



Special Issue on ‘Performance Evaluation, Modeling and Optimization of Computer Systems’

José L (Sevi) Sevilla

This Special Issue of *SIMULATION: Transactions of The Society for Modeling and Simulation International* is dedicated to the performance evaluation, modeling and optimization of computer systems. The aim of this special issue is to publish novel research work that contributes to the development of methodologies, techniques and tools, as well as their application in state-of-the-art computer systems.

Some of the papers included were submitted specially for this issue, while others are extended versions of a selection of the best papers accepted and presented at the 2009 International Symposium on Performance Evaluation of Computer and Telecommunication Systems, SPECTS 2009. These papers have undergone a new round of reviews to guarantee that all of the papers included in this Special Issue have been rigorously peer-reviewed.

The five articles of this Special Issue cover a variety of topics, including: performance evaluation of scheduling algorithms in clusters and distributed systems; reconstruction of load properties from observable traffic; modeling of workflows in heterogeneous distributed systems; and simulation of the performance of Flash memories. This wide range of topics reflects the many facets of this field, which has significantly progressed in the last years.

The first paper, ‘Gang Scheduling in a two-cluster system implementing migrations and periodic feedback’, is authored by Zafeirios C Papazachos and Helen D Karatza. They consider scheduling gangs in a homogeneous multi-cluster system, in the presence of critical sporadic jobs that require immediate execution upon their arrival. Several scheduling algorithms, responsible for allocating the available system resources to the existing jobs, are examined, and their performance is studied using a simulation model. This study includes the effect of migrating tasks from a local queue to the head of another queue, both locally and through the grid. In addition, the paper studies the performance of a grid dispatching algorithm which allocates gangs to the available clusters based on periodic feedback from the clusters.

In the second paper, ‘Reconstructing arrival processes to discrete queueing systems by inverse load transformation’, by Stephan Heckmüller and Bernd E Wolfinger, the

problem of recovering arrival processes to discrete G/D/1 queueing systems from observable traffic is studied. Independent normally distributed arrivals and autoregressive processes are considered, which are commonly used as a model for highly aggregated traffic such as that of computer networks. Closed-form expressions are obtained for the normal distribution, while for the case of autoregressive processes, a novel use of the Tobit regression model (commonly used in social and economic sciences) allows estimates of the process parameters to be obtained with high accuracy. The estimation procedures are generalized to distributions other than the normal distribution by using the Buckley–James estimator. Finally, the authors demonstrate through simulation that the proposed methods can be used in realistic scenarios such as polling mechanisms in wireless local area networks.

In the third paper, ‘Modeling and simulation of distributed computing workflows in heterogeneous network environments’, Qishi Wu and Yi Gu propose a simulation system to study the execution dynamics of distributed computing workflows. The simulator, named SDEDS (Simulation of Dynamic Execution of Distributed Systems), is designed for the performance evaluation of large-scale workflow systems, and it is implemented using multi-threaded programming and an event-driven control structure. It simulates the dynamic execution process of distributed systems with data execution on computer nodes and data transfer along network links, and evaluates the network performance of workflow scheduling or mapping algorithms. Experimental observations collected in real networks and theoretical analysis are used to confirm the simulation-based performance measurements.

The fourth paper, ‘An effective duplication-based task-scheduling algorithm for heterogeneous systems’, is authored by Mahsa Hosseinzadeh and Hadi Shahriar Shahhoseini. They present a novel task scheduling

Corresponding author:

José L (Sevi) Sevilla, Associate Professor, Department of Computer Architecture University of Seville, Seville, Spain.
Email: sevi@atc.us.es

algorithm called Earliest Starting and Finishing time Duplication-based (ESFD) algorithm, which is a combination of list-scheduling-based and duplication-based mechanisms. It is intended to schedule a directed acyclic graph onto a heterogeneous distributed system, and considers all parameters associated with the processors, the task, and its immediate predecessors. The ESFD algorithm is a three-phase algorithm with a dynamic phase to assign a priority to each task. Simulation results on both random graphs and real application graphs are used to evaluate the ESFD algorithm as well as to compare it with well-known algorithms. Simulation results show that the ESFD algorithm outperforms these algorithms, with the improvements being related to both minimizing the idle times and reducing the communication costs between processors. The proposed algorithm is significantly efficient in the case of high communication costs between processors in heterogeneous computing systems.

Finally, in the fifth paper, entitled 'Investigating Flash memory wear levelling and execution modes' and authored by Soraya Zertal and Peter Harrison, the impact of wear

levelling on a Flash storage package and its access operations' execution modes is investigated. The authors propose simple static mapping functions that ensure good wear levelling for uniformly distributed and hot-spot accesses to the storage space. Also, for each access mode, different preemptive and non-preemptive priority schemes are considered with a range of Input/Output profiles representing main real applications behavior. The impact of the execution modes on the performance of the Flash memory is studied through simulation both qualitatively and quantitatively.

I would like to finish with thanks to all of the authors and reviewers for their valuable contributions to this special issue. I would also like to thank Professor Mohammad S Obaidat, founder of the International Symposium on Performance Evaluation of Computer and Telecommunication Systems, for his dedication and advice. I hope that the papers selected in this special issue will become useful resources for researchers and practitioners in the area of Performance Evaluation, Modeling and Optimization of Computer Systems.