

THE 7TH ACM/SPEC
INTERNATIONAL CONFERENCE
ON PERFORMANCE ENGINEERING

CONFERENCE PROGRAM

MARCH 12-16, 2016
DELFT, THE NETHERLANDS

ICPE 2016



Association for
Computing Machinery



Standard Performance
Evaluation Corporation

Program Chairs' Welcome

The 7th ACM/SPEC International Conference on Performance Engineering (ICPE 2016) takes place in Delft in The Netherlands in March 2016. The conference grew out of the ACM Workshop on Software Performance (WOSP since 1998) and the SPEC International Performance Engineering Workshop (SIPEW since 2008), with the goal of integrating theory and practice in the field of performance engineering. It is a great pleasure for us to offer an outstanding technical program this year, which we believe will allow researchers and practitioners to present their visions and latest innovation, and to exchange ideas within the community.

Overall, we received 89 high quality submissions across all three tracks. The main Research Track attracted 57 submissions with 19 accepted (33% acceptance rate) for presentation at the conference. Among them were 16 full papers and three short papers. Each paper received at least three reviews from experienced program committee members. In the Work-In-Progress and Vision Track, six out of 15 contributions were selected. The Industry and Experience Track received 17 submissions, of which seven were selected for inclusion in the program. The accepted papers were organized into five research track sessions, two industry track sessions, and one WiP and vision track session. Three best paper candidates were also selected: two research papers and one industry paper.

We are proud to have three excellent keynote speakers as part of our technical program:

- Bianca Schroeder from University of Toronto, Canada, presenting "Case studies from the real world: The importance of measurement and analysis in building better systems"
- Wilhelm Hasselbring from Kiel University, Germany, discussing "Microservices for Scalability"
- Angelo Corsaro, Chief Technology Officer at PrismTech, talking about "Cloudy, Foggy and Misty Internet of Things"

In addition, the program includes four tutorials, a doctoral symposium, a poster and demo track, the SPEC Distinguished Dissertation Award, and three interesting workshops, including the International Workshop on Large-Scale Testing (LT), the 2nd International Workshop on Performance Analysis of Big data Systems (PABS), and the 2nd Workshop on Challenges in Performance Methods for Software Development (WOSP-C).

The program covers traditional ICPE topics such as software and systems performance modeling and prediction, analysis and optimization, characterization and profiling, as well as application of performance engineering theory and techniques to several practical fields, including distributed systems, cloud computing, storage, energy, big data, virtualized systems and containers.

We'd like to thank all the authors who submitted their innovative work to ICPE this year. In addition, we thank all the program committee members and subreviewers for volunteering their time for the benefit of the community and their hard work in providing quality reviews for the submitted papers. Finally, we'd like to thank all the participants who will attend ICPE in person this year, since we will rely on you to make this event interactive, engaging, and thought-provoking for everyone involved.

We look forward to meeting all of you in Delft in March 2016!

Steffen Becker

ICPE 2016 Program Co-Chair
Chemnitz University of Technology, Germany

Xiaoyun Zhu

ICPE 2016 Program Co-Chair
Futurewei Technologies, USA

Jerry Rolia

ICPE 2016 Industrial Co-Chair
Hewlett Packard Labs, USA

Manoj Nambiar

ICPE 2016 Industrial Co-Chair
Tata Consultancy Services, India

General Chairs' Welcome

It is our great pleasure to welcome you to the 7th ACM/SPEC International Conference on Performance Engineering and to the beautiful city of Delft. If one looks at the ancient canals lined with merchants' houses, the old churches, and the splendid town hall, one sees that Delft's rich history is still very much alive. At the same time, though, it is a very modern and vibrant city. Ten percent of its 100,000 inhabitants are students and the university has attracted a large number of technology-oriented companies.

This year's International Conference on Performance Engineering (ICPE) continues its tradition of being the premier forum for the integration of theory and practice in the field of performance engineering. ICPE is an annual joint meeting that has grown out of the ACM Workshop on Software Performance (WOSP) and the SPEC International Performance Engineering Workshop (SIPEW). It brings together researchers and industry practitioners to share ideas, discuss challenges, and present results of both work-in-progress and state-of-the-art research on performance engineering of software and systems.

Putting together *ICPE'16* was a team effort. We first thank the authors for providing the content of the program. We are grateful to the PC Chairs and the program committee who worked very hard in reviewing papers and providing feedback for authors. Finally, we thank our hosting University, Delft University of Technology and our sponsors, ACM SIGMETRICS, and SPEC.

We hope that you will find this program interesting and thought-provoking and that the conference will provide you with a valuable opportunity to share ideas with other researchers and practitioners from institutions around the world and to learn about Delft's rich history.

Alberto Avritzer
ICPE 2016 General Chair
Sonatype, Inc., MD, USA

Alexandru Iosup
ICPE 2016 General Chair
Delft University of Technology,
The Netherlands

Program at a Glance

Saturday, March 12, 2016

	LT 2016 Workshop <i>Room: Senate Hall</i>
08:00 - 09:00	Registration @ Committee Room 2
09:00 - 10:35	Session 1
10:35 - 11:00	Coffee Break @ Beatrix
11:00 - 12:30	Session 2
12:30 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 15:30	Session 3
15:30 - 16:00	Coffee Break @ Beatrix
16:00 - 17:40	Session 4
	PABS 2016 Workshop <i>Room: Committee Room 4</i>
08:00 - 09:00	Registration @ Committee Room 2
09:00 - 10:20	Session 1
10:20 - 10:40	Coffee Break @ Beatrix
10:40 - 12:30	Session 2
12:30 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 15:50	Session 3
15:50 - 16:10	Coffee Break @ Beatrix
16:10 - 17:30	Session 4
	WOSP-C 2016 Workshop <i>Room: Committee Room 3</i>
08:00 - 09:00	Registration @ Committee Room 2
09:00 - 10:30	Session 1
10:30 - 11:00	Coffee Break @ Beatrix
11:00 - 12:30	Session 2
12:30 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 15:30	Session 3
15:30 - 16:00	Coffee Break @ Beatrix
16:00 - 17:30	Session 4

Program at a Glance

Sunday, March 13, 2016

	Tutorials	
08:00 - 09:00	Registration @ Committee Room 2	
09:00 – 12:30	Andre Bondi. Incorporating Software Performance Engineering Methods Practice into the Software Development Life Cycle	<i>Verena Bitto, Philipp Lengauer. Building Custom, Efficient and Accurate Memory Monitoring Tools for Java Applications</i>
12:30 - 14:00	Lunch @ Frans van Hasselt Hall	
14:00 - 17:30	Rean Griffith, Anne Holler, Xiaoyun Zhu. Application Performance Management in Virtualized Datacenters	Giuliano Casale, Simon Spinner, Weikun Wang. Automated Parameterization of Performance Models from Measurements

Morning Coffee Break @ Beatrix **10:30 – 11:00**

Afternoon Coffee Break @ Beatrix **15:30 – 16:00**

Program at a Glance

Monday, March 14, 2016

	Main Conference <i>Room: Senate Hall</i>
08:00 - 09:00	Registration @ Committee Room 2
08:45 - 09:00	Welcome Message
09:00 - 10:00	Keynote 1
10:00 - 10:30	Best Paper Candidate
10:30 - 11:00	Coffee Break @ Beatrix
11:00 - 12:30	Research Track Session 1: Performance Modeling
12:30 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 15:15	Industry Track Session 1: Modeling and Prediction
15:15 - 15:30	Research Track
15:30 - 16:00	Coffee Break @ Beatrix
16:00 - 17:00	Work in Progress Track
17:00 - 18:00	Doctoral Symposium
18:00 - 19:30	Demos and Posters Track

Tuesday, March 15, 2016

	Main Conference <i>Room: Senate Hall</i>
08:30 - 09:00	Registration @ Committee Room 2
09:00 - 10:00	Keynote 2
10:00 - 10:30	Best Paper Candidate
10:30 - 11:00	Coffee Break @ Beatrix
11:00 - 12:15	Industry Track Session 2: Analysis and Optimization
12:15 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 15:30	Research Track Session 2: Distributed Systems & Cloud
15:30 - 16:00	Coffee Break @ Beatrix
16:00 - 16:30	Best Paper Candidate
16:30 - 17:45	Research Track Session 3: Characterization and Profiling
19:30 - 22:00	Social Dinner (including Best Paper and other awards)

Program at a Glance

Wednesday, March 16, 2016

	Main Conference <i>Room: Senate Hall</i>
08:30 - 09:00	Registration @ Committee Room 2
09:00 - 10:00	Keynote 3
10:00 - 10:30	SPEC Distinguished Dissertation Award
10:30 - 11:00	Coffee Break @ Beatrix
11:00 - 12:45	Research Track Session 4: Monitoring and Analysis
12:45 - 14:00	Lunch @ Frans van Hasselt Hall
14:00 - 14:30	Research Track Session 5: Data Intensive Computing
14:30 - 15:10	SPEC Research Group Annual Meeting
15:10 - 15:30	Closing Session

International Workshop on Large-Scale Testing

Many large-scale software systems (e.g., e-commerce websites, telecommunication infrastructures and enterprise systems, etc.) must service hundreds, thousands or even millions of concurrent requests. Many field problems of these systems are due to their inability to scale to field workloads, rather than feature bugs. In addition to conventional functional testing, these systems must be tested with large volumes of concurrent requests (called the load) to ensure the quality of these systems. Large-scale testing includes all different objectives and strategies of testing large-scale software systems using load. Examples of large-scale testing are live upgrade testing, load testing, high availability testing, operational profile testing, performance testing, reliability testing, stability testing and stress testing.

Large-scale testing is a difficult task requiring a great understanding of the system under test. Practitioners face many challenges such as tooling (choosing and implementing the testing tools), environments (software and hardware setup) and time (limited time to design, test, and analyze). Yet, little research is done in the software engineering domain concerning this topic. Moreover, prior large-scale testing research have been largely focused on telecommunication applications and web-based e-commerce systems. Industry has been focused primarily on creating tools to automatically drive specified load into the system under test (e.g., LoadRunner or Apache JMeter). Large-scale testing is gaining more importance, as an increasing number of systems (on-premise and/or cloud-based systems) are designed to serve thousands or millions of users.

Program Outline

Session 1:

09:00 - 09:05	Welcome & Introduction
09:05 - 10:05	Keynote I: Automated Analysis of Load Test Results of Systems with Equilibrium or Transient Behaviour , Andrè Bondi (Software Performance and Scalability Consulting LLC)
10:05 - 10:35	Upcoming Challenges in Large Scale Performance Testing , Manoj Karunakaran Nambiar (Tata Consultancy Services)
10:35 - 11:00	Coffee Break

Session 2:

- 11:00 - 11:30 **Recent Experiences and Future Challenges Using Automatic Performance Modelling to Complement Testing,**
Paul Brebner (Performance Assurance Pty Ltd)
- 11:30 - 12:00 **How to Load Test a SOA - a Practitioners View,**
Daniel Tertilt (Qupe GmbH)
- 12:00 - 12:30 **Monitoring-Based Testing of Elastic Cloud Computing Applications,**
Michel Albonico (Universidade Tecnológica Federal do Paraná),
Jean-Marie Mottu (Université de Nantes) and Gerson Sunyé
(Université de Nantes)
- 12:30 - 14:00 Lunch Break

Session 3:

- 14:00 - 15:00 Keynote II: **Performance Testing in Software Development: Getting the Developers on Board,**
[Lubomír Bulej](#) (Charles University in Prague)
- 15:00 - 15:30 **Performance Regression Analysis: Challenges and Opportunities,**
Cor-Paul Bezemer (Queen's University)
- 15:30 - 16:00 Coffee Break

Session 4:

- 16:00 - 17:30 Joint Discussion with WOSP-C
- 17:30 - 17:40 Closing of LT 2016

PABS 2016 Workshop – Saturday, March 12

Workshop on Performance Analysis of Big Data Systems

Today, big data systems deal with volume, velocity, variety and veracity of the application data which may be deployed on dedicated single high performance systems (such as Netezza), dedicated commodity based clusters or shared architectures such as clouds. We witness an explosive growth in the complexity, diversity, number of deployments and capabilities of big data processing systems such as Map-Reduce, Hbase, Hive, Cassandra, Big Table, Hyracks, Dryad, Pregel and Mongo DB. The big data system may use new operating system designs, advanced data processing algorithms, parallelization of application, high performance computing architectures such as GPUs etc. and clusters to improve the performance. Looking at the volume of data to mine, and complex architectures, one may need to monitor, analyze, identify or predict bottlenecks to optimize the system and improve its performance.

The workshop on performance analysis of big data systems (PABS) aims at providing a platform for scientific researchers, academicians and practitioners to discuss techniques, models, benchmarks, tools and experiences while dealing with performance issues in big data systems. The primary objective is to discuss performance bottlenecks and improvements during big data analysis using different paradigms, architectures and technologies such as Map-Reduce, Hbase, MPP, Big Table, NOSQL, graph based models and any other new upcoming paradigms. We propose to use this platform as an opportunity to discuss systems, architectures, tools, and optimization algorithms that are parallel in nature and hence make use of advancements to improve the system performance. This workshop shall focus on the performance challenges imposed by big data systems and on the different state-of-the-art solutions proposed to overcome these challenges.

Program Outline

Session 1:

- | | |
|---------------|--|
| 09:00 - 09:05 | Welcome note by workshop Co-chair |
| 09:05 - 10:20 | Keynote - Challenges in Truly Scaling Services
<i>Manish Gupta, VP and Director, Xerox Research Center, India</i> |
| 10:20 - 10:40 | Coffee Break |

Session 2:

- | | |
|---------------|--|
| 10:40 - 11:20 | Paper Presentation - Towards the Prediction of the Performance and Energy Efficiency of Distributed Data Management Systems
<i>Raik Niemann, Institute for Information Systems, Germany</i> |
|---------------|--|

11:20 - 12:30 Invited Talk - Big Data Applications Performance Assurance
Boris Zibitsker, CEO BEZNext

12:30 - 14:00 Lunch Break

Session 3:

14:00 - 15:10 Invited Talk - Performance Engineering for In-Memory
Databases: Models, Experiments and Optimization
Giuliano Casale, Sr. Lecturer, Imperial College London, UK

15:10 - 15:50 Paper Presentation - A Constraint Programming Based Energy
Aware Resource Management Middleware for Clouds Processing
MapReduce Jobs with Deadlines
Adam Gregory et. al. , Carleton University, Canada

15:50 - 16:10 Coffee Break

Session 4:

16:10 - 17:20 Tutorial - Challenges for Big Data Application Performance
Tuning and Prediction
Rekha Singhal, Sr. Scientist, TCS Innovation Labs, India

17:20 - 17:30 Closing remarks by workshop Co-Chair

Workshop on Challenges in Performance Methods for Software Development

There are new challenges to product performance due to changes in software and in the development process. Faster development means less time for performance planning. The need for scalability and adaptability increases the pressure while introducing new sources of delay in the use of middleware. Model-driven engineering, component engineering and software development tools offer opportunities, but their exploitation requires effort and carries risks.

This second edition of WOSP-C will explore the challenges and seek to identify the most promising lines of attack on them, through a combination of research and experience papers, vision papers describing new initiatives and ideas, and discussion sessions. Papers describing new projects and approaches are particularly welcome. As implied by the title, the workshop focus is on methods usable anywhere across the life cycle, from requirements to design, testing and evolution of the product. The discussions will attempt to map the future of the field.

Program Outline

Session 1: Contributed Papers on Performance-aware Design

09:00 - 09:05	Introduction and Chair's Welcome
	Challenges in Applying Control Theory to Software Performance Engineering for Adaptive Systems Davide Arcelli and Vittorio Cortellessa (Universita' dell'Aquila)
09:05 - 10:30	DiffLQN: Differential Equation Analysis of Layered Queuing Networks Tabea Waizmann and Mirco Tribastone (IMT Institute for Advanced Studies Lucca)
	Simulation of techniques to improve the utilization of cloud elasticity in workload-aware adaptive software Diego Pérez Palacin, Raffaella Mirandola, and Marco Scoppetta (Politecnico di Milano)
10:30 - 11:00	Coffee Break

Session 2: Contributed Papers on Performance Measurement

- A Reference Architecture for Online Performance Model Extraction in Virtualized Environments
Simon Spinner, Jürgen Walter, and Samuel Kounev (University of Würzburg)
- Challenges with Applying Performance Testing Methods for Systems Deployed on Shared Environments with Indeterminate Competing Workloads
11:00 - 12:30 Andre Bondi (Software Performance and Scalability Consulting LLC)
- Execution Time Compensation for Cloud Applications by Subtracting Steal Time based on Host-Level Sampling
Masao Yamamoto and Kohta Nakashima (Fujitsu Laboratories Ltd.)
- Performance Mimicking Benchmarks for Multi-tier Applications
Subhasri Duttagupta, Mukund Kumar, (both Tata Consulting Services), and Varsha Apte (Indian Institute of Technology - Bombay)
- 12:30 - 14:00 Lunch Break

Session 3:

- 14:00 - 15:30 **Invited Talk by Murray Woodside: Some Software Performance Challenges in 2016**
Chair: Andre Bondi (Software Performance and Scalability Consulting LLC)
- 15:30 - 16:00 Coffee Break

Session 4:

- 16:00 - 17:30 Session 4: Break-out discussion groups (together with LT workshop)
- 17:30 - 17:40 Closing and Workshop Summary
Including 5 minute presentation per break-out group

Tutorials – Sunday, March 13

Incorporating Software Performance Engineering Methods Practice into the Software Development Life Cycle

Author: Andre Bondi (Software Performance and Scalability Consulting LLC)

Tutorial date/time: Sunday, March 13, 2016 (3 hours + break, 09:00-12:30)

Abstract: Too often, attention is only paid to performance concerns after functional testing, when it is usually too late to remedy disabling performance problems. We describe how early attention to performance concerns and early planning of performance requirements and performance testing can prevent debacles like the early rollout of healthcare.gov while enhancing reliability and scalability, and while addressing other cross-cutting concerns. We show how performance engineering methods may be integrated into all phases of the software lifecycle, from the conception of a system to requirements specification, architecture, testing, and finally to deployment. By reviewing the architecture of a system before design and implementation take place, we reduce the risk of designing and developing a system that contains inherent performance vice. Modeling can be used to justify architectural and scheduling decisions such as the use of scheduling rules. The outputs of performance tests planned with reference to performance models enable us to identify concurrent programming issues and other issues that would not be apparent in unit testing.

Building Custom, Efficient and Accurate Memory Monitoring Tools for Java Applications

Authors: Verena Bitto (Johannes Kepler University, Linz, Austria), Philipp Lengauer (Johannes Kepler University, Linz, Austria)

Tutorial date/time: Sunday, March 13, 2016 (3 hours + break, 09:00-12:30)

Abstract: Traditional monitoring techniques can distort application behavior significantly. In this tutorial, we will provide an overview about state-of-the-art monitoring techniques and their impact on memory behavior. It will show lightweight techniques that can be used to build a custom but efficient monitoring tool. We will include topics ranging from capturing interesting events, serializing and processing the data offline, dealing with large amounts of data, to visualizing it. AntTracks is a custom memory monitoring tool built into the Hotspot Java Virtual Machine. It achieves very low run-time overhead (4.68%) although capturing all allocations as well as object movements. We will thus use it exemplarily throughout the tutorial to demonstrate the discussed techniques.

Application Performance Management in Virtualized Datacenters

Authors: Rean Griffith (VMware, USA), Anne Holler (Independent Researcher) and Xiaoyun Zhu (Futurewei Technologies, Inc., USA).

Tutorial date/time: Sunday, March 13, 2016 (3 hours + break, 14:00-17:30)

Abstract: Virtualized private or public cloud datacenters provide flexible access to computing resources, but their use can present challenges to dynamically meeting application performance goals efficiently. In this tutorial, we will first examine the kinds of resource schedulers currently available for the datacenter, along with their use cases. We will then present techniques for the automatic scaling of applications horizontally and vertically to maintain their Service Level Objectives with right-sized encapsulation. And finally, we will discuss analytics pipelines for the telemetry data of workloads running in the datacenter.

Automated Parameterization of Performance Models from Measurements

Authors: Giuliano Casale (Imperial College London, UK); Simon Spinner (University of Würzburg, Germany); Weikun Wang (Imperial College London, UK)

Tutorial date/time: Sunday, March 13, 2016 (3 hours + break, 14:00-17:30)

Abstract: The goal of this tutorial is to present the problem of estimating parameters of performance models from measurements of real systems and discuss algorithms that can support researchers and practitioners in this task. The focus will be on performance models based on queueing systems, where the estimation of request arrival rates and service demands is a required input to the model. The tutorial will review existing estimation methods for service demands, ranging from regression-based methods to maximum likelihood techniques, and present models to characterize time-varying arrival processes. The tutorial will also demonstrate the use of relevant tools that automate demand estimation, such as LibRede, FG, and M3A.

Main Conference – Monday, March 14

Welcome Session

08:45 - 09:00 Welcome Message

Keynote 1

Chair: Alexandru Iosup, Delft University of Technology, The Netherlands

09:00 - 10:00 Bianca Schroeder (University of Toronto, Canada). Case studies from the real world: The importance of measurement and analysis in building better systems

Best Paper Candidate

Chair: Steffen Becker, Chemnitz University of Technology, Germany

10:00 - 10:30 Weikun Wang, Giuliano Casale, Ajay Kattapur and Manoj Nambiar. *Maximum Likelihood Estimation of Closed Queueing Network Demands from Queue Length Data*

10:30 - 11:00 Coffee Break

Research Track Session 1: Performance Modeling

Chair: Andre Bondi, Software Performance and Scalability Consulting LLC

David Maplesden, Ewan Tempero, John Hosking and John Grundy. *A Cost/Benefit Approach to Performance Analysis*

Piotr Rygielski, Viliam Simko, Felix Sittner, Doris Aschenbrenner, Samuel Kounev and Klaus Schilling. *Automated Extraction of Network Traffic Models Suitable for Performance Simulation*

11:00 - 12:30

Juan Pablo Sandoval Alcocer, Alexandre Bergel and Marco Tulio Valente. *Learning From Source Code History to Identify Performance Failures*

Alireza Khoshkbarforoushha and Rajiv Ranjan. *Resource and Performance Distribution Prediction for Large Scale Analytics Queries*

12:30 - 14:00 Lunch

Main Conference – Monday, March 14

Industry Track Session 1: Modeling and Prediction

Chair: Klaus-Dieter Lange, Hewlett Packard Enterprise, USA

Paul Brebner. *Automatic Performance Modelling from Application Performance Monitoring (APM) Data: An Experience Report*

14:00 - 15:15

Patrick Pegus li, Benoy Varghese, Tian Guo, David Irwin, Prashant Shenoy, Anirban Mahanti, James Culbert, John Goodhue and Chris Hill. *Analyzing the Efficiency of a Green University Data Center*

Qi Wang, Lucy Cherkasova, Jun Li and Haris Volos. *Interconnect Emulator for Aiding Performance Analysis of Distributed Memory Applications*

Research Track

Chair: Klaus-Dieter Lange, Hewlett Packard Enterprise, USA

15:15 - 15:30

Assaf Eisenman, Ludmila Cherkasova, Guilherme Magalhaes, Qiong Cai, Paolo Faraboschi and Sachin Katti. *Parallel Graph Processing: Prejudice and State of the Art*

15:30 – 16:00

Coffee Break

Work in Progress Track

Chair: Vittorio Cortellessa, Universita' dell'Aquila, Italy

Jürgen Walter, André van Hoorn, Heiko Koziolok, Dušan Okanović and Samuel Kounev. *Asking "What?", Automating the "How?": The Vision of Declarative Performance Engineering*

Thomas Begin and Alexandre Brandwajn. *Predicting the System Performance by Combining Calibrated Performance Models of its Components - A Preliminary Study*

16:00 - 17:00

Mathias Menninghaus and Elke Pulvermüller. *Towards Using Code Coverage Metrics for Performance Comparison on the Implementation Level*

Dheeraj Chahal, Rupinder Virk and Manoj Nambiar. *Pre-deployment Performance Estimation of IO Intensive Workloads*

Divya Gupta, Lucas Perronne and Sara Bouchenak. *BFT-Bench: A Framework to Evaluate BFT Protocols*

Main Conference – Monday, March 14

Sourav Medya, Ludmila Cherkasova, Guilherme Magalhaes, Kivanch Ozonat, Chaitra Padmanabha, Jiban Sarma and Imran Sheikh. *Towards Performance and Scalability Analysis of Distributed Memory Programs on Large-Scale Clusters*

17:00 - 18:00 Coffee Break

Doctoral Symposium

Holger Knoche. *Sustaining Runtime Performance while Incrementally Modernizing Transactional Monolithic Software towards Microservices*

17:00 - 18:00

David Georg Reichelt and Stefan Kühne. *Empirical Analysis of Performance Problems at Code Level*

Demos and Posters Track

Ji Xue, Evgenia Smirni, Robert Birke and Lydia Y. Chen. *PROST: Prediction for Resource Usages with Spatial and Temporal Dependencies*

Nick Principe and Vernon Miller. *SPEC SFS 2014: Application-Level Storage Performance Evaluation*

Christoph Heger, André van Hoorn, Dušan Okanović, Stefan Siegl and Alexander Wert. *diagnoseIT: Expert-Guided Automatic Diagnosis of Performance Problems in Enterprise Applications*

18:00 - 19:30 Vincenzo Ferme and Cesare Pautasso. *Integrating Faban with Docker for Performance Benchmarking*

David Schmidt, Lisa Roderick and Andrew Bond. *Inspiring Virtualization Innovation - SPECvirt_dc*

Ahmed Ali-Eldin, Alexey Ilyushkin, Bogdan Ghit, Nikolas Herbst, Alessandro Papadopoulos and Alexandru Iosup. *Benchmarking Auto-Scaling Algorithms*

Yasir Ibrahim and Waleed Ahmed. *New Technique Using Multiple Symmetric keys for Multilevel Encryption*

Alexandru Iosup, Samuel Kounev and Kai Sachs. *SPEC Research Group's Cloud Working Group (poster)*

Main Conference – Monday, March 14

Nikolas Roman Herbst, Andreas Weber, Samuel Kounev and Henning Groenda. *BUNGEE: An Elasticity Benchmark for Self-Adaptive IaaS Cloud Environments* (poster)

Jürgen Walter, André van Hoorn, Heiko Koziol, Dušan Okanović and Samuel Kounev. *The Vision of Declarative Performance Engineering* (poster)

Main Conference – Tuesday, March 15

Keynote 2

Chair: Murray Woodside, Carleton University, Canada

09:00 - 10:00 Wilhelm Hasselbring (Kiel University, Germany). *Microservices for Scalability*

Best Paper Candidate

Chair: Manoj Nambiar, Tata Consultancy Services, India

10:00 - 10:30 Michael Steindorfer and Jurgen Vinju. *Performance Modeling of Maximal Sharing*

10:30 - 11:00 Coffee Break

Industry Track Session 2: Analysis and Optimization

Chair: Manoj Nambiar, Tata Consultancy Services, India

Jóakim von Kistowski, Hansfried Block, John Beckett, Cloyce Spradling, Klaus-Dieter Lange and Samuel Kounev. *Variations in CPU Power Consumption*

11:00 - 12:15 Luyang Wang, Pallab Bhattacharya, Yao-Min Chen, Shrinivas Joshi and James Cheng. *End-to-End Java Security Performance Enhancements for Oracle SPARC Servers*

Amit Kalele and Anubhav Jain. *Accelerating The Optimal Trade-Off Circular Harmonic Function Filter Design on Multicore Systems*

12:15 - 14:00 Lunch

Research Track Session 2: Distributed Systems & Cloud

Chair: Giuliano Casale, Imperial College London, UK

Tatsuma Matsuki and Naoki Matsuoka. *A Resource Contention Analysis Framework for Diagnosis of Application Performance Anomalies in Consolidated Cloud Environments*

14:00 - 15:30 Supreeth Subramanya, Zain Mustafa, David Irwin and Prashant Shenoy. *Beyond Energy-Efficiency: Evaluating Green Datacenter Applications for Energy-Agility*

Zhan Qiu, Juan F. Perez and Peter Harrison. *Tackling Latency via Replication in Distributed Systems*

Main Conference – Tuesday, March 15

Mark Grechanik, Qi Luo, Denys Poshyvanyk and Adam Porter.
Enhancing Rules For Cloud Resource Provisioning Via Learned Software Performance Models

15:30 - 16:00 Coffee Break

Best Paper Candidate

Chair: Xiaoyun Zhu, Futurewei Technologies, USA

16:00 - 16:30 Niklas Carlsson. *Optimized eeeBond: Energy Efficiency with non-Proportional Router Network Interfaces*

Research Track Session 3: Characterization and Profiling

Chair: Wilhelm Hasselbring, Kiel University, Germany

Sarat Sreepathi, Ed D'Azevedo, Bobby Philip and Patrick Worley.
Communication Characterization and Optimization of Applications
Using Topology-Aware Task Mapping on Large Supercomputers

16:30 - 17:45 Manjula Peiris and James Hill. Automatically Detecting "Excessive Dynamic Memory Allocations" Software Performance Anti-Pattern

Philipp Lengauer, Verena Bitto and Hanspeter Mössenböck. Efficient and Viable Handling of Large Object Traces

Social Dinner (including Best Paper and other awards)

19:30 - 22:00 Start at the Aula at 18:15 for a touristic tour ending at the dinner location De Lindenhof (participants can sign in at the registration desk) or start at the Grote Markt at 19:00 to walk together to the dinner location (for those participants who wants to go to the hotel after the session).

Main Conference – Wednesday, March 16

Keynote 3

Chair: Manoj Nambiar, Tata Consultancy Services, India

09:00 - 10:00 Angelo Corsaro (PrismTech). *Cloudy, Foggy and Misty Internet of Things*

SPEC Distinguished Dissertation Award

10:00 - 10:30 SPEC Distinguished Dissertation Award

10:30 - 11:00 Coffee Break

Research Track Session 4: Monitoring and Analysis

Chair: Lubomír Bulej, Charles University, Czech Republic

Peter Hofer, David Gnedt, Andreas Schörghener and Hanspeter Mössenböck. *Efficient Tracing and Versatile Analysis of Lock Contention in Java Applications on the Virtual Machine Level*

Vojtěch Horký, Jaroslav Kotrč, Peter Libič and Petr Tuma. *Analysis of Overhead in Dynamic Java Performance Monitoring*

11:00 - 12:45

Jesun Firoz, Martina Barnas, Marcin Zalewski and Andrew Lumsdaine. *The value of variance*

Sara Fioravanti, Simone Mattolini, Fulvio Patara and Enrico Vicario. *Experimental performance evaluation of different data models for a reflection software architecture over NoSQL persistence layers*

12:45 - 14:00

Lunch

Research Track Session 5: Data Intensive Computing

Chair: Samuel Kounev, University of Würzburg, Germany

14:00 - 14:30 Ravjot Singh, Cor-Paul Bezemer, Weiyi Shang and Ahmed E. Hassan. *Optimizing the Performance-Related Configurations of Object-Relational Mapping Frameworks Using a Multi-Objective Genetic Algorithm*

Main Conference – Wednesday, March 16

SPEC Research Group Annual Meeting

Chair: Samuel Kounev, University of Würzburg, Germany

14:30 - 15:10

SPEC Research Group (RG) serves as a platform for collaborative research efforts in the area of quantitative system evaluation and analysis, fostering interaction between industry and academia. The group develops methodologies, techniques and tools for measurement, load testing, profiling, workload characterization, dependability and efficiency evaluation of computing systems. SPEC RG currently has four working groups working in the areas of cloud computing, big data, DEVOPS and security benchmarking. More information can be found at <http://research.spec.org>.

Closing Session

Chair: Samuel Kounev, University of Würzburg, Germany

15:10 - 15:30

ICPE 2017

The ACM/SPEC International Conference on Performance Engineering (ICPE) provides a forum for the integration of theory and practice in the field of performance engineering. ICPE is an annual joint meeting that has grown out of the ACM Workshop on Software Performance (WOSP) and the SPEC International Performance Engineering Workshop (SIPEW). It brings together researchers and industry practitioners to share ideas, discuss challenges, and present results of both work-in-progress and state-of-the-art research on performance engineering of software and systems. ICPE 2017 will be held in L'Aquila (Italy), from 22 through 26 of April.

L'Aquila is the capital city of the Abruzzi region in Italy, and it is located approximately 100 kilometers east from Rome, with which it is connected by a highway through the mountains. Laid out within medieval walls upon a hillside in the middle of a narrow valley, tall snow-capped mountains of the Gran Sasso massif flank the town. A maze of narrow streets, lined with Baroque and Renaissance buildings and churches, open onto elegant piazzas. Home of the University of L'Aquila, it is a lively college town. The Department of Computer Science and Engineering, and Mathematics (DISIM, <http://www.disim.univaq.it>) is very active in the international research context, and in particular the Software Engineering and Architecture Group includes members that have contributed to the birth and the growth of WOSP first and then ICPE conferences.

The host institution will be the University of L'Aquila. The contact person for ICPE 2017 is Vittorio Cortellessa (<http://www.di.univaq.it/cortelle/>), who will be General Co-Chair along with Walter Binder from University of Lugano (<http://www.inf.usi.ch/faculty/binder/>). Anne Koziolk from Karlsruhe Institute of Technology (http://sdq.ipd.kit.edu/people/anne_koziolk/), and Evgenia Smirni from College of William and Mary (<http://www.cs.wm.edu/~esmirni/>) will be Program Co-Chairs.

About SIGSOFT

The ACM Special Interest Group on Software Engineering (SIGSOFT) focuses on issues related to all aspects of software development and maintenance. Areas of special interest include: requirements, specification and design, software architecture, validation, verification, debugging, software safety, software processes, software management, measurement, user interfaces, configuration management, software engineering environments, and CASE tools. SIGSOFT is run by a volunteer Executive Committee composed of officers elected every three years, and assisted by a professional program director employed by the ACM.

Newsletter

Software Engineering Notes is the bi-monthly ACM SIGSOFT newsletter. For further information, see <http://www.acm.org/sigsoft/SEN/>.

About SIGMETRICS

SIGMETRICS is the ACM Special Interest Group (SIG) for the computer systems performance evaluation community.

SIGMETRICS promotes research in performance analysis techniques as well as the advanced and innovative use of known methods and tools. It sponsors conferences, such as its own annual conference (SIGMETRICS), publishes a newsletter (Performance Evaluation Review), and operates a mailing list linking researchers, students, and practitioners interested in performance evaluation.

Target areas of performance analysis include file and memory systems, database systems, computer networks, operating systems, architecture, distributed systems, fault tolerant systems, and real-time systems. In addition, members are interested in developing new performance methodologies including mathematical modeling, analysis, instrumentation techniques, model verification and validation, workload characterization, simulation, statistical analysis, stochastic modeling, experimental design, reliability analysis, optimization, and queuing theory.

For further information, please visit <http://www.sigmetrics.org/>.

About SPEC

The **Standard Performance Evaluation Corporation (SPEC)** was formed in 1988 to establish industry standards for measuring computer performance. Since then, SPEC has become the largest and most influential benchmark consortium in the world.

SPEC currently offers more than 20 industry-standard benchmarks and tools for system performance evaluation in a variety of application areas. Thousands of SPEC benchmark licenses have been issued to companies, resource centers, and educational institutions globally. Organizations using these benchmarks have published more than 30,000 peer-reviewed performance reports.

SPEC Benchmarks

SPEC offers a range of computer benchmarks and performance evaluation tools. The latest releases include SPECcapc for PTC Creo 3.0, SPEC SFS 2014,, SPECcapc for 3ds Max 2015 and SPEC Accell V1.0

Besides working on updating of many existing SPEC benchmarks and performance evaluation tools, several new projects are in development:

- Service Oriented Architecture (SOA) benchmark suite - measuring performance for typical middleware, database, and hardware deployments.
- A benchmark for measuring compute-intensive performance of handheld devices.
- SPECsip_Application benchmark suite, a system-level benchmark for application servers, HTTP, and SIP load generators.

SPEC welcomes interested conference attendees who would like to attend the co-located SPEC subcommittee meetings as guest **for free**. Please register online <URL: <https://www.regonline.com/register/login.aspx?eventID=1657216&MethodID=0&Even tsessionID=>>>. If you have any questions, e-mail info(at)spec(dot)org or contact Mr. Charles McKay or Ms. Dianne Rice on site during the conference.

ICPE 2016 Sponsors & Supporters



ACM SIGMETRICS
special interest group on performance evaluation

