



# SPEC Research<sup>SM</sup> Group Newsletter

## CONTENTS OF THIS ISSUE

- 2 SPEC Research Group Officers
- 2 SPEC Research Working Groups
- 2 Foreword from Chairs and Editors
- 3 SPEC Research Group Mission Statement
- 3 SPEC Announcements
- 4 Winner of the SPEC Distinguished Dissertation Award 2013
- 4 SPEC Distinguished Dissertation: Dynamic Server Provisioning for Data Center Power Management
- 5 SPEC RG Cloud Working Group
- 5 SPEC RG IDS Benchmarking Working Group
- 6 SPA—A Tool for Analyzing I/O Performance in Virtualized Environments
- 7 Preliminary Announcement of the 6th ACM<sup>1</sup>/SPEC International Conference on Performance Engineering ICPE 2015

## SPEC DISTINGUISHED DISSERTATION: DYNAMIC SERVER PROVISIONING FOR DATA CENTER POWER MANAGEMENT

We proudly present an abstract of this year's edition of the SPEC Distinguished Dissertation Award. Dr. Anshul Gandhi from Carnegie Mellon University, (now at IBM T. J. Watson Research Center) was awarded for his work on "Dynamic Server Provisioning for Data Center Power Management".

*Read more on page 4*

## REPORT OF THE CLOUD WORKING GROUP

The first working group of SPEC Research continues its work on Cloud Benchmarking. Currently, among others, the Cloud Working Group develops extensions to the CUP formalism for describing usage patterns for Cloud environments. Read further and get involved.

*Read more on page 5*

## REPORT OF THE IDS BENCHMARKING WORKING GROUP

The youngest working group of SPEC Research presented its first results and tools. The recently presented tool is called HInjector. HInjector is a framework for injecting hypercall attacks during regular operation of a paravirtualized guest virtual machine in a Xen-based virtualized environment.

*Read more on page 5*

## NEXT ICPE (2015) IN AUSTIN, TEXAS, USA

Interested in submitting a paper to the next edition of the ICPE? We present a preliminary announcement of the the 6th ACM<sup>1</sup>/SPEC International Conference on Performance Engineering.

*Read more on page 7*



## CONTACT

Standard Performance Evaluation Corporation (SPEC)  
7001 Heritage Village Plaza, Suite 225  
Gainesville, VA 20155, USA

### SPEC Research Group

Chair: Samuel Kounev (rgchair@spec.org)  
Web: <http://research.spec.org>

---

1 ACM approval pending

---

## SPEC RESEARCH GROUP OFFICERS

### Chair:

**Samuel Kounev**, Karlsruhe Institute of Technology (KIT), Germany, rgchair@spec.org

### Vice-chair:

**Kai Sachs**, SAP AG, Germany

### Secretary:

**Klaus-Dieter Lange**, HP, USA

### Release Manager:

**Qais Noorshams**, KIT, Germany

### Steering committee:

**J. Nelson Amaral**, University of Alberta, Canada

**Alexandru Iosup**, TU Delft, Netherlands

**Lizy John**, University of Texas at Austin, USA

**Samuel Kounev**, KIT, Germany

**Klaus-Dieter Lange**, HP, USA

**Matthias Müller**, RWTH Aachen, Germany

**Meikel Poess**, Oracle Corporation, USA

**Kai Sachs**, SAP AG, Germany

**Seetharami R. Seelam**, IBM TJ Watson Research Center, USA

**Petr Tůma**, Charles University of Prague, Czech Republic

### Publicity Officer:

**André van Hoorn**, University of Stuttgart, Germany

### Newsletter Editor:

**Piotr Rygielski**, KIT, Germany

## SPEC RESEARCH WORKING GROUPS

### Cloud Working Group

#### Chair:

**Alexandru Iosup**, TU Delft, Netherlands

#### Vice-Chair:

**Erich Nahum**, IBM TJ Watson Research Center, USA

#### Secretary:

**Aleksandar Milenkoski**, KIT, Germany

#### Release Manager:

**Nikolas Herbst**, KIT, Germany

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

### IDS Benchmarking Working Group

#### Chair:

**Marco Vieira**, University of Coimbra, Portugal

#### Vice-Chair:

**Alberto Avritzer**, Siemens Corporate Research, USA

#### Secretary:

**Aleksandar Milenkoski**, KIT, Germany

#### Release Manager:

**Aleksandar Milenkoski**, KIT, Germany

<http://research.spec.org/working-groups/ids-benchmarking-working-group.html>

## WELCOME TO THE SPEC RESEARCH GROUP NEWSLETTER

We are delighted to present to you the 3rd issue of the SPEC Research Group Newsletter. This regular publication provides information on latest developments, news and announcements relevant to the benchmarking and quantitative system evaluation communities. Our newsletter is part of our mission to foster the exchange of knowledge and experiences between industry and academia in the field of quantitative system evaluation and analysis.

The past eight months have been very intense and exciting for the SPEC RG. To the major activities and milestones reached, we include: presenting the SPEC Distinguished Dissertation Award to Dr. Anshul Gandhi; further expansion of the CUP (Cloud Usage Patterns) formalism introduced by the Cloud Working Group; The HInjector framework designed and implemented by the IDS Working Group; publishing the SPA tool (Storage Performance Analysis).

We have been actively working on preparation, planning and organization of the ICPE 2014 and ICPE 2015 conferences. We hope that the vivid exchange of ideas during the upcoming ICPE 2014 will be a great motivation for the next year of scientific and engineering work.

This issue of the Newsletter is a short one but we hope that you will enjoy reading it. We welcome and encourage your contributions for articles and suggestions for future coverage.

Samuel Kounev (KIT), Kai Sachs (SAP AG), Piotr Rygielski (KIT)

SPEC, the SPEC logo and the names SPECviewperf, SPECcapc, SPECwpc, SPECvirt\_sc, SERT, and SPEC PTDaemon are trademarks of the Standard Performance Evaluation Corporation. The SPEC Research Logo and the name SPEC Research are service marks of SPEC. Additional company, product and service names mentioned herein may be the trademarks or service marks of their respective owners. Copyright © 1988-2014 Standard Performance Evaluation Corporation (SPEC). Reprinted with permission. All rights reserved.

---

## **SPEC RESEARCH GROUP MISSION STATEMENT**

The SPEC Research Group (RG) is one of the four groups of the Standard Performance Evaluation Corporation (SPEC). Its mission is to promote innovative research in the area of quantitative system evaluation and analysis by serving as a platform for collaborative research efforts fostering the interaction between industry and academia in the field.

The scope of the group includes computer benchmarking, performance evaluation, and experimental system analysis considering both classical performance metrics such as response time, throughput, scalability and efficiency, as well as other non-functional system properties included under the term dependability, e.g., availability, reliability, and security. The conducted research efforts span the design of metrics for system evaluation as well as the development of methodologies, techniques and tools for measurement, load testing, profiling, workload characterization, dependability and efficiency evaluation of computing systems.

Samuel Kounev (KIT)

## **SPEC ANNOUNCEMENTS**

### **SPECviewperf13 Released**

*December 18, 2013*

SPEC's Graphics Performance Characterization group (SPECgpc) has released SPECviewperf 12, an all-new version of the worldwide standard for measuring graphics performance on systems running under the OpenGL and DirectX application programming interfaces.

[www.spec.org](http://www.spec.org)

### **PTDamenon Updated to Version 1.6.2**

*December 16, 2013*

An updated release of PTDaemon, v1.6.2 is now available to all licensees of SPECpower\_ssj2008 and SERT. This release adds support for the Newtons4th PPA5x0. All SPEC benchmark and tool submissions that are made using PTDaemon will require V1.6.2 from June 10th, 2014 onwards.

[www.spec.org](http://www.spec.org)

### **SPEC Releases Workstation Performance Benchmark**

*November 14, 2013*

SPEC's Workstation Performance Characterization group has released SPECwpc V1.0, designed to mea-

sure all key aspects of workstation performance based on diverse professional applications. A first set of results is available, and the benchmark is available for download in versions for both commercial and non-commercial use.

<http://spec.org/gwpg/downloadindex.html>

[www.spec.org](http://www.spec.org)

### **SERT v1.0.2 Released**

*September 19, 2013*

SERT V1.0.2 is now available. This is a point release which includes several enhancements, including support for 32-bit ARM processors, updated documentation and general bugfixes. Current licensees are entitled to a free update.

<http://spec.org/sert/>

[www.spec.org](http://www.spec.org)

### **SPECvirt\_sc2010 Goes on Retirement**

*August 8, 2013*

SPECvirt\_sc2010 will be retired on February 26, 2014 in favor of its successor, SPECvirt\_sc2013. New result submissions will be accepted through February 26, 2014. After that time, no further submissions will be accepted and support will be terminated. As SPECvirt\_sc2010 results must be reviewed by SPEC before publication, independent publication of results post-benchmark retirement is not allowed.

[http://spec.org/virt\\_sc2010/](http://spec.org/virt_sc2010/)

[www.spec.org](http://www.spec.org)

### **SPECapc Benchmark for Siemens NX 8.5 Updated**

*July 24, 2013*

SPEC's Application Performance Characterization project group has released an updated SPECapc benchmark for Siemens NX 8.5. New features in SPECapc for NX 8.5 include quality control improvements, faster run-times and automatic collection of system information for result submissions for publication on the SPEC website.

[http://spec.org/gwpg/publish/nx8.5\\_rel.html](http://spec.org/gwpg/publish/nx8.5_rel.html)

[www.spec.org](http://www.spec.org)

---

## WINNER OF THE SPEC DISTINGUISHED DISSERTATION AWARD 2013

December 23, 2013

The Research Group of the Standard Performance Evaluation Corporation (SPEC) selected the annual research prize to be awarded to a Ph.D. student whose thesis is regarded to be an exceptional, innovative contribution in the scope of the SPEC Research Group. The committee had—again—a demanding task to select the winner among the high-quality nominations. The criteria for the selection are the overall contribution in terms of scientific originality, practical relevance, impact, and quality of the presentation.

After an in-depth process the committee selected the following nomination: Anshul Gandhi was nominated by Professor Mor Harchol-Balter from Carnegie Mellon University with the topic “Dynamic Server Provisioning for Data Center Power Management.”

The winner will receive \$1000, which will be awarded at the ICPE 2014 International Conference in Dublin, Ireland. Selection committee for 2013:

- Chair: Matthias Müller, RWTH Aachen University, Germany
- Walter Binder, University of Lugano, Switzerland
- Edgar Gabriel, University Houston, USA
- Samuel Kounev, Karlsruhe Institute of Technology (KIT), Germany
- John Murphy, University College Dublin, Ireland
- Marco Vieira, University of Coimbra, Portugal
- Gerhard Wellein, Friedrich-Alexander-University Erlangen-Nürnberg

The SPEC Distinguished Dissertation Award was established in 2011 to recognize outstanding dissertations within the scope of the SPEC Research Group in terms of scientific originality, scientific significance, practical relevance, impact, and quality of the presentation. The scope of SPEC’s Research Group includes computer benchmarking, performance evaluation, and experimental system analysis in general, considering both classical performance metrics such as response time, throughput, scalability and efficiency, as well as other non-functional system properties included under the term dependability, e.g., availability, reliability, and security.

Contributions of interest span the design of metrics for system evaluation as well as the development of methodologies, techniques and tools for measurement, load testing, profiling, workload characterization, dependability and efficiency evaluation of computing systems.

Matthias Müller (RWTH Aachen University)

## SPEC DISTINGUISHED DISSERTATION: DYNAMIC SERVER PROVISIONING FOR DATA CENTER POWER MANAGEMENT

January 29, 2014

Data centers play an important role in today’s IT infrastructure. However, their enormous power consumption makes them very expensive to operate. Sadly, much of the power used by data centers is wasted because of poor capacity management, leading to low server utilization.

In order to reduce data center power consumption, researchers have proposed several dynamic server provisioning approaches. However, there are many challenges that hinder the successful deployment of dynamic server provisioning, including: (i) unpredictability in workload demand, (ii) switching costs when setting up new servers, and (iii) unavailability of data when provisioning stateful servers. Most of the existing research in dynamic server provisioning has ignored, or carefully sidestepped, these important challenges at the expense of reduced benefits. In order to realize the full potential of dynamic server provisioning, we must overcome these associated challenges.

This thesis provides new research contributions that explicitly address the open challenges in dynamic server provisioning. We first develop novel performance modeling tools [1,7,8] to estimate the effect of these challenges on response time and power. In doing so, we also address several long-standing open questions in queueing theory, such as the analysis of multi-server systems with switching costs [1,8]. We then present practical dynamic provisioning solutions [2-6,9] for multi-tier data centers, including novel solutions that allow scaling the stateful caching tier [4], and solutions that are robust to load spikes [2,3]. Our implementation results using realistic workloads and request traces on a 38-server multi-tier testbed demonstrate that dynamic server provisioning can successfully meet typical response time guarantees while significantly lowering power consumption.

While this thesis focuses on server provisioning for reducing power in data centers, the ideas presented herein can also be applied to: (i) private clouds, where unneeded servers can be repurposed for “valley-filling” via batch jobs, to increase server utilization, (ii) community clouds, where unneeded servers can be given away to other groups, to increase the total throughput, and (iii) public clouds, where unneeded virtual machines can be released back to the cloud, to reduce rental costs.

---

[1] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, and Alan Scheller-Wolf. "Exact Analysis of the M/M/k/setup Class of Markov Chains via Recursive Renewal Reward". In Proceedings of Sigmetrics 2013, pages 153-166.

[2] Anshul Gandhi, Mor Harchol-Balter, Ram Raghunathan, and Michael Kozuch. "AutoScale: Dynamic, Robust Capacity Management for Multi-Tier Data Centers". In Transactions on Computer Systems, Volume 30, Issue 4, Article 14.

[3] Anshul Gandhi, Timothy Zhu, Mor Harchol-Balter and Michael Kozuch. "SOFTScale: Stealing Opportunistically For Transient Scaling". In Proceedings of Middleware 2012, pages 142-163.

[4] Timothy Zhu, Anshul Gandhi, Mor Harchol-Balter and Michael Kozuch. "Saving Cash by Using Less Cache". In Proceedings of HotCloud 2012.

[5] Anshul Gandhi, Mor Harchol-Balter and Michael Kozuch. "Are sleep states effective in data centers?". In Proceedings of IGCC 2012, pages 113-122.

[6] Anshul Gandhi, Yuan Chen, Daniel Gmach, Martin Arlitt and Manish Marwah. "Minimizing Data Center SLA Violations and Power Consumption via Hybrid Resource Provisioning" (Best Paper Award). In Proceedings of IGCC 2011, pages 49-56.

[7] Anshul Gandhi, Varun Gupta, Mor Harchol-Balter and Michael Kozuch. "Optimality Analysis of Energy-Performance Trade-off for Server Farm Management". In Performance Evaluation, Volume 67, Issue 11, pages 1155-1171.

[8] Anshul Gandhi, Mor Harchol-Balter and Ivo Adan. "Server Farms with Setup Costs". In Performance Evaluation, Volume 67, Issue 11, pages 1123-1138.

[9] Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, and Charles Lefurgy. "Optimal Power Allocation in Server Farms". In Proceedings of Sigmetrics 2009, pages 157-168.

Anshul Gandhi, (Carnegie Mellon University, IBM T. J. Watson Research Center)

## CLoud WORKING GROUP

February 7, 2014

The SPEC RG Cloud Working Group is pleased to report a continuation, over 2013, of our commitment to cloud benchmarking, quantitative evaluation, and experimental analysis, with relevance for both academia and industry. In 20 online meetings, we have gathered guests from 18 organizations, and advanced work on four major activities. The group welcomes the additions of Salesforce, Google Inc., and MITRE to the list of active contributors. We have maintained overhead of under 10% of our meeting time, effectively maximizing our contribution to Cloud Working Group's collective mission to facilitate and foster leading research in cloud computing through experience sharing, constructive feedback, agreement on definitions, specification of metrics, and dissemination of novel methods, techniques and tools for quantitative evaluation and analysis of cloud-based systems. For 2014, the group actively seeks new participants and activities.

Experience sharing was a key activity for the Cloud Working Group in 2013 and will continue to be for the

upcoming years. Through online and face-to-face meetings, we have exchanged best-practices and collaborated on cloud benchmarking methodologies, performance evaluation tools and benchmarking suites, performance metrics, and practical results from the performance evaluation of various commercial cloud services. We have hosted representatives of the KIT (DE) and TU Delft (NL) universities; SAP (DE), Salesforce (US/SE), and IBM (US).

Over the course of 2013, we have renewed our effort towards a formalism for cloud usage patterns, which will help experts and novices, both from the industry and from the academia, to share information about IaaS, PaaS, and SaaS usage. Other main activities include work on cloud surveying existing cloud benchmarking metrics and tools, benchmarking PaaS for graph processing, and benchmarking PaaS for time-based analytics. The group welcomes new participants in these and other activities, including but not limited to benchmarking virtualized data centers (including interference), benchmarking elastic environments, benchmarking elastic or hierarchical storage, benchmarking interactive or collaborative distributed networked environments, and benchmarking online social networking platforms.

Continuing a successful 2012, in 2013 the Cloud Working Group has been represented by its members at many international academic and industrial venues. Internally within SPEC, our group has maintained contact with the OSG's implementation efforts, and contributed to the review and subsequent publication of several benchmarking tools.

To conclude, 2013 was a successful year for the Cloud Working Group. We are eager to continue pursuing our mission and strengthen our contribution in 2014 and the years to come.

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

Alexandru Iosup (TU Delft)

## IDS BENCHMARKING WORKING GROUP

January 29, 2014

The SPEC RG IDS Benchmarking Working Group successfully concluded its agenda for 2013 and faces 2014 with a renewed commitment. In 2013, the Working Group further established its research agenda having several works published or currently under consideration for publishing.

The SPEC RG IDS Benchmarking Working Group is currently researching injection of hypercall attacks for live testing of IDSeS designed for operation in virtualized environments. Hypercalls are software traps from

---

a kernel of a paravirtualized guest Virtual Machine (VM) to the underlying Virtual Machine Monitor (VMM). The hypercall interface is a relevant attack vector that a malicious VM user may use to conduct attacks of high severity against the VMM. The Working Group is designing and developing HInjector, a framework for injecting hypercall attacks during regular operation of a paravirtualized guest VM in a Xen-based virtualized environment. The on-going work on hypercall attack injection was presented to the security research community at ACSAC 2013 [1].

As part of the work on hypercall attack injection, the Working Group is currently conducting a field study on vulnerabilities of hypercalls by analyzing and characterizing such vulnerabilities. Due to the lack of public technical information on hypercall vulnerabilities, the analysis is performed by reverse-engineering patches of such vulnerabilities and implementing a proof-of-concept code. The analysis of hypercall vulnerabilities enables the configuration of HInjector so that it injects representative hypercall attacks. In addition, it enables the injection of representative vulnerabilities in the source code of VMMs, which may be used for the generation of attacks targeting the injected vulnerabilities for the purpose of evaluating IDSes.

As future work, the Working Group plans on defining metrics that take into consideration on-demand resource provisioning in virtualized environments (e.g., CPU and memory hotplugging, memory ballooning).

The Working Group is expanding its membership body with new members, as shown with the recent addition of SESM s.c.a.r.l.

[1] Aleksandar Milenkoski, Bryan D. Payne, Nuno Antunes, Marco Vieira, and Samuel Kounev. "HInjector: Injecting Hypercall Attacks for Evaluating VMI-based Intrusion Detection Systems (poster paper)." In the 2013 Annual Computer Security Applications Conference (ACSAC 2013), New Orleans, Louisiana, USA, 2013. Applied Computer Security Associates (ACSA), Maryland, USA.

<http://research.spec.org/working-groups/ids-benchmarking-working-group.html>

Aleksandar Milenkoski (KIT), Marco Vieira (University of Coimbra), Alberto Avritzer (Siemens Corporate Research, USA)

## SPA—A TOOL FOR ANALYZING I/O PERFORMANCE IN VIRTUALIZED ENVIRONMENTS

January 29, 2014

The Storage Performance Analyzer (SPA) [1] is a tool for analyzing the I/O performance in virtualized environments and can be used for performance measuring, monitoring, and modeling [2, 3, 4]. The SPA framework basically consists of a benchmark harness that

coordinates and controls the execution of benchmarks and monitors as well as a tailored analysis library used to process and evaluate the collected data.

**Measuring:** Using integrated I/O benchmarks, SPA coordinates the execution of benchmark runs on possibly multiple targets (e.g., on co-located virtual machines) to obtain measurements of the I/O performance. Currently, we have integrated two benchmarks into our framework, but further benchmarks can be integrated as required. After the measurement setup has been configured, the SPA framework first configures the benchmarks, then it executes the target workload, and it finally collects the results.

**Monitoring:** During the measurement process, the system environment as well as specific targets can be monitored to observe the I/O performance behavior. SPA can activate operating system monitors, such as blktrace and iostat, as well as self-defined monitors, e.g., logging the amount of files the benchmarks produce. The number of executed monitors is not limited and more monitors can be included if needed. If any monitors are activated, SPA starts all monitors before the benchmarks are started. After all benchmarks are finished, SPA stops all monitors and collects the data.

**Modeling:** The measurement and monitoring results can be processed and analyzed for a variety of purposes, e.g., to identify performance bottlenecks and performance interference effects. The results can also be used for performance modeling. SPA includes analysis libraries enabling fully-automated tuning and modeling using statistical regression techniques.

While the approach is tailored and pre-packaged for benchmarking and monitoring the I/O performance of storage systems in both native and virtualized environments, SPA is not limited to a specific domain and can be extended to integrate other benchmarks and monitoring tools. SPA is open-source and available for common operating systems.

[1] The Storage Performance Analyzer (SPA). <http://sdqweb.ipd.kit.edu/wiki/SPA>.

[2] Qais Noorshams, Samuel Kounev, and Ralf Reussner. Experimental Evaluation of the Performance-Influencing Factors of Virtualized Storage Systems. In Proceedings of the 9th European Performance Engineering Workshop (EPEW '12).

[3] Qais Noorshams, Dominik Bruhn, Samuel Kounev, and Ralf Reussner. Predictive Performance Modeling of Virtualized Storage Systems using Optimized Statistical Regression Techniques. In Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering (ICPE '13).

[4] Qais Noorshams, Kiana Rostami, Samuel Kounev, Petr Tuma, and Ralf Reussner. I/O Performance Modeling of Virtualized Storage Systems. In Proceedings of the IEEE 21st International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '13).

<http://research.spec.org/en/tools.html>

Qais Noorshams (KIT), Samuel Kounev (KIT)

---

# PRELIMINARY ANNOUNCEMENT OF THE 6TH ACM<sup>1</sup>/SPEC INTERNATIONAL CONFERENCE ON PERFORMANCE ENGINEERING ICPE 2015, AUSTIN, TEXAS, USA

February 26, 2014

The 6th International Conference on Performance Engineering, a Joint Meeting of WOSP/SIPEW sponsored by ACM SIGMETRICS and ACM SIGSOFT in cooperation with SPEC will be held in Austin, Texas, USA, during January 31–Feb 4, 2015. The goal of the International Conference on Performance Engineering (ICPE) is to integrate theory and practice in the field of performance engineering by providing a forum for sharing ideas and experiences between industry and academia. ICPE is a joint meeting of the ACM Workshop on Software and Performance (WOSP) and the SPEC International Performance Evaluation Workshop (SIPEW). The conference brings together researchers and industry practitioners to share and present their experience, to discuss challenges, and to report state-of-the-art and in-progress research on performance engineering.

## Topics of Interest

Topics of interest include, but are not limited to, the following: Software and Performance; Performance and software development processes; Model-driven performance engineering; Performance modeling and prediction; Performance measurement and experimental analysis; Benchmarking, configuration, sizing and capacity planning; System management and optimization; Performance in cloud, virtualized and multi-core systems; Performance and power; Performance modeling and evaluation in emerging domains.

## Important Submission Dates

|                                    |                                               |
|------------------------------------|-----------------------------------------------|
| Research papers                    | 13 Jul 2014 (abstract)<br>20 Jul 2014 (paper) |
| Industrial / Experience papers     | 31 Aug 2014 (abstract)<br>7 Sep 2014 (paper)  |
| Tutorial Proposals                 | 14 Oct 2014                                   |
| Poster and Demo papers             | 14 Oct 2014                                   |
| Work-in-progress and Vision papers | 16 Nov 2014                                   |

## Organizing Committee

### General Chairs

Lizy Kurian John, The University of Texas at Austin, USA  
Connie U. Smith, L&S Computer Technology, Inc. USA

### Program Chairs

Kai Sachs, SAP AG, Germany  
Catalina M. Lladó University of Balearic Islands, Spain

### Industrial Chair

Herb Schwetman, Oracle Labs, USA

### Tutorial Chair

Catia Trubiani, Univeristy of L'Aquila, Italy

### Demos and Posters Chairs

William Knottenbelt, Imperial College, UK  
Yaomin Chen, Oracle, USA

### Publication Chair

Simona Bernardi, C. U. De la Defensa

### Finance Chair

Anoush Najarian, Mathworks, USA

### Publicity Chairs

Amy Spellmann, The 451 Group, USA  
André van Hoorn, University of Stuttgart, Germany

### Registration Chair

Ram Krishnan, University of Texas at San Antonio

### Award Chairs

Andre B. Bondi, Siemens Corporation, USA  
Paul V. Gratz, Texas A&M University, USA

### Local Organization Chair

Nasr Ullah, Samsung, USA

### Web Chair

André van Hoorn, University of Stuttgart, Germany

## Submission Guidelines

Authors are invited to submit original, unpublished papers that are not being simultaneously considered in another forum. A variety of contribution styles for papers are solicited including: basic and applied research, industrial and experience reports, and work-in-progress/vision papers. Different acceptance criteria apply for each category; please visit our web page for details. At least one author of each accepted paper is required to register at the full rate, attend the conference and present the paper. Presented papers will be published in the ICPE 2015 conference proceedings that will be published by ACM and included in the ACM Digital Library.

<http://icpe2015.ipd.kit.edu/>

Lizy Kurian John (The University of Texas at Austin, USA)

---

<sup>1</sup> ACM approval pending