

# Scheduling in hybrid clouds

Luiz F. Bittencourt, *Member, IEEE*, Edmundo R. M. Madeira, *Member, IEEE*, and Nelson

L. S. da Fonseca, *Senior Member, IEEE*

Institute of Computing – University of Campinas – UNICAMP

Av. Albert Einstein, 1251, 13083-852 – Campinas – São Paulo – Brazil

e-mail: {bit, edmundo, nfonseca}@ic.unicamp.br

## Abstract

Schedulers for cloud computing determine on which processing resource jobs of a workflow should be allocated. In hybrid clouds, jobs can be allocated either on a private cloud or on a public cloud on a pay per use basis. The capacity of the communication channels connecting these two types of resources impact the makespan and the cost of workflows execution. This paper introduces the scheduling problem in hybrid clouds presenting the main characteristics to be considered when scheduling workflows, as well as a brief survey of some of the scheduling algorithms used in these systems. To assess the influence of communication channels on job allocation, we compare and evaluate the impact of the available bandwidth on the performance of some of the scheduling algorithms.

## I. INTRODUCTION

Cloud computing has attracted an increasing number of users because it offers computational capabilities as services on a pay-per-use basis. Studies conducted by Gartner<sup>1</sup> estimate a potential amount of U\$150 billion market for cloud computing by 2013. Such a huge market is the consequence of a business model that offers high performance and low costs. Indeed, a survey<sup>2</sup> of 3,645 users of cloud computing services conducted by the Computer Sciences Corporation in eight countries between October 2011 and November 2011, reported that improved data center efficiency and lower operational costs are the main reasons for the adoption of cloud computing solutions.

Cloud providers offer computing and storage resources, and platforms for software development and execution, as well as software interfaces accessible throughout the network. Three models of cloud services

<sup>1</sup><http://www.gartner.com/id=914826>

<sup>2</sup>[http://assets1.csc.com/newsroom/downloads/CSC\\_Cloud\\_Usage\\_Index\\_Report.pdf](http://assets1.csc.com/newsroom/downloads/CSC_Cloud_Usage_Index_Report.pdf)

are commonly available: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS) [1]. In SaaS, the clients use applications but cannot control the host environment. Google Apps and Salesforce.com are examples of this model. In PaaS, the platform is typically an application framework and clients use a hosting environment for their applications. Examples of PaaS are the Google App Engine and Amazon Web Services. In IaaS, the clients use computing resources such as processing power and storage and they can also control the environment and the deployment of applications. Amazon Elastic Compute Cloud (EC2), Globus Nimbus Toolkit, and Eucalyptus are good examples of this service model. In summary, clients can use/run applications from a SaaS cloud; both develop and run their applications on a development platform provided by a PaaS cloud; or extend their computational capacity by leasing computing resources from an IaaS cloud.

Moreover, clients can execute most applications using their own computing infrastructure (private cloud), and yet lease service from a cloud provider (public cloud) on demand. It was reported that 48% of United States government agencies have moved at least one workflow to a cloud provider following the federal cloud computing strategy published in February 2011. The operation of such hybrid cloud involves two fundamental questions: (i) what resources should be leased and (ii) which tasks should be executed on the leased resources. These answers are provided by a scheduler, a fundamental component of distributed computing systems including clouds and grids [2].

These questions are answered considering the capacities of the communication links connecting the available resources. Slow communication channels increase delays, thus increasing the execution time (makespan) of applications, with bounds typically negotiated in service level agreements. Understanding the impact of network delays and the costs on scheduling decisions is, thus, fundamental for cloud service provisioning.

In line with that, this paper provides a brief survey of scheduling algorithms for hybrid clouds and the impact of communication networks on scheduling decisions. First, the problem of scheduling tasks and services in clouds is explained (Section II), which is then followed by a comparison of scheduling algorithms for hybrid clouds (section III). At last, the impact of communication links on schedules is assessed (Section IV).

## II. SCHEDULING IN CLOUDS

Applications and services can be decomposed into set of smaller components, called jobs. For example, an application that processes a large image can decompose this image into smaller ones for parallel

processing by distinct jobs. The logical sequence of the jobs of an application is called workflow, which is commonly represented by directed acyclic graph (DAG). The nodes of a DAG represent the jobs of a workflow while arcs represent their data dependencies. A job can be executed only after the data it depends on has been produced and sent to the resource where it will be executed. Such applications can be found in a variety of fields, such as Physics (astronomy, thermodynamics), Bioinformatics (DNA sequencing, proteomics), Chemistry (protein dynamics), and Computer Science (computer vision, image processing).

The workflow of two real applications are illustrated: The Montage [3] (Figure 1(a)) and the Laser Interferometer Gravitational Wave Observatory (LIGO) [4] (Figure 1(b)). Montage consists of an image application that creates mosaics of the sky in astronomy research. The size of the workflow depends on the squared degree size of the part of the sky to be generated, and it can produce an output of 86 TB data set involving 17 hierarchical workflows, each with 900 sub workflows<sup>3</sup>. LIGO is a project used to detect gravitational waves through a network of gravitational-wave detectors, and its workflow often requires on the order of a terabyte of data to produce meaningful results<sup>3</sup>.




Fig. 1. Examples of workflow applications.

The computational demands of applications such as Montage and LIGO can easily overwhelm the available computational power of private clouds. Moreover, their execution time can be prohibitive. The cloud computing paradigm is quite effective for dealing with such problem by providing virtually unbounded on-demand resources.

Figure 2 illustrates a hybrid IaaS cloud composed of the resources of the private cloud as well as those of one or more public IaaS clouds. A hybrid cloud scheduler must decide which resources should be leased

<sup>3</sup>Source: <http://pegasus.isi.edu/>

from the public clouds to guarantee the execution of the workflow within the specified maximum execution time (deadline). After the submission of workflow by a user, a broker runs the scheduling algorithm to start the decision making process. Besides deciding which resources will be used, the scheduler also determines which part of the workflow will run in each cloud provider.




Fig. 2. Hybrid cloud infrastructure and workflow submission.

One challenging issue in hybrid clouds is how interfaces can be provided to interact automatically with different existing public clouds, so that the broker can gather information about resources and the workflow executed and monitored in a variety of public cloud infrastructures. Some projects, such as the JClouds ([www.jclouds.org](http://www.jclouds.org)) try to solve this problem by providing portable abstractions to several existing cloud providers. Another challenge involves the consideration of security requirements of the applications which can reduce the pool of potential hosts for scheduling jobs.

The scheduling problem involved is known to be NP-Complete in general, including the scheduling of workflows in heterogeneous computer systems discussed in this paper. Scheduling algorithms often utilize heuristics and optimization techniques to try to obtain a near optimal schedule.

The input of scheduling algorithm must include the DAG that represents the workflow of jobs and their dependencies, as well as information about the target system, including the processing capacity of each resource and the available capacities of the network links. This information is obtained from a resource

information repository in the private cloud. Moreover, scheduling algorithms for clouds are usually cost-aware, so that the information about the cost per time unit of usage of each resource must be available. Providing with this information, the scheduling algorithm is capable of estimating the workflow makespan and its execution costs.

Figure 3 illustrates the inputs necessary for a scheduler and gives an example of a schedule. The DAG contains information on the computational requirements of its jobs as well as information on the amount of bytes to be transmitted to resolve each data dependency. In the example, we use millions of instructions (MI) for specifying job computational costs, MI/s for resource processing capacities, Megabytes for data dependencies, and Megabits/s for link bandwidths. The scheduler combines all of this information to compute how long each job takes to run on each resource, how long each data transmission would take according to the resource assigned to each job, and how much a given job assignment would cost. The scheduling algorithm runs when a workflow is submitted for execution, and the necessary resources are allocated on demand to run the workflow.

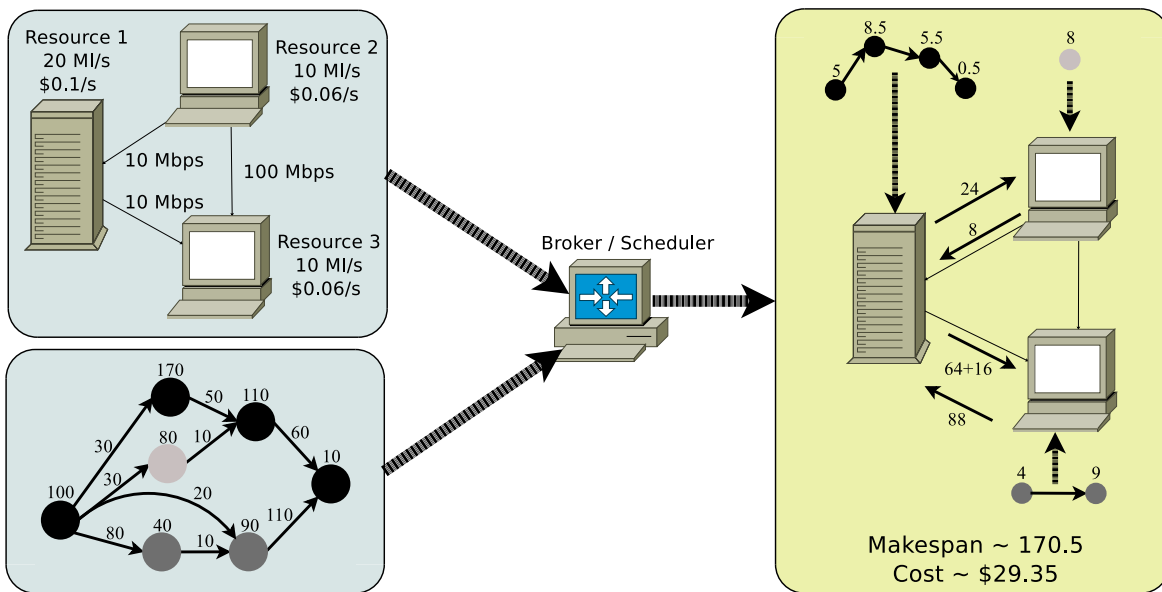


Fig. 3. Scheduler inputs, with information about the DAG and the target system, and a possible resulting schedule.

The right-hand side of Figure 3 illustrates a hypothetical schedule for this example. The critical path of the DAG was scheduled on resource 1 ( $R1$ ), which is the fastest available. Data transmission takes 24 seconds between  $R1$  and  $R2$  in order to fulfill data requirements for the job scheduled on  $R2$ . This job returns the results to  $R1$  so it can be utilized by the third job on the critical path, which demands a transmission of 8 seconds. Longer data transmissions occur between  $R1$  and  $R3$ , since jobs running on those resources have stronger data dependencies. By considering all the computations performed and all the

transmission delays, the resulting makespan is 170.5 seconds and cost is 29.35 dollars. The consideration of communication delays is important for the minimization of costs since a job which receives data from its predecessors needs to be active in that resource, thus consuming processing time and money.

### III. SCHEDULERS FOR CLOUDS

Cloud computing evolved from grid computing, service oriented computing, and virtualization paradigms. This means that scheduling algorithms developed for these type of systems can be also used in clouds. Scheduling algorithms can be distinguished by their main characteristics, such as:

- **Target system:** the system for which the scheduling algorithm was developed, which can be a heterogeneous system, a grid, or a cloud computing system.
- **Optimization criterion:** Makespan and cost are the main metrics specified by cloud user and considered by schedulers in the decision making process.
- **Multi-core awareness:** Computer systems can have multiple cores which should be considered by scheduling algorithms in resource selection.
- **On-demand resources:** Resources can be leased either on-demand or for long terms. The on-demand leasing of resources is treated by the scheduling algorithm as a “single expense” during the execution of the workflow.
- **Reserved resources:** The algorithm should consider the use of a resource reserved for a long term.
- **Levels in service level agreement (SLA):** The scheduling algorithm should consider that SLAs can be organized hierarchically. SLAs with a single-level allow clients and providers to interact directly to negotiate resource capacities and prices. When multiple levels, the scheduling algorithm can run in an intermediate facility between the IaaS cloud provider and the final client. By doing so, costs can be decreased.

Table I lists various workflow scheduling algorithms and compares their characteristics and applicability for cloud scheduling. Although not all scheduling algorithms used in clouds were conceived for these systems, recently some scheduling algorithms specially designed to hybrid clouds have been proposed [5], [6]. Next, we briefly describe some scheduling algorithms commonly used.

The Heterogeneous Earliest Finish Time (HEFT) [7] scheduling algorithm was designed for heterogeneous computing systems. Since it was developed before the advent of cloud computing and utility grids, it does not consider monetary costs. Its objective is to minimize the workflow makespan.

TABLE I  
SCHEDULING ALGORITHMS CHARACTERISTICS

Algorithm	Target system	Optimization criteria	Multi-core aware	On-demand resources	Reserved resources	SLA levels
HEFT [7]	Heterogeneous	minimize makespan	No	No	No	No
MDP [8]	Utility Grid	minimize cost within deadline	No	Yes	No	Single-level
PCP [9]	Utility Grid	minimize cost within deadline	No	Yes	No	Single-level
Pandey [10]	Cloud	minimize cost	No	Yes	No	Single-level
Wu [11]	Cloud	minimize cost	No	Yes	No	Single-level
HCOC [5]	Cloud	minimize cost within deadline	Yes	Yes	No	Single-level
Genez [6]	Cloud	minimize cost within deadline	Yes	Yes	Yes	Two-level

The deadline-driven cost-minimization algorithm [8], or the Deadline-Markov Decision Process (MDP), breaks the DAG into partitions, assigning a maximum finishing time for each partition according to the deadline set by the user. Based on this time, each partition is scheduled for that resource which will result in the lowest cost and earliest estimated finishing time. This algorithm works with on-demand resource reservation.

Abrishami et al. [9] presented the Partial Critical Paths (PCP) algorithm which schedules the workflow in a backwards fashion. Constraints are added to the scheduling process when such scheduling of jobs in a partial critical path fails, so that the algorithm will be re-started. This algorithm presents the same characteristics as does MDP, although it involves greater time complexity, since a relatively large number of re-schedulings can be demanded during the execution of the algorithm.

The self-adaptive global search optimization technique called particle swarm optimization (PSO) is utilized to schedule workflows in the algorithm proposed in [10] and which was developed to work in clouds with a single-level SLAs and on-demand resource leasing. It considers neither multi-core resources nor workflow deadlines, but focuses solely on monetary cost minimization.

The Hybrid Cloud Optimized Cost (HCOC) algorithm [5], schedules workflows in hybrid clouds by first attempting costless local scheduling using HEFT. If the local scheduling cannot meet the deadline, the algorithm selects jobs for scheduling in resources from the public cloud. When selecting resources from the public cloud, the HCOC algorithm considers the relation between the number of parallel jobs being scheduled and the number of cores of each resource as well as deadlines, performance, and cost. As with MDP algorithm, the objective is to minimize the financial cost obeying the deadlines stipulated

by the user in a single-level SLA contract.

In [6], the workflow scheduling problem was formulated as an integer linear program that considers the leasing of reserved and on-demand resources from multiple IaaS providers according to a two-level SLA. The scheduler can run either in a SaaS or PaaS cloud provider, and receive workflow execution requests with deadlines from its clients (first SLA level), but it can also lease resources from multiple IaaS providers (second SLA level).

A common characteristic of the above mentioned algorithms is that they do not take into account the fluctuation of the prices of resource allocation due to the varying demand of resources. In order to reduce costs, a scheduler could allocate or even re-allocate jobs when prices are low.

#### IV. IMPACT OF AVAILABLE BANDWIDTH ON SCHEDULING

The available bandwidth in channels connecting processing resources of the hybrid cloud impacts the makespan and the cost of a schedule. This section discusses this impact as well as the effectiveness of the HEFT [7], MDP [8], and HCOC [5] in scheduling workflows in hybrid clouds. HEFT is a well-known scheduling algorithm for heterogeneous computing systems and it aims at makespan minimization. The MDP scheduling algorithm was designed for utility grids and it is often used in the literature for comparisons with cost-based algorithms. Moreover, the HCOC scheduling algorithm is a recent development specially for hybrid clouds. By evaluating these algorithms, we can analyze the adequacy of scheduling algorithms designed for clouds with those designed for other systems but used for scheduling in cloud systems.

In the evaluation, the capacity of intra cloud channels as well as the inter clouds channels were varied so that their impact on the scheduling efficacy could be evaluated.

Three thousand different Montage and LIGO DAGs were generated with computation demands varying in the interval  $[5 \times 10^5, 4 \times 10^6]$  millions of instructions (MI), and communication demands varying randomly in the interval  $[60, 500]$  Megabytes. The deadline for completion of the workflow was set to 2.5 times the duration taken to compute the critical path of the DAG. Such value was set since it was observed in previous experiments that it leads to the highest number of workflow completions by the three scheduling algorithms. If the predicted makespan is shorter than the deadline, the solution given by HEFT is adopted, avoiding cost increase due to the leasing of resources.

The number of processing resources in the private cloud was varied from 1 to 10 with the processing capacities for each randomly taken from the interval  $[10^4, 10^5]$  MIPS. For the public cloud, there were 4

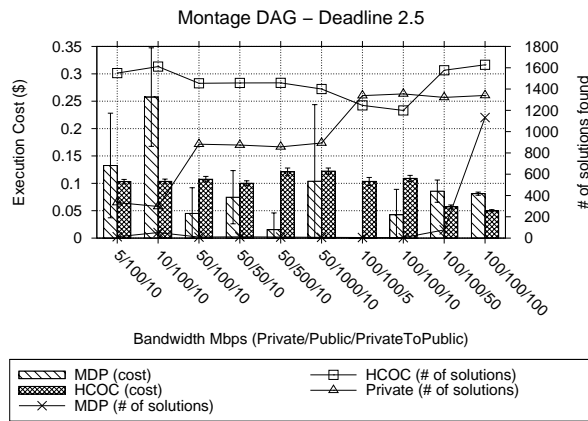


types of resources with computing unit capacities and leasing costs equivalent to the Amazon EC2 *small*, *large*, *extra large*, and *extra large high CPU* on-demand instance types. In the experiments undertaken, each computing unit was randomly drawn from the interval  $[10^4, 7 \times 10^4]$  MIPS, referring to the *small* instance type computing power. The 4 types of resources were: *small* (1 core of 1 computing unit, \$0.085 per hour); *large* (2 cores of 2 computing units each, \$0.34 per hour); *extra large* (4 cores of 2 computing units each, \$0.68 per hour); *extra large high cpu* (8 cores of 2.5 compute units each, \$0.68 per hour). Moreover, in the experiments, the topologies of private cloud and public cloud networks were fully connected graphs, and the private cloud was connected to each public cloud by an inter cloud link.

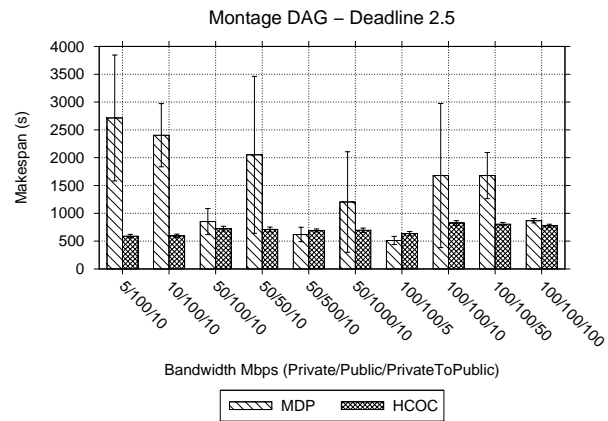
Figure 4 shows the makespan, execution cost and number of solutions found for the Montage and LIGO DAGs by the three scheduling algorithms evaluated. The bars show the average cost and average makespan with 95% confidence interval. The three curves represent the total number of successful schedules (i.e., the number of schedules with a makespan lower than the deadline) achieved by each algorithm. The increase in the number of solutions found by the HEFT algorithm as a function of the increase in available bandwidth in private cloud reveals that the private bandwidth is of major importance for scheduling workflows in the private cloud. Moreover, an increase in bandwidth reduces the costs when the target system is a hybrid cloud.

The number of solutions found by the HCOC and MDP algorithms increases when the bandwidth between the private and public clouds increases, since new solutions for problem instances which cannot be satisfied in the private cloud, can now be found using the resources of the public cloud. The small influence of public cloud bandwidth in the number of solutions found by the HCOC and MDP algorithms can be explained by the tendency of these algorithms to group dependent jobs to minimize communication costs.

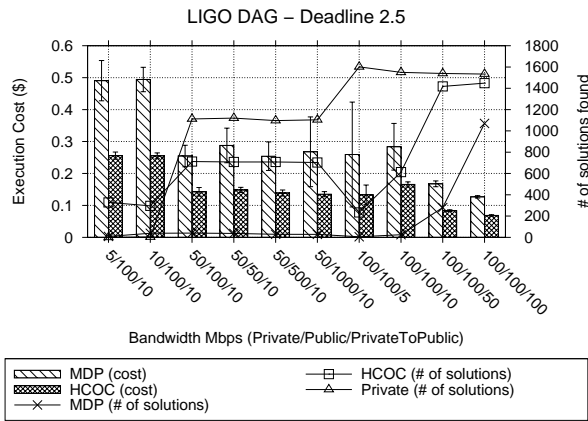
The inter cloud available bandwidth reduces the makespan. Moreover, when numerous solutions are found in the hybrid cloud processing, the average cost of the schedule decreases with an increase in the bandwidth between the private and public clouds. Schedules which are unfeasible in private cloud processing can be implemented at a low cost when the inter cloud bandwidth increases. The average makespan value tends to be close to the deadline, since the scheduling algorithms try to minimize cost as long as the makespan is kept lower than the deadline. Makespan closer to the deadline can use less expensive (and slower) resources from public clouds, and more resources from the private cloud which hardly cost anything compared to using public clouds. Results using other DAGs (Random, AIRSN spatial



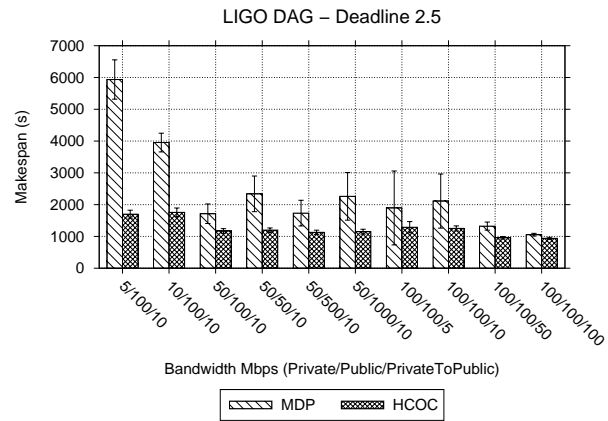
(a) Montage - average cost



(b) Montage - average makespan



(c) LIGO - average cost



(d) LIGO - average makespan

Fig. 4. Results for the Montage and LIGO DAGs.

normalization, the example DAG in MDP paper, and Chimera virtual data) not discussed here, reinforce this conclusion.

Comparison of the makespan produced by the MDP algorithm with that given by the HCOC algorithm shows that the former, designed for utility computing, underperforms the HCOC, which was specially designed for hybrid clouds.

Moreover, the cost demanded by HCOC for processing the LIGO workflow is lower than that demanded by MDP. The low number of solutions found by MDP when processing the Montage DAG do not allow a significant comparison of the costs for this DAG.

## V. CONCLUSION

This paper has discussed the problem of scheduling applications and services in hybrid clouds. A brief comparison of some of the scheduling algorithms used in hybrid grids has been provided, with the results suggesting a prominent importance of communication capacity when scheduling workflows in hybrid

clouds, especially that of the communication channels between the private and public clouds. Such channels are usually located in the Internet backbone, which increases the challenges involved in the development of communication-aware scheduling algorithms since the available bandwidth of these links fluctuates widely. Moreover, given the importance of the inter-cloud communication channels, the development of communication-aware or even communication-driven scheduling algorithms is of paramount importance to provide quality of service and competitive costs for hybrid clouds. Furthermore, an efficient algorithm developed for efficient use on utility grids (MDP) may not be as efficient for hybrid clouds as are those algorithms developed specially for such systems (HCOC). This superior performance of the HCOC scheduler may partially be due to its multi-core awareness, which is clearly a characteristic requiring consideration in hybrid cloud computing.

#### ACKNOWLEDGMENTS

The authors would like to thank FAPESP (2009/15008-1) and CNPq for the financial support.

#### REFERENCES

- [1] M. Dikaiakos, D. Katsaros, P. Mehra, G. Pallis, and A. Vakali, "Cloud computing: Distributed internet computing for IT and scientific research," *IEEE Internet Computing*, vol. 13, no. 5, Sep.-Oct. 2009, pp. 10–13.
- [2] D. M. Batista and N. L. S. da Fonseca, "A survey of self-adaptive grids," *IEEE Communications Magazine*, vol. 48, no. 7, Jul. 2010, pp. 94–100.
- [3] Y. Gil, E. Deelman, M. Ellisman, T. Fahringer, G. Fox, D. Gannon, C. Goble, M. Livny, L. Moreau, and J. Myers, "Examining the challenges of scientific workflows," *IEEE Computer*, vol. 40, no. 12, Dec. 2007, pp. 24–32.
- [4] E. Deelman, C. Kesselman, G. Mehta, L. Meshkat, L. Pearlman, K. Blackburn, P. Ehrens, A. Lazzarini, R. Williams, and S. Koranda, "Grifhyn and ligo, building a virtual data grid for gravitational wave scientists," in *11th IEEE International Symposium on High Performance Distributed Computing*, 2002, pp. 225–234.
- [5] L. F. Bittencourt and E. R. M. Madeira, "HCOC: a cost optimization algorithm for workflow scheduling in hybrid clouds," *Journal of Internet Services and Applications*, vol. 2, no. 3, Dec 2011, pp. 207–227.
- [6] T. A. L. Genez, L. F. Bittencourt, and E. R. M. Madeira, "Workflow Scheduling for SaaS / PaaS Cloud Providers Considering Two SLA Levels," in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Apr. 2012.
- [7] H. Topcuoglu, S. Hariri, and M.-Y. Wu, "Performance-effective and low-complexity task scheduling for heterogeneous computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 13, no. 3, 2002, pp. 260–274.
- [8] J. Yu, R. Buyya, and C. K. Tham, "Cost-based scheduling of scientific workflow applications on utility grids," in *International Conference on e-Science and Grid Computing*, Jul. 2005, pp. 140–147.
- [9] S. Abrishami, M. Naghibzadeh, and D. Epema, "Cost-driven scheduling of grid workflows using partial critical paths," in *11th IEEE/ACM International Conference on Grid Computing (GRID)*, Oct. 2010, pp. 81–88.

- [10] S. Pandey, L. Wu, S. Guru, and R. Buyya, "A particle swarm optimization-based heuristic for scheduling workflow applications in cloud computing environments," in *24th IEEE International Conference on Advanced Information Networking and Applications (AINA)*, Apr. 2010, pp. 400–407.
- [11] Z. Wu, Z. Ni, L. Gu, and X. Liu, "A revised discrete particle swarm optimization for cloud workflow scheduling," in *International Conference on Computational Intelligence and Security (CIS)*, Dec. 2010, pp. 184–188.

## BIOGRAFIES

Luiz F. Bittencourt received his Bachelor's degree in Computer Science from the Federal University of Parana, Brazil, in 2004. He received his Masters degree in 2006 and PhD degree in 2010 from the University of Campinas (UNICAMP), Brazil. He is currently an Assistant Professor at UNICAMP, and his main interests are in the areas of virtualization and scheduling in grids and clouds.

Edmundo R. M. Madeira is a Full Professor at the University of Campinas (UNICAMP), Brazil. He received his Ph.D. in Electrical Engineering from UNICAMP. He has published over 120 papers in national and international conferences and journals. He was the General Chair of the 7th Latin American Network Operation and Management Symposium (LANOMS'11), and he is Technical Program Co-chair of the IEEE LatinCloud'12. His research interests include network management, future Internet and cloud computing.

Nelson L. S. da Fonseca obtained his Ph.D. degree from the University of Southern California in 1994. He is Full Professor at the State University of Campinas. He has published 300+ refereed papers and supervised 50+ graduate students. He is currently ComSoc Vice President Member Relations. He served as EiC of the IEEE Communications Surveys and Tutorials, Member-at-Large of the ComSoc Board of Governors, Director of Latin America Region and Director of On-Line Services.