Volume 1, Issue 1
September 2012

SPEC Research℠ Group Newsletter

CONTENTS OF THIS ISSUE

REPOSITORY OF PEER-REVIEWED TOOLS

A repository of peer-reviewed tools for quantitative system evaluation and analysis is now available at the SPEC RG web site. Among the first tools to pass the review process are Faban and Kieker.

Read more on page 4

NEW SPEC RESEARCH WORKING GROUP

SPEC Research announces new working group: “Benchmarking architectures for intrusion detection in virtualized environments“. The goal of the working group is to foster and facilitate innovative research through exchange of ideas and experiences. On a long-term basis, the working group aims to develop a customizable and flexible IDS benchmarking framework which employs representative cloud workloads, accurate metrics and rigorous benchmarking methodology.

Read more on page 5

CALL FOR NOMINATIONS – SPEC DISTINGUISHED DISSERTATION AWARD

The SPEC Distinguished Dissertation Award (formally called SPEC Benchmarking Research PhD Award) aims to recognize outstanding doctoral 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 main prize – $1000 – will be awarded at the ICPE 2013 in Prague.

Read more on page 8

THE ‘BEST PAPER’ AWARDS FROM THE ICPE 2012 CONFERENCE

At the 2012 International Conference on Performance Engineering (ICPE 2012) in Boston, USA, three outstanding papers were distinguished by the Program and Awards Chairs. The awards were presented in three different categories: Best Innovation, Industry-Related, and Student Paper Award.

Read more on page 10

CONTACT

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

SPEC Research Group
Chair: Samuel Kounev (kounev@kit.edu)
Web: http://research.spec.org

SPEC RESEARCH GROUP OFFICERS

Chair:
Samuel Kounev, Karlsruhe Institute of Technology (KIT), Germany, kounev@kit.edu
Vice-chair:
Kai Sachs, SAP AG, Germany
Secretary:
Seetharami R. Seelam, IBM TJ Watson Research Center, USA
Release Manager:
Qais Noorshams, KIT, Germany
Steering committee:
J. Nelson Amaral, University of Alberta, Canada
Rema Hariharan, AMD, USA
Wilhelm Hasselbring, University of Kiel, Germany
Alexandru Iosup, TU Delft, Netherlands
Lizzy John, University of Texas at Austin, USA
Samuel Kounev, KIT, Germany
Klaus Lange, HP, USA
Meikel Poess, Oracle Corporation, USA
Kai Sachs, SAP AG, Germany
Seetharami R. Seelam, IBM TJ Watson Research Center, USA
Pert Tůma, Charles University of Prague, Czech Republic
Newsletter Editor:
Piotr Rygielski, KIT, Germany

SPEC RESEARCH WORKING GROUPS

Cloud Working Group

Chair:
Erich Nahum, IBM TJ Watson Research Center, USA
Vice-Chair:
Alexandru Iosup, TU Delft, Netherlands
Secretary:
Aleksandar Milenkoski, KIT, Germany
Release Manager:
Michal Papež, KIT, Germany

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

Working Group: Benchmarking architectures for intrusion detection in virtualized environments

Members:
Alberto Avritzer, Siemens Corporate Research, USA
Marco Vieira, University of Coimbra, Portugal
Aleksandar Milenkoski, KIT, Germany
Samuel Kounev, KIT, Germany

At the time of publishing, the officers have been not elected yet.

WELCOME TO THE SPEC RESEARCH GROUP NEWSLETTER

We are proud to launch the first edition of the SPEC Research Group Newsletter. This regular publication will provide you with 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.

This first issue includes among others a report from the 2nd Research Group (RG) Face-to-Face Meeting held on April 26th in Boston, a report on the 3nd ACM/SPEC International Conference on Performance Engineering (ICPE 2012) co-sponsored by the SPEC RG, a report on the current activities and latest developments in the RG Cloud working group as well as in the newly formed working group on intrusion detection, abstracts of selected papers from ICPE 2012, extended abstracts of the first two dissertations to win the SPEC Distinguished Dissertation Award, and finally brief descriptions of the first tools published in our new peer-reviewed tool repository.

We hope you will enjoy reading this issue and welcome and encourage your contributions for articles and suggestions for future coverage.

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

SPEC RESEARCH GROUP MISSION STATEMENT

The SPEC Research Group (RG) is an approved group of the Standard Performance Evaluation Corporation (SPEC) with the mission 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

Server Efficiency Rating Tool (SERT)

August 21, 2012

SERT is being developed for the US Environmental Protection Agency (EPA) to be part of the next generation of ENERGY STAR for Servers qualification process.

The SERT has been developed by the SPECpower committee, comprising representatives of many of key server and software companies, as well as veteran SPEC benchmark developers and supporting contributors. The SERT has drawn from experience gained in developing and supporting the SPECpower_ssj2008 benchmark.

Unlike most SPEC products the SERT is specifically designed not to be a comparative benchmark, in that it does not provide a composite power/performance score. Instead it is composed of multiple workloads, optimized to exercise CPU, memory and storage I/O subsystems. It produces detailed information relating to each subsystem's respective power and performance characteristics across multiple load levels, ranging from idle to 100% usage.

As well as working with the EPA SPEC is engaging with other international regulatory bodies to encourage adoption of the SERT. Creating an internationally available test platform is aimed at reducing the testing burden on the server industry, and enabling meaningful and consistent comparisons between products sold across the world. For more information about SERT, please see http://www.spec.org/sert/

Klaus Lange (SPECpower Committee Chair)

How Did This Get Published? Pitfalls in Experimental Evaluation of Computing Systems

June 21, 2012

Nelson J. Amaral gave a talk during the LCTES 2012 (Languages, Compilers, Tools and Theory for Embedded Systems) Special Session: Benchmarking and Performance Evaluation. The LCTES conference was held in conjunction with PLDI 2012 (Programming Language Design and Implementation) in Beijing, June 12–13. The talk was very well attended and received by many senior people in the PLDI community. The slides from the talk are available online to download.

http://webdocs.cs.ualberta.ca/~amaral/Amaral-LCTES2012.pptx

Nelson J. Amaral (University of Alberta), Piotr Rygielski (KIT)

New OSG Cloud Group in SPEC

June 16, 2012

The Standard Performance Evaluation Corp. (SPEC) has formed a new group to research and recommend application workloads for benchmarking cloud computing performance.

The group functions under SPEC's Open Systems Group (OSG) and is working in cooperation with other SPEC committees and subcommittees to define cloud benchmark methodologies, determine and recommend application workloads, identify cloud metrics for existing SPEC benchmarks, and develop new cloud benchmarks. Current participants in OSGCloud include AMD, Dell, IBM, Ideas International, Intel, Karlsruhe Institute of Technology, Oracle, Red Hat and VMware. Long-time SPEC benchmark developer Yun Chao is a supporting contributor. The group collaborates with the SPEC Research Cloud group, which is working on gaining a broader understanding of cloud behavior and performance issues.

While cloud performance takes into account many of the same characteristics as current SPEC benchmarks – throughput, response time and power, for example – it also brings into play new metrics such as elasticity, defined as how quickly a service can adapt to changing customer needs.

"Cloud computing is on the rise and represents a major shift in how servers are used and how their performance is measured," says Rema Hariharan, chair of OSGCloud. "We want to assemble the best minds to define this space, create workloads, augment existing SPEC benchmarks, and develop new cloud-based benchmarks."

The group targets three types of users for the workloads and benchmarks it will create: Hardware and software vendors providing products that enable cloud services. Cloud providers – entities that offer cloud services, such as IaaS, PaaS or SaaS. Business customers who would use benchmark results to help them select cloud providers.

OSGCloud has already developed a 50-page report that details its objectives, benchmark considerations, characteristics of a cloud benchmark, and tools for creating metrics.

"We have gotten off to a great start in creating guidelines for benchmarks with clearly defined, standardized metrics, but we would like to see wider participation, especially from cloud providers and users," says Hariharan.

www.spec.org

SPEC Announces New Benchmark

June 4, 2012

SPEC's Application Performance Characterization project group announces the release of SPECapc for PTC Creo® 2.0. The benchmark provides eight workflows that exercise all aspects of system performance when running the popular application. Composite scores are generated for graphics, shaded graphics, wireframe graphics, CPU and I/O performance.

www.spec.org

Over 100 Attendees at ICPE 2012

May 7, 2012

Over 100 participants from industry and academia coming from more than a dozen countries spread across four continents attended ICPE 2012 in Boston, USA. The technical program included 19 full research papers, 2 short papers, 2 industrial/experience papers

and 10 work-in-progress papers selected from over 80 submissions over all tracks. It also featured 2 keynote speakers of international repute, 9 poster/demo presentations, 3 tutorials and 2 invited talks from the winners of the SPEC Distinguished Dissertation Award. The next ICPE conference will take place in April 2013 in Prague, Czech Republic.

www.spec.org

New Version of SPECjEnterprise2010

April 5, 2012

SPECjEnterprise2010 v1.03 is now available. This update includes performance-neutral changes to address known issues. Beginning May 1, 2012, all result submissions must be made using version 1.03. Licensees of earlier versions of the SPECjEnterprise2010 software may request a free update to 1.03 by contacting SPEC.

www.spec.org

Research Group Elections

March 23, 2012

Between 5th and 23rd of March 2012, the members of SPEC Research Group voted to elect the Streering Committee and officers of SPEC Research. The following individuals have been elected to the SPEC Research Group Steering Committee: J. Nelson Amaral (University of Alberta), Rema Hariharan (AMD), Willhelm Hasselbring (University of Kiel), Alexandru Iosup (TU Delft), Lizy John (University of Texas at Austin), Samuel Kounev (KIT), Klaus Lange (HP), Meikel Poess (Oracle), Kai Sachs (SAP AG), Seetharami Seelam (IBM TJ Watson Research Center), and Petr Tuma (Charles University Prague).

The steering committee elected the following RG Officers: Chair, Vice-Chair, Secretary and Release Manager. The results of the election are the following; Chair: Samuel Kounev (KIT), Vice-Chair: Kai Sachs (SAP AG), Secretary: Seetharami Seelam (IBM TJ Watson Research Center), Release Manager: Qais Noorshams (KIT).

Piotr Rygielski (KIT)

Repository of Peer-Reviewed Tools Launched

February 2, 2012

A repository of peer-reviewed tools for quantitative system evaluation and analysis is now available at the SPEC RG web site. As a contribution to the community, SPEC RG addresses the need for a collection of tools that have undergone a thorough review process by multiple independent experts to ensure high quality and relevance. The review process covers important quality factors, including maturity, availability and usability.

The repository is intended to have a broad scope, covering system evaluation and analysis with respect to 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. In particular, tools for measurement, profiling, workload characterization, load testing, stress testing and resilience testing are solicited.

Among the first tools to pass the review process are Faban and Kieker. Faban is a framework to automate server performance tests in multi-tier environments, already used by the SPECjEnterprise2010 and SPECsipInfrastructure2011 benchmarks. The tool was originally developed at Sun Microsystems and is distributed under the CDDL 1.0 license. Kieker is a framework for monitoring and analysis of the performance of software systems. Kieker was developed by a research team from the University of Kiel and is distributed under an Apache 2.0 License.

More tools are currently under review and further submissions are encouraged. After a careful review process, the tools are published on the SPEC RG website. The tools are redistributed without modification and the submitter retains the associated rights and responsibilities. As part of the acceptance, SPEC RG may ask the authors to make changes or enhance certain features or aspects but SPEC RG itself does not make any changes.

The members of SPEC RG review the submitted tools using a wide spectrum of criteria – starting with code-related requirements such as source structure and documentation, and ending with practical issues including licensing and support. Thus, the tools are expected to perform as advertised in most configurations. The goal of the review process is to deliver tools that are robust and result in a productive experience for those that commit resources to use them.

http://research.spec.org/tools.html

www.spec.org

NEW SPEC RESEARCH WORKING GROUP: BENCHMARKING ARCHITECTURES FOR INTRUSION DETECTION IN VIRTUALIZED ENVIRONMENTS

July 31, 2012

The cloud computing paradigm, enabled by virtualization technology, continuously gains on popularity since it brings many benefits such as pay-per-use business model and on-demand resource provisioning. However, the on-going wide migration to cloud systems is challenged by many concerns among which security is a major one.

An IDS (intrusion detection system) is a common security mechanism which features detection of malicious events while monitoring regular host and/or network activities. Therefore, many academic and industrial organizations are conducting extensive research on novel IDSes specifically designed to operate in virtualized cloud environments. As the amount and the popularity of such IDSes increase, benchmarking IDSes for cloud environments becomes imperative since it provides insight and deeper understanding of their behavior and performance.

In order to contribute towards addressing the increasing demand for representative and rigorous IDS benchmarks for cloud platforms, the SPEC Research Group has formed a new working group on benchmarking architectures for intrusion detection in virtualized environments. The goal of the working group is to foster and facilitate innovative research through exchange of ideas and experiences. Its membership body currently includes representatives of Siemens Corporate Research (USA), University of Coimbra (Portugal) and Karlsruhe Institute of Technology (Germany). Interested people and organizations are strongly encouraged to join the new group.

The group has already initiated activities such as in-depth analysis of the existing IDS benchmarking efforts, identification of relevant challenges and requirements, and development of benchmarking approaches. On a long-term basis, the working group aims to develop a customizable and flexible IDS benchmarking framework which employs representative cloud workloads, accurate metrics and rigorous benchmarking methodology. The group may also widen its research scope in terms of designing and developing novel security mechanisms applicable in cloud environments, such as effective attack detection techniques and security system architectures.

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

FACE-TO-FACE MEETING OF SPEC RESEARCH GROUP STEERING COMMITTEE

April 26, 2012

The SPEC Research Group (SPEC RG) held a very productive face-to-face meeting during the Boston SPEC meeting on April 26th, 2012. The meeting was attended by 23 members (19 in person and 4 via teleconference). The chair of the SPEC RG, Samuel Kounev from the Karlsruhe Institute of Technology (KIT) commenced the meeting evaluating the first year of the SPEC RG. 35 member organizations from academia and industry joined the RG within the first year. First elections were held and a Steering Committee including Officers was finalized.

Since its inauguration the SPEC RG has been meeting twice every month and the first face-to-face meeting of the SPEC RG was held in Karlsruhe in conjunction with ICPE 2011. The group felt that coinciding face-to-face meetings with the ICPE conference is desirable. Thanks to the dedicated members of the group, the technical infrastructure of the RG and its public web site were setup and populated with relevant information. The first working group of the RG (RG Cloud) focusing on cloud computing was launched mid 2011.

The flagship conference ICPE was established as a regular event with a 12 month cycle. With over 120 participants coming from 20 different countries and 70 organizations worldwide, ICPE 2011 has been a great success. A peer-reviewed tool repository was launched with an acceptance procedure. Two of the initial submissions passed our rigorous review process and were published. The SPEC Distinguished Dissertation Award was established with some very high quality nominations for 2011. The winners of the award presented their work at ICPE 2012.

At this meeting the Steering Committee evaluated the above mentioned accomplishments and self critically discussed ways to strive in 2012. One criticism was around the biweekly meetings. To focus the meetings we will include target amounts of specific topics, increase the time spent on technical issues to 2/3 of the meeting time and make meeting minutes available immediately after the meeting.

As ICPE is now officially a joint endeavor by both WOSP and SIPEW, the ICPE Steering Committee will consist of six people from SPEC RG and six people from WOSP. The SPEC RG Steering Committee discussed how the six people from SIPEW should be elected. It was decided that we will hold official elections by May 5th.

A large portion of the meeting was taken by discussing future activities. One of such activities was to build a repository of traces from real-life systems (performance, security, etc). After Samuel Kounev introduced the idea a lively discussion started about legal issue to share such data. This discussion did not come to a conclusion and will be continued during subsequent conference calls. A second topic was a repository to ensure reproducibility of conference results. A discussion about

the topic resulted in Petr Tůma to write a proposal to be circulated via e-mail.

Samuel Kounev continued the meeting by proposing new working groups. Possible new working group would focus on Security Benchmarking, Big Data and Dependability Benchmarking. Security Benchmarking will be pursued by KIT, Siemens and the University of Coimbra. Meikel will attend a workshop on Big Data (http://clds.ucsd.edu/wbdb2012/) and advertise the SPEC RG as a potential consortia to drive the development of a benchmark. Marco Vieira of the University of Coimbra together with Alberto Avritzer will work on a proposal on Dependability Benchmarking.

Meikel Poess (Oracle Corporation)

SPEC DISTINGUISHED DISSERTATION AWARD WINNERS 2011

December 22, 2011

The Research Group of the Standard Performance Evaluation Corp. (SPEC) has awarded its first two Distinguished Dissertation Awards to Luk Van Ertvelde of Ghent University in Belgium and Kai Sachs of Technische Universität Darmstadt in Germany.

Luk Van Ertvelde was nominated by Professor Lieven Eeckhout from Ghent University. He receives the award for his dissertation titled "Workload Generation for Microprocessor Performance Evaluation."

Kai Sachs was nominated by Professor Alejandro Buchmann from Technische Universität Darmstadt. He receives the award for his dissertation titled "Performance Modeling and Benchmarking of Event-Based Systems."

The awards have been presented at the ICPE 2012 International Conference which will take place on April 22–25, 2012 in Boston, Mass., USA.

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.

"It was a demanding task to select the winners among the high-quality nominations," says the chair of the selection committee, Prof. Wilhelm Hasselbring of the University of Kiel in Germany. "It was great to see such outstanding work for the first awards program from the SPEC Research Group."

Members of the selection committee include: J. Nelson Amaral, University of Alberta, Canada; Wilhelm Hasselbring, University of Kiel, Germany (Chair); Seetharami R. Seelam, IBM Watson Research Center, USA; Petr Tuma, Charles University of Prague, Czech Republic; Walter Bays, Oracle Corp., USA.

www.spec.org

SPEC President Walter Bays presents the award to Luk Van Ertvelde.

Award Winner Dissertation: Workload Generation for Microprocessor Performance Evaluation

December 22, 2011

Although the availability of standardized benchmark suites has streamlined the process of performance evaluation, computer architects and engineers still face several important benchmarking challenges.

Benchmarks should be representative of the applications that are expected to run on the target computer system; however, it is not always possible to compose a set of representative benchmarks. The main reason is that standardized benchmark suites are typically derived from open-source programs – because industry hesitates to share proprietary applications – which may not be representative of the real-world applications of interest.

The amount of redundancy within and across benchmarks should be as small as possible. However, contemporary benchmark suites execute trillions of instructions to stress microprocessors in a meaningful way. As a result, it is infeasible to simulate entire benchmark suites using detailed cycle-accurate simulators.

Benchmarks should enable micro-architecture, architecture and compiler research and development. Although existing benchmark suites satisfy this requirement, this is often not the case for benchmark reduction techniques because they typically operate at the assembly-level.

In this dissertation, I propose three novel benchmark generation and reduction techniques to address these challenges.

Code mutation [1] is a novel methodology that mutates a proprietary application to complicate reverse-engineering so that it can be distributed as a benchmark among industry and academia. These benchmark mutants hide the functional semantics of proprietary applications, while exhibiting similar performance characteristics. Consequently, they can be used as proxies for proprietary software to help drive performance evaluation by third parties.

Code mutation conceals the intellectual property of an application, but it does not lend itself to the generation of short-running benchmarks. Sampled simulation on the other hand reduces the simulation time of a benchmark by only simulating a small sample from a complete benchmark execution in a detailed manner. In sampled simulation, the performance bottleneck is the establishment of the micro-architectural state (particularly the state of the caches) at the beginning of each sampling unit, often referred to as the cold-start problem. I address this problem by proposing a new cache warm-up methodology, namely NSL-BLRL [2, 3],

which reduces sampled simulation time by an order of magnitude compared to the state-of-the-art.

Although code mutation can be used in combination with sampled simulation to generate short-running workloads that can be distributed to third parties without revealing intellectual property, the limitation is that this approach operates at assembly-level. This excludes them from being used for architecture and compiler research. We therefore propose a novel benchmark synthesis methodology and framework [4, 5] that aims at generating small but representative benchmarks in a high-level programming language, so that they can be used to explore both the architecture and compiler spaces.

[1] Luk Van Ertvelde and Lieven Eeckhout, "Dispersing Proprietary Applications as Benchmarks through Code Mutation", in Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2008, 201–210

[2] Luk Van Ertvelde, Filip Hellebaut, Lieven Eeckhout and Koen De Bosschere, "NSL-BLRL: Efficient Cache Warmup for Sampled Processor Simulation", in Proceedings of the Annual Simulation Symposium (ANSS), 2006, 87–96

[3] Luk Van Ertvelde, Filip Hellebaut and Lieven Eeckhout, "Accurate and Efficient Cache Warmup for Sampled Processor Simulation through NSL-BLRL", in The Computer Journal, Vol. 51, No. 2, 192–206, March 2008

[4] Luk Van Ertvelde and Lieven Eeckhout, "Benchmark Synthesis for Architecture and Compiler Exploration", in Proceedings of the IEEE International Symposium on Workload Characterization (IISWC), 2010, 106–116

[5] Luk Van Ertvelde and Lieven Eeckhout, "Workload Reduction and Generation Techniques", in IEEE Micro, Nov/Dec 2010, Vol.30, No.6

Luk Van Ertvelde (Ghent University, Belgium)

Award Winner Dissertation: Performance Modeling and Benchmarking of Event-Based Systems

December 22, 2011

Event-based systems (EBS) are increasingly used as underlying technology in many mission critical areas and large-scale environments, such as environmental monitoring and location-based services. Moreover, novel event-based applications are typically highly distributed and data intensive with stringent requirements for performance and scalability. Common approaches to address these requirements are benchmarking and performance modeling. However, there was a lack of general performance modeling methodologies for EBS as well as test harnesses and benchmarks using representative workloads for EBS. Therefore, this thesis focused on approaches to benchmark EBS as well as the development of a performance modeling methodology. In this context, novel extensions for queueing Petri nets (QPNs) were proposed. The motivation was to support the development and maintenance of EBS that meet certain Quality-of-Service (QoS) requirements.

To address the lack of representative workloads we developed the first industry standard benchmark for EBS jointly with the Standard Performance Evaluation Corporation (SPEC) in whose development and specification the author was involved as a chief benchmark architect and lead developer. Our efforts resulted in the SPECjms2007 standard benchmark. Its main contributions were twofold: based on the feedback of industrial partners, we specified a comprehensive standardized workload with different scaling options and implemented the benchmark using a newly developed complex and flexible framework. Using the SPECjms2007 benchmark we introduced a methodology for performance evaluation of message-oriented middleware platforms and showed how the workload can be tailored to evaluate selected performance aspects. The standardized workload can be applied to other EBS. For example, we developed an innovative research benchmark for publish/subscribe-based communication named jms2009-PS based on the SPECjms2007 workload. The proposed benchmarks are now the de facto standard benchmarks for evaluating messaging platforms and have already been used successfully by several industrial and research organizations as a basis for further research on performance analysis of EBS.

Kai Sachs (right) receives the award from Walter Bays.

To describe workload properties and routing behavior we introduced a novel formal definition of EBS and their performance aspects. Furthermore, we proposed an innovative approach to characterize the workload and to model the performance aspects of EBS. We used operational analysis techniques to describe the system traffic and derived an approximation for the mean event delivery latency. We showed how more detailed performance models based on QPNs could be built and used to provide more accurate performance prediction. It is the first general performance modeling methodology for EBS and can be used for an in-depth performance analysis as well as to identify potential bottlenecks. A further contribution is a novel terminology for performance modeling patterns targeting common aspects of event-based applications using QPNs.

To improve the modeling power of QPNs, we defined several extensions of the standard QPNs. They allow us to build models in a more flexible and general way and address several limitations of QPNs. By introducing an additional level of abstraction, it is possible to distinguish between logical and physical layers in models. This enables to flexibly map logical to physical resources and thus makes it easy to customize the model to a specific deployment. Furthermore, we addressed two limiting aspects of standard QPNs: constant cardinalities and lack of transition priorities.

Finally, we validated our methodology to model EBS in two case studies and predicted system behavior and performance under load successfully. As part of the first case study we extended SIENA, a well-known

distributed EBS, with a runtime measurement framework and predicted the runtime behavior including delivery latency for a basic workload. In the second case study, we developed a comprehensive model of the complete SPECjms2007 workload. To model the workload we applied our performance modeling patterns as well as our QPN extensions. We considered a number of different scenarios with varying workload intensity (up to 4,500 transactions / 30,000 messages per second) and compared the model predictions against measurements. The results demonstrated the effectiveness and practicality of the proposed modeling and prediction methodology in the context of a real-world scenario.

This thesis opens up new avenues of frontier research in the area of event-based systems. Our performance modeling methodology can be used to build self-adaptive EBS using automatic model extraction techniques. Such systems could dynamically adjust their configuration to ensure that QoS requirements are continuously met.

Kai Sachs (TU Darmstadt, SAP AG)

CALL-FOR-NOMINATIONS: SPEC DISTINGUISHED DISSERTATION AWARD 2012

July 25, 2012

The SPEC Distinguished Dissertation Award (formally called SPEC Benchmarking Research PhD Award) aims to recognize outstanding doctoral 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 the SPEC 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.

The winner will receive $1000, which will be awarded at the ICPE 2013 International Conference. A nomination should include:

- A nomination letter (including the name of the student, the title of the dissertation, the institution where the dissertation was defended, and the date of the defense).
- A one page report that outlines the outstanding contribution of the dissertation.
- The dissertation itself including a one page extended abstract of the dissertation. (If the dissertation is written in language other than English, it may be accompanied by publications, in English, describing the same research as the dissertation.)

The SPEC Distinguished Dissertation Award is open to dissertations that have been defended between October 1, 2011, and September 30, 2012. If there are several outstanding submissions, the committee may split the award between them. The submission deadline is September 30, 2012.

Nominations are welcome at any time before the final submission deadline. Nominations, or questions about the application process, should be sent by e-mail to: nominations@spec.org Self-nominations are not accepted.

Selection committee: Chair: Wilhelm Hasselbring, University of Kiel, Germany; David A. Bader, Georgia Institute of Technology, USA; Walter Bays, Oracle Corporation, USA; Walter Binder, University of Lugano, Switzerland; Rema Hariharan, AMD, USA; Samuel Kounev, Karlsruhe Institute of Technology (KIT), Germany; Diwakar Krishnamurthy, University of Calgary, Canada; Klaus Lange, HP, USA; John Murphy, University College Dublin, Ireland; Marco Vieira, University of Coimbra, Portugal.

Qais Noorshams (KIT)

KIEKER: A FRAMEWORK FOR APPLICATION PERFORMANCE MONITORING AND DYNAMIC SOFTWARE ANALYSIS

July 9, 2012

Kieker is a framework for monitoring and analyzing the runtime behavior of concurrent and distributed software systems. It is maintained by the Software Engineering Group at the University of Kiel, Germany. Kieker development started in 2006 as a small tool for monitoring response times of Java software operations. Since then, it has evolved into a powerful and extensible framework for application performance management and dynamic software analysis, developed and employed for various purposes in research, teaching, and industrial practice. Application areas include performance and dependency analysis, online capacity management, simulation, as well as extraction of architecture-level performance and usage models. Being designed for continuous monitoring in production environments, it imposes only a low performance overhead, as evaluated experimentally with micro- and macro-benchmarks. Kieker has also been used for dynamic analysis of production systems in industrial contexts.

The framework is structured into a monitoring and an analysis component. On the monitoring side, monitoring probes collect measurements represented as monitoring records, which a monitoring writer passes to a configured monitoring log or stream. On the analysis side, monitoring readers import monitoring records of interest from the monitoring log/stream and pass them to a configurable pipe-and-filter architecture of analysis plugins. Focusing on application-level analysis, Kieker includes monitoring probes for collecting timing and

trace information from distributed executions of software operations. Moreover, probes for sampling system-level measures, e.g., CPU utilization and memory usage, are included. In addition to instrumentation and monitoring support for Java, adapters for other platforms, such as .NET, COM, and COBOL, have been developed recently. Kieker supports monitoring logs and streams utilizing file systems, databases, messaging queues, etc. A number of plugins for extracting, analyzing, and visualizing software architectural models such as calling-dependency graphs, call trees, and sequence diagrams are included. Utilizing the framework's extensibility, custom components can be developed easily, if required.

Kieker is open-source software, freely available at http://kieker-monitoring.net. In 2011, Kieker was reviewed and accepted for distribution as part of the SPEC Research Group's repository of peer-reviewed tools for quantitative system evaluation and analysis (http://research.spec.org/projects/tools.html). At ICPE 2012, we gave an invited poster/demo presentation of the framework.

http://dx.doi.org/10.1145/2188286.2188326

André van Hoorn (Christian-Albrechts-University of Kiel)

FABAN: A PLATFORM FOR PERFORMANCE TESTING OF SOFTWARE

June 21, 2012

Faban is a free and open source platform for performance testing of software applications. Faban can be used in performance, scalability and load testing of almost any type of server application including complex multi-tiered (web/cache/database type) applications. At the same time, Faban can also be used to develop a micro-benchmark that tests a single server such as a ftp server. If the application accepts requests on a network, Faban can test and measure it.

Faban provides an API to develop the workloads or benchmarks (called the Faban Driver Framework) and a runtime execution environment (known as the Faban Harness). It automates the handling of client threads, timing, metrics collection and production of reports.

Faban Driver Framework is an API-based framework and component model to help develop new workloads rapidly. The driver framework controls the lifecycle of the benchmark run as well as the stochastic model used to simulate users. It provides built-in support for many servers (e.g., Apache httpd, MySQL, Oracle RDBMS, memcached) automating their start/stop and stats collection. It provides a well documented interface to add support for any other server.

Faban Harness is a tool to automate running of workloads. Faban provides an easy to use web interface to configure and queue runs, and includes functionality to view, compare graph run outputs and archive results. For more information, visit the Faban website http://faban.org .

Shanti Subramanyam

REPORT FROM THE 3RD ACM/SPEC INTERNATIONAL CONFERENCE ON PERFORMANCE ENGINEERING (ICPE 2012)

August 7, 2012

The ACM/SPEC ICPE 2012 took place in the Omni Parker House, Boston from the 22nd to 25th of April. Local general chair David Kaeli emphasized in his opening remarks the historical meaning of the hotel: it was the workplace for well-known persons such as Hô Chí Minh and Malcolm X. John F. Kennedy announced his candidacy for the Congress in the press room, which was also used for the conference.

This place provided the right spirit for the event. More than 100 participants from industry and academia came together to discuss the latest research results in the field of performance engineering. The conference program covered a variety of topics, including benchmarking, energy efficiency, performance modeling, performance and software development processes, performance evaluation of cloud, data-intensive and adaptive systems. The technical program included 19 full research papers, 2 short papers, 2 industrial/experience papers and 10 work-in-progress papers selected from over 80 submissions over all tracks. The program was enriched by social events, a poster & demo exhibition, and the presentation of the winners of the SPEC Distinguished Dissertation Award (see page 6). The presentations of the keynotes were two highlights of the conference. In "Assuring the Trustworthiness of the Smarter Electric Grid", Prof. William H. Sanders discussed the challenges in assuring the trustworthiness – performance, dependability, and security – of the merging smart grid. The second keynote was presented by Dr. Amnon Naamad from EMC and discussed "New Challenges in Performance Engineering" with regards to new technologies such as virtualization and cloud, their impact on the IT, and tools helping to exploit the technologies.

The PC-Chairs had the difficult challenge to select the winners of the best paper awards. Best paper awards were presented in the following categories: Best Innovation, Best Industry-Related and Best Student Paper (for more details see page 10).

Kai Sachs (SAP AG)

ANNOUNCEMENT OF ICPE 2013

June 17, 2012

The 4th ACM/SPEC International Conference on Performance Engineering (ICPE) will take place on April 21–24, 2013, in Prague, Czech Republic.

ICPE is a joint meeting of the Association for Computing Machinery (ACM) and the Standard Performance Evaluation Corporation (SPEC), which provides a forum for the integration of theory and practice in the field of performance engineering. ICPE brings together researchers and industry practitioners that share experiences, discuss challenges, and report work-in-progress and state-of-the-art results in performance engineering of software and systems.

Historically, ICPE has grown from the ACM Workshop on Software and Performance (WOSP) and the SPEC International Performance Evaluation Workshop (SIPEW). This unique connection to both industrial and academic background gives ICPE participants a singular advantage of connecting the theoretical and practical perspectives on performance engineering – which is also reflected in the variety of contribution styles, including: Traditional research papers on both basic and applied research, covering any of the wide range of topics related to performance engineering, to be reviewed by the scientific program committee; Industrial and experience papers covering novel applications and innovative implementations of performance engineering, interesting performance results and industrial experience, to be reviewed by the industrial program committee; Work-in-progress and vision papers provide space for reporting on promising preliminary results that have not yet been fully validated, emerging research challenges and long-term research directions.

The conference proceedings will be published by ACM and included in the ACM Digital Library. Authors are also invited to submit auxiliary technical artifacts of their work (experimental data, evaluation scripts), thus increasing the potential for further use, citation or expansion of their results.

Among the solicited paper topics one can find the following. In the category of 'Performance and Software Development Processes' the topics include: Performance engineering for systems including, but not limited to: smart grids, cloud platforms, storage systems, sensor nets, embedded and real-time systems, control systems, multi-tier systems, event-driven systems; Techniques to elicit and incorporate performance, availability, power and other extra-functional requirements throughout the software and system life cycle; Agile performance-test-driven development; Performance-requirements engineering and design for performance predictability; Software performance patterns and anti-patterns.

In the category of 'Performance modeling and prediction' the topics include: Languages, annotations, tools and methodologies to support model-based performance engineering; Analytical, simulation, statistical, AI-based, and hybrid modeling/prediction methods; Automated model building; Model validation and calibration techniques.

In the category of 'Performance measurement, testing and experimental analysis' the topics include: Performance measurement, monitoring, and workload characterization techniques; Test planning, tools for performance load testing, measurement, profiling and tuning; Methodologies for performance testing and functional testing; Reproducibility of performance studies.

In the category of 'Benchmarking, configuration, sizing, and capacity planning' the topics include: Benchmark design and benchmarking methods, metrics, and suites; Development of new, configurable, and/or scalable benchmarks; Use of benchmarks in industry and academia; System configuration, sizing and capacity planning techniques.

In the category of 'System management/optimization' the topics include: Use of models for run-time application configuration/management; Online performance prediction and model parameter estimation; Adaptive resource management.

In the category of 'Performance in Cloud, virtualized and multi-core systems' the topics include: Modeling, monitoring, and testing of cloud computing platforms and applications; Performance/management of virtualized machines, storage and networks; Performance engineering of multi-core and parallel systems.

In the category of 'Performance and power' the topics include: Algorithms for combined power and performance management; Instrumentation, profiling, modeling and measurement of power consumption; Power/performance engineering in grid/cluster/cloud/mobile computing systems.

http://icpe2013.ipd.kit.edu/

Petr Tuma (Charles University of Prague)

ICPE 2012: BEST PAPER AWARDS

June 21, 2012

At the 2012 International Conference on Performance Engineering (ICPE 2012) in Boston, USA, three outstanding papers were distinguished by the Program and Awards Chairs. The awards were presented in three different categories: Best Innovation, Industry-Related, and Student Paper Award.

The Best Innovation Paper Award was presented to: Giuliano Casale and Peter Harrison (Imperial College London) for the paper "A Class of Tractable Models for Run-Time Performance Evaluation".

The distinction – Best Industry-Related Paper Award – was presented to: Thijmen de Gooijer, Anton Jansen (ABB Corporate Research, Västerås, Sweden), Heiko Koziolek (ABB Corporate Research, Ladenburg, Germany) and Anne Koziolek (University of Zurich, Switzerland) for the paper "An Industrial Case Study of Performance and Cost Design Space Exploration".

The Best Student Paper Award was presented to: Vlastimil Babka, Peter Libič, Tomáš Martinec and Petr

Tůma (Charles University in Prague, Czech Republic) for the paper "On the Accuracy of Cache Sharing Models". Congratulations to award winners!

Piotr Rygielski (KIT)

ICPE 2012: ABSTRACTS OF SELECTED PAPERS

July 21, 2012

Giuliano Casale, Peter Harrison: A Class of Tractable Models for Run-Time Performance Evaluation

Run-time resource allocation requires the availability of system performance models that are both accurate and inexpensive to solve. We here propose a new methodology for run-time performance evaluation based on a class of closed queueing networks. Compared to exponential product-form models, the proposed queueing networks also support the inclusion of resources having first-come first-served scheduling under non-exponential service times. Motivated by the lack of an exact solution for these networks, we propose a fixed point algorithm that approximates performance indexes in linear time and linear space with respect to the number of requests considered in the model. Numerical evaluation shows that, compared to simulation, the proposed models solved by fixed-point iteration have errors of about 1%–6%, while, on the same test cases, exponential product-form models suffer errors even in excess of 100%. Execution times on commodity hardware are of the order of a few seconds or less, making the proposed methodology practical for runtime decision-making.

http://dx.doi.org/10.1145/2188286.2188299

Thijmen de Gooijer, Anton Jansen, Heiko Koziolek, Anne Koziolek: An Industrial Case Study of Performance and Cost Design Space Exploration

Determining the trade-off between performance and costs of a distributed software system is important as it enables fulfilling performance requirements in a cost-efficient way. The large amount of design alternatives for such systems often leads software architects to select a suboptimal solution, which may either waste resources or cannot cope with future workloads. Recently, several approaches have appeared to assist software architects with this design task. In this paper, we present a case study applying one of these approaches, i.e. PerOpteryx, to explore the design space of an existing industrial distributed software system from ABB. To facilitate the design exploration, we created a highly detailed performance and cost model, which was instrumental in determining a cost-efficient architecture solution using an evolutionary algorithm. The case study demonstrates the capabilities of various modern performance modeling tools and a design space exploration tool in an industrial setting, provides lessons learned, and helps other software architects in solving similar problems.

http://dx.doi.org/10.1145/2188286.2188319

Vlastimil Babka, Peter Libič, Tomáš Martinec, Petr Tůma: On the Accuracy of Cache Sharing Models

Memory caches significantly improve the performance of workloads that have temporal and spatial locality by providing faster access to data. Current processor designs have multiple cores sharing a cache. To accurately model a workload performance and to improve system throughput by intelligently scheduling workloads on cores, we need to understand how sharing caches between workloads affects their data accesses.

Past research has developed analytical models that estimate the cache behavior for combined workloads given the stack distance profiles describing these workloads. We extend this research by presenting an analytical model with contributions to accuracy and composability – our model makes fewer simplifying assumptions than earlier models, and its output is in the same format as its input, which is an important property for hierarchical composition during software performance modeling.

To compare the accuracy of our analytical model with earlier models, we attempted to reproduce the reported accuracy of those models. This proved to be difficult. We provide additional insight into the major factors that influence analytical model accuracy.

http://dx.doi.org/10.1145/2188286.2188294

Alexander Wert, Jens Happe, Dennis Westermann: Integrating Software Performance Curves with the Palladio Component Model

Software performance engineering for enterprise applications is becoming more and more challenging as the size and complexity of software landscapes increases. Systems are built on powerful middleware platforms, existing software components, and 3rd party services. The internal structure of such a software basis is often unknown especially if business and system boundaries are crossed. Existing model driven performance engineering approaches realize a pure top down prediction approach. Software architects have to provide a complete model of their system in order to conduct performance analyses. Measurement-based approaches depend on the availability of the complete system under test. In this paper, we propose a concept for the combination of model-driven and measurement-based performance engineering. We integrate software

performance curves with the Palladio Component Model (PCM) (an advanced model based performance prediction approach) in order to enable the evaluation of enterprise applications which depend on a large software basis.

http://dx.doi.org/10.1145/2188286.2188339

Sadeka Islam, Kevin Lee, Alan Fekete, Anna Liu: How a Consumer Can Measure Elasticity for Cloud Platforms

One major benefit claimed for cloud computing is elasticity: the cost to a consumer of computation can grow or shrink with the workload. This paper offers improved ways to quantify the elasticity concept, using data available to the consumer. We define a measure that reflects the financial penalty to a particular consumer, from under-provisioning (leading to unacceptable latency or unmet demand) or over-provisioning (paying more than necessary for the resources needed to support a workload). We have applied several workloads to a public cloud; from our experiments we extract insights into the characteristics of a platform that influence its elasticity. We explore the impact of the rules used to increase or decrease capacity.

http://dx.doi.org/10.1145/2188286.2188301

Rasha Tawhid, Dorina Petri: User-Friendly Approach for Handling Performance Parameters during Predictive Software Performance Engineering

A Software Product Line (SPL) is a set of similar software systems that share a common set of features. Instead of building each product from scratch, SPL development takes advantage of the reusability of the core assets shared among the SPL members. In this work, we integrate performance analysis in the early phases of SPL development process, applying the same reusability concept to the performance annotations. Instead of annotating from scratch the UML model of every derived product, we propose to annotate the SPL model once with generic performance annotations. After deriving the model of a product from the family model by an automatic transformation, the generic performance annotations need to be bound to concrete product-specific values provided by the developer. Dealing manually with a large number of performance annotations, by asking the developer to inspect every diagram in the generated model and to extract these annotations is an error-prone process. In this paper we propose to automate the collection of all generic parameters from the product model and to present them to the developer in a user-friendly format (e.g., a spreadsheet per diagram, indicating each generic parameter together with guiding information that helps the user in providing concrete binding values). There are two kinds of generic parametric annotations handled by our approach: product-specific (corresponding to the set of features selected for the product) and platform-specific (such as device choices, network connections, middleware, and runtime environment). The following model transformations for (a) generating a product model with generic annotations from the SPL model, (b) building the spreadsheet with generic parameters and guiding information, and (c) performing the actual binding are all realized in the Atlas Transformation Language (ATL).

http://dx.doi.org/10.1145/2188286.2188304

Mike G. Tricker, Klaus-Dieter Lange, Jeremy A. Arnold, Hansfried Block, Christian Koopmann: The implementation of the server efficiency rating tool

The Server Efficiency Rating Tool (SERT) has been developed by Standard Performance Evaluation Corporation (SPEC) at the request of the US Environmental Protection Agency (EPA), prompted by concerns that US data centers consumed almost 3% of all energy in 2010. Since the majority was consumed by servers and their associated heat dissipation systems the EPA launched the ENERGY STAR Computer Server program, focusing on providing projected power consumption information to aid potential server users and purchasers. This program has now been extended to a world-wide audience.

This paper expands upon the one published in 2011, which described the initial design and early development phases of the SERT. Since that publication, the SERT has continued to evolve and has entered the first Beta phase in October 2011 with the goal of being released in 2012. This paper describes more of the details of how the SERT is structured. This includes how components interrelate, how the underlying system capabilities are discovered, and how the various hardware subsystems are measured individually using dedicated worklets.

http://dx.doi.org/10.1145/2188286.2188307

4th ACM/SPEC International Conference on Performance Engineering

ICPE 2013

Prague - Czech Republic - April 21-24

The International Conference on Performance Engineering (ICPE) provides a forum for the integration of theory and practice in the field of performance engineering. ICPE 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 and present their experiences, discuss challenges, and report state-of-the-art and in-progress research on performance engineering of software and systems. Topics of interest include:

Performance and software development processes

- Techniques to elicit and incorporate performance, availability, power and other extra-functional requirements throughout the software and system lifecycle
- Agile performance-test-driven development
- Performance-requirements engineering and design for software performance predictability
- Software performance modeling, patterns and anti-patterns

Performance modeling and prediction

- Languages, annotations, tools and methodologies to support model-based performance engineering
- Analytical, simulation, statistical, AI-based, and hybrid modeling/prediction methods
- Automated model discovery and model building
- Model validation and calibration techniques

Performance measurement and experimental analysis

- Performance measurement, monitoring, and workload characterization techniques
- Test planning, tools for performance, load testing, measurement, profiling and tuning
- Model extraction for functional or partly functional systems
- Methodologies for performance testing and for functional testing
- Reproduction and reproducibility of performance studies

System management/optimization

- Use of models for run-time configuration/management
- Online performance prediction and model parameter estimation
- Adaptive resource management

Benchmarking, configuration, sizing, and capacity planning

- Benchmark design and benchmarking methods, metrics, and suites
- Development of new, configurable, and/or scalable benchmarks
- Use of benchmarks in industry and academia
- System configuration, sizing and capacity planning techniques

Performance in Cloud, virtualized and multi-core systems

- Modeling, monitoring, and testing of cloud platforms and applications
- Performance/management of virtualized machines, storage and networks
- Performance engineering of multi-core and parallel systems

Performance and power

- Algorithms for combined power and performance management
- Instrumentation, profiling, modeling and measurement of power consumption
- Power/performance engineering in grid/cluster/cloud/mobile computing systems

Performance modeling and evaluation in other domains such as:

- Web-based systems, e-business, web services, SOAs
- Transaction-oriented and event-based systems
- Embedded and autonomous systems
- Real-time and multimedia systems
- Peer-to-peer, mobile and wireless systems

Multiple different kinds of papers are sought: basic and applied research, industrial/experience reports, and work-in-progress/vision papers. Different acceptance criteria apply for each category. The conference proceedings will be published by ACM and included in the ACM Digital Library. The best conference paper will receive a Best Paper Award.

Important Dates

Research paper submissions:	**24 Sep 2012**	Poster and demo papers submissions:	**13 Nov 2012**
Research paper notification:	**14 Dec 2012**	Tutorial proposals submissions:	**10 Nov 2012**
Industrial/experience paper submissions:	**30 Oct 2012**	Work-in-progress/vision paper submissions	**14 Jan 2013**

Organizing Committee

General Chairs
Petr Túma, Charles University, Czech Republic
Giuliano Casale, Imperial College London, UK
Program Chairs
J. Nelson Amaral, University of Alberta, Canada
Tony Field, Imperial College of London, UK
Industrial Chair
Seetharami R. Seelam, IBM Research, USA

Tutorial Chair
Mirco Tribastone, LMU Munich, Germany
Demos and Posters Chair
Kaustubh Joshi, AT&T Labs Research, US
Publicity Chairs
John Murphy, University College Dublin, Ireland
Kai Sachs, SAP AG, Germany

Finance Chair
Lubomír Bulej, Charles University, Czech Republic
Publication Chair
Evgenia Smirni, College of William and Mary, USA
Awards Chairs
Lucy Cherkasova, HP Labs, USA
Vittorio Cortellessa, Universita' di L'Aquila, Italy

Program Committee

Martin Arlitt, HP Labs, USA, University of Calgary, Canada
Alberto Avritzer, Siemens Corporate Research, USA
Steffen Becker, University of Paderborn, Germany
Anne Benoit, ENS Lyon – LIP, France
Steve Blackburn, ANU, Australia
Andre Bondi, Siemens Corporate Research, USA
Edson Borin, Universidade de Campinas, Brazil
Jeremy Bradley, Imperial College London, UK
Lydia Chen, IBM Zurich, Switzerland
Lucy Cherkasova, Hewlett-Packard Laboratories, USA
Lawrence Chung, University of Texas at Dallas, USA
Susanna Donatelli, Universita' di Torino, Italy
Sandhya Dwarkadas, University of Rochester, USA
Dick Epema, Delft University of Technology, Netherlands
Dror Feitelson, Hebrew University, Israel
Sebastian Fischmeister, University of Waterloo, Canada
Stephen Gilmore, University of Edinburgh, UK
Lars Grunske, TU Kaiserslautern, Germany
Wilhelm Hasselbring, University of Kiel, Germany
Michael Hind, IBM T.J. Watson Research Center, USA
Robert Hundt, Google Inc., USA
Alexandru Iosup, TU Delft, The Netherlands
Stephen Jarvis, University of Warwick, UK
Carlos Juiz, University of the Balearic Islands, Spain
David Kaeli, Northeastern University, USA
William Knottenbelt, Imperial College London, UK
Samuel Kounev, Karlsruhe Institute of Technology, Germany
Heiko Koziolek, ABB Corporate Research, Germany
Anirban Mahanti, NICTA, Australia
Pat Martin, Queen's University, Canada
Daniel Menasce, George Mason University, USA
Raffaela Mirandola, Politecnico di Milano, Italy
David Pearce, Victoria University of Wellington, N. Zealand
Dorina Petriu, Carleton University, Canada
Meikel Poess, Oracle Corporation, USA
Lawrence Rauchwerger, Texas A&M University, USA
Ralf Reussner, Karlsruhe Institute of Technology, Germany
George Riley, Georgia Institute of Technology, USA
Alma Riska, College of William and Mary, USA
Jerry Rolia, Carleton University, USA
Peter Sweeney, IBM T.J. Watson Research Center, USA
Mirco Tribastone, LMU Munich, Germany
Catia Trubiani, Universita' di L'Aquila, Italy
Carey Williamson, University of Calgary, Canada
Murray Woodside, Carleton University, Canada
Peng Wu, IBM T. J. Watson Research Center, USA
Xiaoyun Zhu, VMware, USA