., AND BECK, N. A Comparison of Eleven Static Heuristics for Mapping a Class of Independent Tasks onto Heterogeneous Distributed Computing Systems. *Journal of Parallel and Distributed Computing* 61 (2001), 810–837.
- [7] BUYYA, R., AND MURSHED, M. GridSim: A Toolkit for the Modeling and Simulation of Distributed Resource Management and Scheduling for Grid Computing. *The Journal of Concurrency and Computation: Practice and Experience (CCPE)* (2002). to appear.
- [8] CALVERT, K. L., DOAR, M. B., AND ZEGURA, E. W. Modeling internet topology. *IEEE Communications Magazine* 35, 6 (June 1997), 160–163.
- [9] CASANOVA, H. Simgrid: A Toolkit for the Simulation of Application Scheduling. In *Proceedings of the IEEE International Symposium on Cluster Computing and the Grid (CCGrid'01)* (May 2001), pp. 430–437.
- [10] CASANOVA, H., LEGRAND, A., ZAGORODNOV, D., AND BERMAN, F. Heuristics for scheduling parameter sweep applications in grid environments. In *Heterogeneous Computing Workshop* (2000), IEEE Computer Society Press, pp. 349–363.
- [11] CASANOVA, H., AND MARCHAL, L. A network model for simulation of grid application. Tech. Rep. 2002-40, LIP, Oct. 2002.
- [12] CHEN, Y., WINSLETT, M., KUO, S., CHO, Y., SUBRAMANIAM, M., AND SEAMONS, K. E. Performance modeling for the Panda array I/O library. In *Proceedings of Supercomputing '96* (1996), ACM Press and IEEE Computer Society Press.
- [13] CHIU, D. N. Some observations on fairness of bandwidth sharing. Tech. rep., Sun Microsystems, 1999.
- [14] COWIE, J., DODSON, B., ELKENBRACHT-HUIZING, R.-M., LENSTRA, A. K., MONTGOMERY, P. L., AND ZAYER, J. A world wide number field sieve factoring record: on to 512 bits. In *Advances in Cryptology - Asiacrypt '96* (1996), K. Kim and T. Matsumoto, Eds., vol. 1163 of *LNC3*, Springer Verlag, pp. 382–394.
- [15] DICKENS, P. M., HEIDELBERGER, P., AND NICOL, D. M. A distributed memory LAPSE: Parallel simulation of message-passing programs. In *Proceedings of the 8th Workshop on Parallel and Distributed Simulation (PADS '94)* (1994).
- [16] DOAR, M. A better model for generating test networks. In *Proceedings of Globecom '96* (Nov. 1996).
- [17] ENTROPIA. URL: <http://www.entropia.com>.
- [18] FAERMAN, M., BIRNBAUM, A., CASANOVA, H., AND BERMAN, F. Resource Allocation for Steerable Parallel Parameter Searches. In *Proceedings of the Grid Computing Workshop* (Nov. 2002).
- [19] FALOUTSOS, M., FALOUTSOS, P., AND FALOUTSOS, C. On power-law relationships of the internet topology. In *SIGCOMM* (1999), pp. 251–262.
- [20] FISHWICK, P. SimPack: Getting Started With Simulation Programming In C And C++. In *Proceedings of the Winter Simulation Conference* (1992), pp. 154–162.
- [21] FLOYD, S., AND FALL, K. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking* 7, 4 (1999), 458–472.
- [22] FOSTER, I., AND KESSELMAN, C., Eds. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, Inc., San Francisco, USA, 1999.
- [23] FOSTER, I., KESSELMAN, C., AND TUECKE, S. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of High Performance Computing Applications* 15, 3 (2001).
- [24] GOMES, F., FRANKS, S., UNGER, B., XIAO, Z., CLEARY, J., AND COVINGTON, A. SimKit: A High Performance Logical Process Simulation Class Library in C++. In *Proceedings of the 1995 Winter Simulation Conference* (December 1995), pp. 706–713.
- [25] HOWELL, F., AND R., M. SimJava: A Discrete Event Simulation Package for Java with Applications in Computer Systems Modelling. In *Proceedings of the First International Conference on Web-based Modelling and Simulation* (Jan 1998).
- [26] KIELMANN, T., BAL, H. E., MAASSEN, J., VAN NIEUWPOORT, R., EYRAUD, L., HOFMAN, R., AND VERSTOEP, K. Programming environments for high-performance grid computing: the albatross project. *Future Generation Computer Systems* (2002).
- [27] KWOK, Y., AND AHMAD, I. Benchmarking and Comparison of the Task Graph Scheduling Algorithms. *Journal of Parallel and Distributed Computing* 59, 3 (1999), 381–422.
- [28] LEGRAND, A., AND LEROUGE, J. Metasimgrid : Towards realistic scheduling simulation of distributed applications. Tech. Rep. 2002-28, LIP, July 2002.
- [29] LIU, J., AND NICOL, D. M. *DaSSF 3.1 User's Manual*, Apr. 2001.
- [30] MASSOULIÉ, L., AND ROBERTS, J. Bandwidth sharing: Objectives and algorithms. In *INFOCOM (3)* (1999), pp. 1395–1403.
- [31] MATHIS, M., SEMKE, J., AND MAHDAVI, J. The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communications Review* 27, 3 (1997).
- [32] MEDINA, A., LAKHINA, A., MATTA, I., AND BYERS, J. BRITE: Universal topology generation from a user's perspective. Tech. Rep. 2001-003, Computer science department Boston University, 1 2001.
- [33] MEDINA, A., MATTA, I., AND BYERS, J. On the origin of power laws in internet topologies. Tech. Rep. 2000-004, Computer Science Department, Boston University, 20, 2000.
- [34] MILLER, A., NAIR, R., AND ZHANG, Z. JSIM: A Java-Based Simulation and Animation Environment. In *Proceedings of*

- the 30th Annual Simulation Symposium (ANSS'97) (April 1997), pp. 31–42.
- [35] The network simulator - ns-2. <http://www.isi.edu/nsnam/ns>.
 - [36] Objective modular network testbed. <http://www.hit.bme.hu/phd/vargaa/omnetpp.htm>.
 - [37] PRIME. URL: <http://www.mersenne.org>.
 - [38] SCHWETMAN, H. CSIM: A C-based, process oriented simulation language. In *Proceedings of the 1986 Winter Simulation Conference* (December 1986), pp. 387–396.
 - [39] SETI. URL: <http://setiathome.ssl.berkeley.edu>.
 - [40] SHAO, G., BERMAN, F., AND WOLSKI, R. Using effective network views to promote distributed application performance. In *International Conference on Parallel and Distributed Processing Techniques and Applications* (June 1999).
 - [41] SMALLE, S., CASANOVA, H., AND BERMAN, F. Tunable Online Parallel Tomography. In *Proceedings of SC'01* (Nov. 2001). best student paper award.
 - [42] SONG, H. J., LIU, X., JAKOBSEN, D., BHAGWAN, R., ZHANG, X., TAURA, K., AND CHIEN, A. A. The microgrid: a scientific tool for modeling computational grids. In *Supercomputing* (2000).
 - [43] SU, A., CASANOVA, H., AND BERMAN, F. Utilizing DAG Scheduling Algorithms for Entity-Level Simulations. In *Proceedings of the 10th Symposium on High Performance Computing 2002 (HPC'02)* (Apr. 2002).
 - [44] SUGIH, J. W. Inet-3.0: Internet topology generator. Tech. Rep. 456-02, Computer science department University of Michigan, 2002.
 - [45] TOH, S. SimC: A C Function Library for Discrete Simulation. In *Proceedings of the 11th Workshop in Parallel and Distributed Simulation* (January 1993), pp. 18–20.
 - [46] ULLMAN, J. NP-complete Scheduling Problems. *Journal of Computer and System Sciences* 10 (1975), 434–439.
 - [47] YANG, Y., AND CASANOVA, H. Multi-Round Algorithm for Scheduling Divisible Workload Applications: Analysis and Experimental Evaluation. Tech. Rep. CS2002-0721, Dept. of Computer Science and Engineering, University of California, San Diego, 2002.