

2017 IEEE 25th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2017)

**Banff, Alberta, Canada
20-22 September 2017**



**IEEE Catalog Number: CFP17010-POD
ISBN: 978-1-5386-2765-5**

**Copyright © 2017 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP17010-POD
ISBN (Print-On-Demand):	978-1-5386-2765-5
ISBN (Online):	978-1-5386-2764-8
ISSN:	1526-7539

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2017 IEEE 25th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems

MASCOTS 2017

Table of Contents

General Chair's Welcome.....	viii
Program Chair's Welcome.....	ix
Organizing Committee.....	x
Program Committee.....	xi
Additional Reviewers.....	xiii

Session 1: GPUs

Performance Evaluation of Priority Queues for Fine-Grained Parallel Tasks on GPUs	1
<i>Nikolai Baudis, Florian Jacob, and Philipp Andelfinger</i>	
Nexus: Bringing Efficient and Scalable Training to Deep Learning Frameworks	12
<i>Yandong Wang, Li Zhang, Yufei Ren, and Wei Zhang</i>	
Characterizing Temperature, Power, and Soft-Error Behaviors in Data Center Systems: Insights, Challenges, and Opportunities	22
<i>Bin Nie, Ji Xue, Saurabh Gupta, Christian Engelmann, Evgenia Smirni, and Devesh Tiwari</i>	

Session 2: Storage / Cloud

Cost-Benefit Tradeoffs of Content Sharing in Personal Cloud Storage	32
<i>Glauber Gonçalves, Alex Borges Vieira, Idilio Drago, Ana Paula Couto da Silva, and Jussara M. Almeida</i>	
MPLEX: In-Situ Big Data Processing with Compute-Storage Multiplexing	43
<i>Joy Rahman and Palden Lama</i>	

Lightweight Replication Through Remote Backup Memory Sharing for In-memory Key-Value Stores	53
<i>Yandong Wang, Li Zhang, Michel Hack, Yufei Ren, and Min Li</i>	

Session 3: Data Centers, Big Data, and HPC

Delay Versus Stickiness Violation Trade-Offs for Load Balancing in Large-Scale Data Centers	63
<i>Qingkai Liang and Sem Borst</i>	
PTPmesh: Data Center Network Latency Measurements Using PTP	73
<i>Diana Andreea Popescu and Andrew W. Moore</i>	
Model-Based Performance Evaluation of Batch and Stream Applications for Big Data	80
<i>Johannes Kroß and Helmut Krcmar</i>	
Toward Managing HPC Burst Buffers Effectively: Draining Strategy to Regulate Bursty I/O Behavior	87
<i>Kun Tang, Ping Huang, Xubin He, Tao Lu, Sudharshan S. Vazhkudai, and Devesh Tiwari</i>	

Session 4A: Modeling

Swimming with Fishes and Sharks: Beneath the Surface of Queue-Based Ethereum Mining Pools	99
<i>Alexei Zamyatin, Katinka Wolter, Sam Werner, Peter G. Harrison, Catherine E. A. Mulligan, and William J. Knottenbelt</i>	
Testing Black-Box SDN Applications with Formal Behavior Models	110
<i>Jiangyuan Yao, Zhiliang Wang, Xia Yin, Xingang Shi, Yahui Li, and Chongrong Li</i>	
Free Energy Approximations for CSMA Networks	121
<i>Benny Van Houdt</i>	

Session 4B: Measurement and Simulation

QDIME: QoS-Aware Dynamic Binary Instrumentation	132
<i>Pansy Arafa, Guy Martin Tchamgoue, Hany Kashif, and Sebastian Fischmeister</i>	
iQ: An Efficient and Flexible Queue-Based Simulation Framework	143
<i>Damian Roca, Daniel Nemirovsky, Marc Casas, Miquel Moreto, Mateo Valero, and Mario Nemirovsky</i>	
Simulating Exploits for the Creation and Refinement of Detection Signatures	150
<i>Victor C. Valgenti, Ya-Wen Lin, Atsuhiko Suzuki, and Min Sik Kim</i>	

Session 5: Power / Energy

Emulating the Power Consumption Behavior of Server Workloads Using CPU Performance Counters	157
<i>Norbert Schmitt, Joakim Von Kistowski, and Samuel Kounev</i>	
Optimal Peak Shaving Using Batteries at Datacenters: Characterizing the Risks and Benefits	164
<i>Neda Nasiriani, George Kesidis, and Di Wang</i>	
An Energy Efficient Demand-Response Model for High Performance Computing Systems	175
<i>Kishwar Ahmed, Jason Liu, and Xingfu Wu</i>	

Session 6: Hardware / Design

SELF: A High Performance and Bandwidth Efficient Approach to Exploiting Die-Stacked DRAM as Part of Memory	187
<i>Yuhua Guo, Qing Liu, Weijun Xiao, Ping Huang, Norbert Podhorszki, Scott Klasky, and Xubin He</i>	
Quantifying the Potential Benefits of On-chip Near-Data Computing in Manycore Processors	198
<i>Jagadish B. Kotra, Diana Guttman, Nachiappan Chidamabaram N., Mahmut T. Kandemir, and Chita R. Das</i>	
DEMM: A Dynamic Energy-Saving Mechanism for Multicore Memories	210
<i>Akbar Sharifi, Wei Ding, Diana Guttman, Hui Zhao, Xulong Tang, Mahmut Kandemir, and Chita Das</i>	
An Analytical Model of Hardware Transactional Memory	221
<i>Daniel Castro, Paolo Romano, Diego Didona, and Willy Zwaenepoel</i>	

Session 7: Cloud / Data Centers and Applications

Modeling and Analysis of Performance Under Interference in the Cloud	232
<i>Scott Votke, Seyyed Ahmad Javadi, and Anshul Gandhi</i>	
Competition and Peak-Demand Pricing in Clouds Under Tenants' Demand Response	244
<i>George Kesidis, Uday V. Shanbhag, Neda Nasiriani, and Bhuvan Urgaonkar</i>	
A Measurement-Based Approach to Performance Prediction in NoSQL Systems	255
<i>Flora Karniavoura and Kostas Magoutis</i>	
Author Index	263