

10th International Conference on Autonomic Computing (ICAC'13)

San Jose, California, USA
26 – 28 June 2013

ISBN: 978-1-7138-1754-3

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2013) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2020)

For permission requests, please contact Usenix Association
at the address below.

Usenix Association
2560 Ninth Street, Suite 215
Berkeley, California, 94710

<https://www.usenix.org/>

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
Email: curran@proceedings.com
Web: www.proceedings.com

ICAC '13:
10th International Conference on Autonomic Computing
June 26–28, 2013
San Jose, CA

Message from the Program Co-Chairs. viii

Wednesday, June 26, 2013

Cloud Management

Application Placement and Demand Distribution in a Global Elastic Cloud: A Unified Approach 1
Hangwei Qian, *VMware, Inc.*; Michael Rabinovich, *Case Western Reserve University*

To Reserve or Not to Reserve: Optimal Online Multi-Instance Acquisition in IaaS Clouds 13
Wei Wang, Baochun Li, and Ben Liang, *University of Toronto*

Elasticity in Cloud Computing: What It Is, and What It Is Not 23
Nikolas Roman Herbst, Samuel Kounev, and Ralf Reussner, *Karlsruhe Institute of Technology*

K-Scope: Online Performance Tracking for Dynamic Cloud Applications 29
Li Zhang, Xiaoqiao Meng, Shicong Meng, and Jian Tan, *IBM T.J. Watson Research Center*

System Resource Management

Adaptive Performance-Aware Distributed Memory Caching 33
Jinho Hwang and Timothy Wood, *The George Washington University*

Exploiting Processor Heterogeneity for Interactive Services 45
Shaolei Ren, *Florida International University*; Yuxiong He, Sameh Elnikety, and Kathryn S. McKinley, *Microsoft Research*

Autonomic Management of Dynamically Partially Reconfigurable FPGA Architectures Using Discrete Control 59
Xin An and Eric Rutten, *Inria Grenoble Rhône-Alpes*; Jean-Philippe Diguët and Nicolas le Griguer, *Lab-STICC*; Abdoulaye Gamatié, *LIRMM*

FMEM: A Fine-grained Memory Estimator for MapReduce Jobs 65
Lijie Xu, *Institute of Software, Chinese Academy of Sciences, and University of Chinese Academy of Sciences*; Jie Liu and Jun Wei, *Institute of Software, Chinese Academy of Sciences*

Virtual Machine Management

AGILE: Elastic Distributed Resource Scaling for Infrastructure-as-a-Service 69
Hiep Nguyen, Zhiming Shen, and Xiaohui Gu, *North Carolina State University*; Sethuraman Subbiah, *NetApp Inc.*; John Wilkes, *Google Inc.*

PACMan: Performance Aware Virtual Machine Consolidation 83
Alan Roytman, *University of California, Los Angeles*; Aman Kansal, *Microsoft Research*; Sriram Govindan, *Microsoft Corporation*; Jie Liu and Suman Nath, *Microsoft Research*

Working Set-based Physical Memory Ballooning 95
Jui-Hao Chiang, *Stony Brook University*; Han-Lin Li and Tzi-cker Chiueh, *Industrial Technology Research Institute*

Coriolis: Scalable VM Clustering in Clouds 101
Daniel Campello and Carlos Crespo, *Florida International University*; Akshat Verma, *IBM Research-India*; Raju Rangaswami, *Florida International University*; Praveen Jayachandran, *IBM Research-India*

Thursday, June 27, 2013

MapReduce Workloads and Key-Value Stores

iShuffle: Improving Hadoop Performance with Shuffle-on-Write107
Yanfei Guo, Jia Rao, and Xiaobo Zhou, *University of Colorado, Colorado Springs*

AUTOPLACER: Scalable Self-Tuning Data Placement in Distributed Key-value Stores 119
João Paiva, Pedro Ruivo, Paolo Romano, and Luís Rodrigues, *INESC-ID Lisboa, Instituto Superior Técnico, and Universidade Técnica de Lisboa*

Adaptive Information Passing For Early State Pruning in MapReduce Data Processing Workflows133
Seokyoung Hong, Padmashree Ravindra, and Kemafor Anyanwu, *North Carolina State University*

Management of Big Data Systems Track

To Auto Scale or Not to Auto Scale145
Nathan D. Mickulicz, Priya Narasimhan, and Rajeev Gandhi, *YinzCam, Inc. and Carnegie Mellon University*

Big Data Exploration via Automated Orchestration of Analytic Workflows153
Alina Beygelzimer, Anton Riabov, Daby Sow, Deepak S. Turaga, and Octavian Udrea, *IBM T. J. Watson Research Center*

ThroughputScheduler: Learning to Schedule on Heterogeneous Hadoop Clusters159
Shekhar Gupta, Christian Fritz, Bob Price, Roger Hoover, and Johan DeKleer, *Palo Alto Research Center;*
Cees Witteveen, *Delft University of Technology*

Real-Time User-Centric Management of Time-Intensive Analytics Using Convergence of Local Functions. ...167
Vinay Deolalikar, *HP-Autonomy Research*

AutoTune: Optimizing Execution Concurrency and Resource Usage in MapReduce Workflows 175
Zhuoyao Zhang, *University of Pennsylvania;* Ludmila Cherkasova, *Hewlett-Packard Labs;* Boon Thau Loo, *University of Pennsylvania*

Self-Aware Internet of Things Track

Self-healing and Optimizing of the HIP-based M2M Overlay Network183
Amine Dhraief, *HANA Research Group, University of Manouba;* Khalil Drira, *LAAS-CNRS, University of Toulouse;* Abdelfettah Belghith, *HANA Research Group, University of Manouba*

Between Neighbors: Neighbor Discovery Analysis in EH-IoTs193
Shruti Devasenapathy, Vijay S. Rao, R. Venkatesha Prasad, and Ignas Niemegeers, *Delft University of Technology;* Abdur Rahim, *CreateNet*

Towards a Generic Architecture and Methodology for Multi-goal, Highly-distributed and Dynamic Autonomic Systems201
Sylvain Frey, *EDF R&D and Télécom ParisTech, CNRS LTCI;* Ada Diaconescu, *Télécom ParisTech, CNRS LTCI;* David Menga, *EDF R&D;* Isabelle Demeure, *Télécom ParisTech, CNRS LTCI*

Learning Deployment Trade-offs for Self-Optimization of Internet of Things Applications213
Arun kishore Ramakrishnan, Nayyab Zia Naqvi, Zubair Wadood Bhatti, Davy Preuveneers, and Yolande Berbers, *KU Leuven*

Friday, June 29, 2013

Self-Protect/Self-Healing

- Autonomic Fail-over for a Software-Defined Container Computer Network**225
Chien-Yung Lee and Yu-Wei Lee, *Industrial Technology Research Institute*; Cheng-Chun Tu, *Stony Brook University and Industrial Technology Research Institute*; Pai-Wei Wang, Yu-Cheng Wang, and Chih-Yu Lin, *Industrial Technology Research Institute*; Tzi-cker Chiueh, *Stony Brook University and Industrial Technology Research Institute*
- Fault Management in Map-Reduce through Early Detection of Anomalous Nodes**235
Selvi Kadirvel, Jeffrey Ho, and José A. B. Fortes, *University of Florida*
- Reliability and Timeliness Analysis of Fault-tolerant Distributed Publish/Subscribe Systems**247
Thadpong Pongthawornkamol and Klara Nahrstedt, *University of Illinois at Urbana-Champaign*; Guijun Wang, *Boeing Research & Technology*
- Mitigating Anonymity Challenges in Automated Testing and Debugging Systems**259
Silviu Andrica and George Candea, *École Polytechnique Fédérale de Lausanne (EPFL)*

Scheduling

- Zoolander: Efficiently Meeting Very Strict, Low-Latency SLOs**265
Christopher Stewart and Aniket Chakrabarti, *The Ohio State University*; Rean Griffith, *VMware*
- Preemptive ReduceTask Scheduling for Fair and Fast Job Completion**279
Yandong Wang, *Auburn University*; Jian Tan, *IBM T.J. Watson Research*; Weikuan Yu, *Auburn University*; Li Zhang and Xiaoqiao Meng, *IBM T.J. Watson Research*
- QoS-Aware Admission Control in Heterogeneous Datacenters**291
Christina Delimitrou, Nick Bambos, and Christos Kozyrakis, *Stanford University*
- Performance Inconsistency in Large Scale Data Processing Clusters**297
Mingyuan Xia and Nan Zhu, *McGill University*; Yuxiong He and Sameh Elnikety, *Microsoft Research Redmond*; Xue Liu, *McGill University*

Power/Temperature-Aware Management

- Temperature Aware Workload Management in Geo-distributed Datacenters**303
Hong Xu, Chen Feng, and Baochun Li, *University of Toronto*
- Power-Aware Throughput Control for Database Management Systems**315
Zichen Xu and Xiaorui Wang, *The Ohio State University*; Yi-Cheng Tu, *University of South Florida*
- Wireless Inference-based Notification (WIN) without Packet Decoding**325
Kevin Chen and H. T. Kung, *Harvard University*