



SPEC ResearchSM Group Newsletter

CONTENTS OF THIS ISSUE

- 2** SPEC Research Group Officers
- 2** SPEC Research Working Groups
- 2** Welcome to the SPEC Research Group Newsletter
- 3** SPEC Research Group Mission Statement
- 3** SPEC Announcements
- 5** Winner of the SPEC Distinguished Dissertation Award 2012
- 6** Call-for-Nominations: SPEC Distinguished Dissertation Award 2013
- 7** FINCoS: A framework for Performance Evaluation of Event Processing Systems
- 7** DiSL: A Domain-Specific Language for Instrumentation-Based Dynamic Program Analysis on the Java Platform
- 8** Report of the SPEC RG IDS Benchmarking Working Group
- 8** Annual Report of the Cloud Working Group
- 9** New Technical Report in RG Cloud Working Group: Cloud Usage Patterns
- 10** Report from the Fourth ACM/SPEC International Conference on Performance Engineering (ICPE 2013) and Collocated Events
- 10** ICPE 2013: Best Paper Awards
- 11** Short Reports from SPEC Face-to-Face Meetings
- 11** SPECTacular Awards for Outstanding Service to SPEC in 2012
- 12** ICPE 2013: Abstracts of Selected Papers
- 14** ICPE 2014: Call for Papers

CALL FOR NOMINATIONS – SPEC DISTINGUISHED DISSERTATION AWARD

The SPEC Distinguished Dissertation Award aims to recognize outstanding doctoral dissertations in the field of computer benchmarking, performance evaluation, and experimental system analysis in general. The main prize – \$1000 – will be awarded at ICPE 2014.

Read more on page 6

NEW TOOLS IN PEER-REVIEWED TOOLS REPOSITORY

The SPEC Research Repository of peer-reviewed tools for quantitative system evaluation and analysis includes two new tools: DiSL and FINCoS.

Read more on page 7

NEW TECHNICAL REPORT: CLOUD USAGE PATTERNS

The SPEC Research Group announces a new technical report on a textual and graphical formalism for describing common cloud usage scenarios. The document is targeted at a wide audience with possibilities of application in various areas of Cloud computing providing a common reference for usage scenarios.

Read more on page 9



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>

THE ‘BEST PAPER’ AWARDS FROM THE ICPE 2013 CONFERENCE

At the 2013 International Conference on Performance Engineering (ICPE 2013) in Prague, Czech Republic, five outstanding papers were distinguished by the Program and Awards Chairs. The awards were presented in two different categories: Best Industry-Related Paper, and Best Research Paper.

Read more on page 10

SPEC RESEARCH GROUP OFFICERS

Chair:

Samuel Kounev, Karlsruhe Institute of Technology (KIT), Germany, kounev@kit.edu

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

Seetharam R. Seelam, IBM TJ Watson Research Center, USA

Petr Tůma, Charles University of Prague, Czech Republic

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 2nd 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 six months have been very intense and exciting for the SPEC RG. Some of the major activities and milestones reached include: launching of a new working group on intrusion detection systems (IDS); publishing the first official RG technical report on cloud usage patterns—a textual and visual formalism for describing common cloud usage scenarios; establishing a long-time cooperation between ACM and SPEC for future ICPE conferences; and finalizing the ICPE constitution. In addition, we had a very successful ICPE 2013 conference and HotTopiCS 2013 workshop, which attracted more than 160 attendees from industry and academia.

The SPEC Distinguished Dissertation Award 2012 attracted again many high-quality nominations making the selection of the winner very challenging. The winner of this years' award is Dr. Shicong Meng from Georgia Tech for his thesis "Monitoring-as-a-Service in the Cloud" supervised by Prof. Ling Liu from the Georgia Institute of Technology.

We would like to use this opportunity to welcome our new members: MIT Lincoln Labs, Hasso Plattner Institute, University of North Florida, and University of Stuttgart.

In this newsletter you will find many interesting articles that describe achievements and major events from the last seven months of SPEC RG's activity. We want to highlight several articles that may be of special interest to the reader. Within the SPEC Announcements you can find the call for new members of the Cloud Working Group. We also invite nominations to the SPEC Distinguished Dissertation Award. We present two new tools that were accepted for the Peer-reviewed Tools Repository. Reports from activities of our working groups will give the insight into our active commitment in the Cloud and IDS research domains. Our first technical report on "Cloud Usage Patterns" is highly relevant to the Cloud Computing community. We close the issue with reports from ICPE 2013, the annual SPEC meetings and the call for papers for ICPE2014.

We hope that 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), Piotr Rygielski (KIT)

SPEC, the SPEC logo and the names SPECjbb, SPECjAppServer, SPECpower_ssj, SPECvirt_sc, SPECweb, 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 © 1988-2013 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

We Want You!

SPEC RG Cloud WG is Looking for New Members

May 8, 2013

If you work in cloud computing and have an interest in performance, you are welcome to join the Cloud Working Group of the SPEC Research Group. We understand performance in a broad sense, including cloud benchmarking, quantitative evaluation, and experimental analysis.

By joining the Cloud Working Group, you will get access to a dedicated group of industry and academic experts with whom you can engage in meaningful discussions about how the performance of cloud-based systems can be measured, you will be able to join the SPEC standardization processes, and you will find a new group of potential employees or potential places to send your students to do internships.

Within the Cloud Working Group, we are committed to rapidly integrating our new members, in either ongoing or emerging activities, in participatory or leadership roles. We also keep the overhead low (under 10% of the scheduled meeting time), so that a vast majority of your and our efforts can be put into meaningful work and content.

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

Alexandru Iosup (TU Delft)

SPEC releases Server Efficiency Rating Tool (SERT)

March 23, 2013

On February 26th, 2013 SPEC released the first version of the Server Efficiency Rating Tool (SERT). Developed at the request of the US Environmental Protection Agency (EPA) the SERT will be used to

generate power efficiency data as input to the newly launched ENERGY STAR for Computer Servers Version 2 program. The SERT was developed by the same team that previously created SPECpower_ssj2008, which has become the industry standard for measuring the power efficiency of servers.

The SERT offers significantly enhanced ease of use, with automated sophisticated hardware configuration discovery, power analyzer range setting and a graphical user interface targeting both naïve and expert users. The expert mode targets researchers and provides configuration flexibility that offers full control over how individual workloads are invoked or combined. The workloads included in the first release test individual subsystems; CPU, memory, storage, and a combined workload.

Initial use of the SERT in the ENERGY STAR for servers V2 program will provide power consumption data to the EPA. These data will also be used by the SPECpower Committee to further enhance the SERT, and specifically will be used to aid in the creation of power efficiency metrics for subsequent releases. The SERT has been designed to be highly flexible and extensible, and its use by the EPA is being reviewed for potential use by international organizations across the world. Widespread adoption of SERT will lead to a more consistent approach to power efficiency measurement across many current and growing markets in the next few years.

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

Hansfried Block, SERT Release Manager

SPECjbb2013 Released

November 11, 2012

SPEC announces the release of SPECjbb2013, developed from the ground up to measure performance based on the latest Java application features. It is expected to be used widely by all those interested in Java server performance, including JVM vendors, hardware developers, Java application developers, researchers and members of the academic community. New features include support for multiple run configurations, exercise of new Java 7 features, and support for virtualization and cloud environments.

<http://www.spec.org/jbb2013/>

www.spec.org

SPEC PTDaemon 1.6.0 Enhancement

March 23, 2013

March 2013 saw the release of PTDaemon V1.6.0, which includes several significant changes. Critical enhancements offer improved support for several accepted power analyzers, together with added options for non-standard baud rates and better handling of power analyzer auto-zero functions. Testing has raised concerns with the use of serial cables that do not have all pins connected leading to failures and missed readings; guidance has been added to the Power and

Temperature Measurement Setup Guide to enable better results to be achieved across the full range of accepted devices.

Support for the Yokogawa WT310 power analyzer and the iButtonLink T-Sense(R) and T-Probe temperature sensors (with the LinkUSBi 1-wire USB interface) have also been added. Three-phase power measurement is now fully supported, with appropriately configured and capable power analyzers, for benchmarks that support three-phase measurements.

This release also includes the first set of changes to provide future Direct-Current measurement support. These changes are still a work-in-progress. Therefore, the appropriate APIs have been changed (for DC only) to ensure that PTDaemon users cannot inadvertently use DC for benchmark results.

http://www.spec.org/power/docs/SPECpower-Device_List.html

Greg Darnell, SPEC PTDaemon Release Manager

SPEC Remembers Alan Adamson

November 11, 2012

The computer-performance benchmarking community mourns the loss of Alan Adamson, a member of the Standard Performance Evaluation Corporation (SPEC) Board of Directors. Alan died in his sleep the night of October 31, 2012.

Alan Adamson came to SPEC as an IBM representative specializing in compiler and Java performance. Within a short period of time, he became vitally involved in almost every aspect of SPEC work. After his retirement from IBM in 2009, he continued to work unpaid for a wide range of SPEC activities.

Alan served as chair of the Java and OSG Steering Committees. He helped establish, and did much of the hard work to create, the SPEC Research Group and the ICPE series of joint conferences with ACM.

At the time of his death, he served as the SPEC Board's inspector of elections, organizing the election process for the 2013 SPEC annual meeting. He was also a supporting contributor to the OSG Power Committee.

Alan was the closest thing to SPEC's official photographer, and a wise adviser for technical, procedural and moral issues. SPEC representatives from around the world will miss him as a friend and colleague. He is survived by his wife Melitta, his mother, and many siblings, nephews, nieces and cousins.

<http://www.spec.org/spec/memories/alan/alanadamson.html>

www.spec.org

Elections of Research Group's Steering Committee and Officers

April 12, 2013

In April 2013, the members of SPEC Research Group voted to elect the Steering Committee and officers of SPEC Research. The following individuals have been elected to the SPEC Research Group Steering Committee to serve until 2014: J. Nelson Amaral

(University of Alberta), Alexandru Iosup (TU Delft), Lizy John (University of Texas, Austin), Klaus-Dieter Lange (HP), Matthias Müller (RWTH Aachen University). Additionally, for two-year cadence, the following individuals were elected to serve until 2015: Samuel Kounev (KIT), 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: Samuel Kounev (KIT), Vice-Chair: Kai Sachs (SAP AG), Secretary: Klaus Lange (HP), Release Manager: Qais Noorshams (KIT).

Piotr Rygielski (KIT), David Schmidt (HP)

SPEC Releases New Server Virtualization Benchmark

May 22, 2013

The Standard Performance Evaluation Corp. (SPEC) has released a new version of its SPECvirt benchmark that reflects the heavy volume and sudden peaks of traffic faced by datacenter servers used for virtualized server consolidation.

The new SPECvirt_sc2013 uses real-world workloads and SPEC's proven performance- and power-measurement methodologies to enable vendors, users and researchers to compare system performance across multiple hardware, virtualization platforms, and applications. It was developed by the SPEC virtualization subcommittee, whose members and contributors include HP, IBM, Intel, Oracle, Red Hat, and VMware.

"Organizations around the world are virtualizing servers to handle significantly heavier and more dynamic traffic," says Paula Smith, chair of the SPEC virtualization subcommittee. "We are responding with workloads that model similar behavior, requiring virtual machines with more memory and virtual CPUs to meet the benchmark's quality-of-service requirements."

SPECvirt_sc2013 provides four workloads based on real-world traffic faced by a web server, a Java application server, an IMAP server and a batch server. A harness on the client side controls the workloads and implements the methodologies for power measurement derived from the SPECpower_ssj2008 benchmark.

The new SPECvirt benchmark better replicates the type of multi-tiered applications and database environments that run on enterprise-class virtualized servers. Each virtualized database server in SPECvirt_sc2013 is shared by as many as four application-server virtual machines.

SPECvirt_sc2013 shares the general benchmark architecture with SPECvirt_sc2010, including the benchmark harness and most of the application workloads and VM types. However, since SPECvirt_sc2010 was developed, many companies are virtualizing significantly heavier workloads than those in the first generation benchmark. In SPECvirt_sc2013, we are providing heavier and more bursty workloads that require virtual machines with more memory and vCPUs and further stress the system's ability to meet the benchmark's quality-of-service requirements. The list below summarizes some of the key changes made for

SPECvirt_sc2013:

- The J2EE Application workload based on SPECjAppServer2004 has been modified to further increase the burstiness of the workload so that the injection rate averages 100 IR vs 20 IR in SPECvirt_sc2010.
- The Database Server VM is now shared by up to four Appserver VMs (fewer if number of tiles is not a multiple of four) to ensure that there is a VM requiring multiple vCPUs (SMP).
- The Web Server workload based on the Support workload in SPECweb2005 has been modified to require SSL (HTTPS) between the clients and the web server, and the load level has been increased to 2500 Simultaneous Sessions (was 500).
- The Mail-Server workload, based on SPECmail2008, has been modified to present a bursty workload by increasing the variability of the think times in the benchmark.
- The Idle Server in SPECvirt_sc2010 has been replaced by the Batch Server that periodically runs a set of jobs that utilize the 401.bzip2 benchmark and harness from SPEC CPU2006. This replacement allows the jobs to be parallelized based on the number of vCPUs in the batch server VM. Like the idle server, this workload is not scored but must meet QoS requirements for completing each set of batch jobs within set limits as well as for being responsive to periodic client polling.

More information and initial benchmark results are available online.

http://www.spec.org/virt_sc2013/

David Schmidt (HP)

Over 160 Attendees at ICPE 2013

May 14, 2013

Over 160 participants from industry and academia attended the 4th ACM/SPEC International Conference on Performance Engineering (ICPE) 2013 in Prague, Czech Republic. The technical program included 20 full research papers, 2 short papers, 8 industrial/experience papers and 10 work-in-progress papers selected from over 80 submissions over all tracks. It also featured 3 keynote speakers of international repute, 6 poster/demo presentations, 4 tutorials and 1 invited talk.

The conference collocated two workshops, HotTopICS and MultiCloud. HotTopics attracted 54 participants with its focus on new and “hot” topics in the field of Cloud Services. MultiCloud had 47 participants and aimed to investigate concepts, models and tools for running applications on multiple and heterogeneous Cloud systems.

The next ICPE conference will take place in March 2014 in Dublin, Ireland.

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

Qais Noorshams (KIT)

WINNER OF THE SPEC DISTINGUISHED DISSERTATION AWARD 2012

January 24, 2013

The Research Group of the Standard Performance Evaluation Corp. (SPEC) selected the 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 for the second time. Nine nominations were submitted and reviewed by the selection committee. The chair of the nomination committee, Prof. Dr. Wilhelm Hasselbring from Kiel University in Germany, mentions that the committee had again a demanding task to select the winner among the nine high-quality nominations and that it is great to see so much exceptional work in scope of the SPEC Research Group. 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, multi-phase selection process the committee selected the following nomination: Shicong Meng who was nominated by Professor Ling Liu from the Georgia Institute of Technology with the topic “Monitoring-as-a-Service in the Cloud”. The winner received \$1000, which was awarded at the ICPE 2013 International Conference in Prague, Czech Republic.

Of the nine nominations, the selection committee decided for an honorable mention of the following three nominations, in alphabetical order:

- Karthik Ganesan who was nominated by Professor Lizy Kurian John from the University of Texas at Austin with the topic “Automatic Generation of Synthetic Workloads for Multicore Systems”. Mentioned for its contributions to the scope of the SPEC Research Group.
- Tilmann Rabl who was nominated by Karl Huppler from the Transaction Processing Performance Council with the topic “Efficiency in Cluster Database Systems”. Mentioned for its potential practical relevance.
- Daniel Sanchez who was nominated by Christos Kozyrakis from Stanford University with the topic “Hardware and Software Techniques for Scalable Thousand Core Systems”. Mentioned for its scientific originality.

Members of the selection committee for 2012 include: Wilhelm Hasselbring, Kiel University, Germany (Chair); David A. Bader, Georgia Institute of Technology, USA; Walter Bays, Oracle Corp., 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-Dieter Lange, HP, USA; John Murphy, University College Dublin, Ireland; Marco Vieira, University of Coimbra, Portugal.

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.

Wilhelm Hasselbring (Kiel University)



Karthik Ganesan receives the distinction from Samuel Kounev.



Tilmann Rabl receives the distinction from Samuel Kounev.

CALL FOR NOMINATIONS: SPEC DISTINGUISHED DISSERTATION AWARD 2013

May 7, 2013

The SPEC Distinguished Dissertation Award aims to recognize outstanding doctoral dissertations in the field of computer benchmarking, performance evaluation, and experimental system analysis in general. Nominated dissertations will be evaluated in terms of scientific originality, scientific significance, practical relevance, impact, and quality of the presentation.

The scope of the SPEC Research Group includes the aforementioned fields, 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 measurements, load testing, profiling, workload characterization, dependability, and efficiency evaluation of computing systems. The winner will receive \$1000, which will be awarded at the ICPE 2014.

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, 2012, and September 30, 2013. If there are several outstanding submissions, the committee may split the award between them. The submission deadline is September 30, 2013. 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 (members that confirmed participation at the time of publishing): Matthias Müller, RWTH Aachen University, Germany, (Chair); Samuel Kounev, Karlsruhe Institute of Technology, Germany; Walter Binder, University of Lugano, Switzerland; John Murphy, University, College Dublin, Ireland; Marco Viera, University of Coimbra, Portugal; Gerhard Wellein, Friedrich-Alexander-University, Erlangen-Nürnberg, Edgar Gabriel, University Houston, USA.

Matthias Müller (RWTH Aachen University)

FINCoS: A FRAMEWORK FOR PERFORMANCE EVALUATION OF EVENT PROCESSING SYSTEMS

April 3, 2013

FINCoS is a set of benchmarking tools for load generation and performance measurement of event processing (EP) systems. It provides a flexible and neutral approach through which users, researchers and engineers can quickly run realistic performance tests on one or more EP platforms without having to code themselves generation of the load, performance measurements and event conversion routines.

The framework can be used to evaluate any event-processing platform capable of exchanging events through the standard JMS API. In addition, an extensible set of custom-code adapters permits to submit loads directly to EP engines, using the client libraries provided by vendors.

FINCoS is not bound to any specific product or workload. Therefore, it can be used both by end-users willing to evaluate candidate EP platforms, and by researchers developing novel benchmarks for the area. For that, the framework offers great flexibility and a wide range of options in the definition of experimental evaluations. Users of FINCoS can choose to configure synthetic workloads, with events being generated by the framework itself, or, instead, use data files containing real-world event traces as input. The framework also allows the evaluation of the response of EP systems to changes. This is obtained by breaking performance tests into one or more phases, each with its own workload characteristics (i.e., event types, input rate, arrival process, etc.)

Load generation and result processing in FINCoS can be distributed across multiple nodes and centrally orchestrated using a graphical application, denominated FINCoS Controller. The results of performance runs can then be visualized, both in real-time and after test completion, using the FINCoS Performance Monitor component. The tool presents performance stats in both tabular and graphical formats—the former displays a snapshot of throughput and latency for each continuous query running at the EP system under test, while the latter shows the evolution of these metrics over time.

FINCoS development started in 2007, as part of the BiCEP benchmarking initiative, at University of Coimbra (<http://bicep.dei.uc.pt>). Its first version was released in 2008 and since then it has been considerably extended and improved. Recently, the framework has undergone a thorough review process, having been accepted to be integrated in the SPEC Research Group's repository of quantitative evaluation and analysis tools. FINCoS is an open-source tool, and can be downloaded free of charge. A user guide and a tutorial explaining how to use the framework are also available in the project web site.

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

Marcelo Mendes (University of Coimbra)

DISL: A DOMAIN-SPECIFIC LANGUAGE FOR INSTRUMENTATION-BASED DYNAMIC PROGRAM ANALYSIS ON THE JAVA PLATFORM

April 8, 2013

Dynamic program analysis (DPA) techniques enable essential software engineering tools, such as profilers, debuggers, and testing tools. DPA tools often rely on code instrumentation, i.e., the insertion of analysis code at selected locations in the analyzed program. As nowadays many programming languages are compiled to bytecode for the Java Virtual Machine (JVM), bytecode instrumentation has become a common technique for building DPA tools for the JVM.

JVM bytecode instrumentation is supported by a variety of bytecode engineering libraries, such as ASM, BCEL, Javassist, or ShrikeBT, to mention some of them. Generally, these tools offer a rather low level of abstraction, which results in high development effort for building DPA tools; the sources of such tools tend to be verbose and difficult to maintain and to extend. On the other hand, bytecode manipulation libraries give the expert developer many opportunities to optimize the inserted code, enabling efficient DPA tools.

Some researchers have proposed the use of aspect-oriented programming (AOP) for creating DPA tools. Thanks to its pointcut/advice mechanism, AOP offers a high abstraction level that allows to concisely specify certain kinds of instrumentations. Consequently, the use of AOP promises to significantly reduce the development effort for building DPA tools. Unfortunately, mainstream AOP languages such as AspectJ lack join points at the level of, e.g., basic blocks of code or individual bytecodes that would be needed for certain DPA tools. In addition, aspect weavers often introduce high overhead due to complex code transformations involved. For example, in AspectJ, the access to dynamic reflective join point information may involve object allocations and iterations (e.g., to store the arguments of a method in an array) that are not visible in the instrumentation code.

DiSL is a new domain-specific language for JVM bytecode instrumentation that reconciles the strengths of low-level bytecode manipulation libraries (i.e., expressiveness and efficiency) with the strength of AOP (i.e., high-level programming abstractions). DiSL allows programmers to rapidly specify custom DPA tools, which tend to be as concise as equivalent tools written in a mainstream AOP language (if this is possible). However, in contrast to prevailing AOP languages, DiSL features an open join point model where any region of bytecodes in a program can become a join point. Moreover, the DiSL weaver produces efficient code that does not incur any unexpected costs; the weaver introduces neither object allocations nor loops

that were not written by the programmer. This makes it especially attractive in performance evaluation tasks, where hidden overheads are particularly disruptive.

DiSL is available open-source under the Apache 2.0 license on OW2.

<http://disl.ow2.org/>

Lukas Marek (Charles University, Czech Republic), Yudi Zheng (University of Lugano, Switzerland), Danilo Ansaloni (University of Lugano, Switzerland), Aibek Sarimbekov (University of Lugano, Switzerland), Lubomir Bulej (Charles University, Czech Republic), Walter Binder (University of Lugano, Switzerland), Petr Tuma (Charles University, Czech Republic), Zhengwei Qi (Shanghai Jiao Tong University, China)

REPORT OF THE SPEC RG IDS BENCHMARKING WORKING GROUP

May 7, 2013

The SPEC RG IDS Benchmarking Working Group was created in July 2012 with the goal of contributing towards addressing the increasing demand for representative and rigorous benchmarks for intrusion detection systems (IDSes) in cloud platforms. Topics of interest include, but are not limited to, the following: security- and performance-related IDS benchmarking metrics, workloads based on vulnerability and attack injection, IDS benchmarking methodologies, IDS benchmarking methods with respect to features and properties of cloud platforms, and attack detection based on performance signatures. The elected officers of the Working Group are: Chair – Marco Vieira (University of Coimbra, Portugal); Vice-Chair – Alberto Avritzer (Siemens Corporate Research, USA); and Secretary/Release Manager – Aleksandar Milenkoski (Karlsruhe Institute of Technology, Germany).

Since its inception, the group has established a research agenda and has produced output that is published or is currently under consideration for publishing. The planned research activities of the group are described in a work-in-progress paper titled “Towards benchmarking intrusion detection systems for virtualized cloud environments” [1]. In addition, the group wrote a Technical Report [2] analyzing state-of-the-art intrusion detection techniques applied in virtualized environments and architectures of VMM (Virtual Machine Monitor)-based IDSes. The Technical Report also describes challenges that apply specifically to benchmarking VMM-based IDSes focusing on workloads and metrics. Moreover, the Working Group produced a survey paper systematizing the accumulated knowledge and current practices in the area of IDS benchmarking by using an IDS benchmarking design space structured into three parts: workload, metrics, and measurement methodology. The survey provides an overview of the current practices in IDS benchmarking by analyzing and comparing benchmarking approaches, tools, and methods related to each part of the design space.

Currently, the SPEC RG IDS Benchmarking Working Group is defining metrics that take into consideration on-demand resource provisioning in virtualized environments (e.g., CPU and memory hotplugging, memory ballooning). Further, vulnerability injection techniques

for generating malicious workloads consisting of attacks against hypervisors are being considered. The group is expanding its membership with new industrial and academic organizations.

[1] Aleksandar Milenkoski and Samuel Kounev. Towards Benchmarking Intrusion Detection Systems for Virtualized Cloud Environments (Extended Abstract). In Proceedings of the 7th International Conference for Internet Technology and Secured Transactions (ICITST 2012), London, United Kingdom, December 10-12, 2012, pages 562-563. Infonomics Society, IEEE. December 2012.

[2] Aleksandar Milenkoski, Samuel Kounev, Alberto Avritzer, Nuno Antunes, and Marco Vieira. On Benchmarking Intrusion Detection Systems for Virtualized (Cloud) Environments. Technical Report SPEC-RG-2013-002 v. 1.0, SPEC Research Group – IDS Benchmarking Working Group, June 2013. To appear.

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

ANNUAL REPORT OF THE CLOUD WORKING GROUP

May 10, 2013

The SPEC RG Cloud Working Group is pleased to report that in 2012 we have renewed our commitment to cloud benchmarking, quantitative evaluation, and experimental analysis, with relevance for both academia and industry. In just over 25 online meetings, we have gathered guests from 16 organizations, hosted 7 expert lectures on relevant topics, and started 2 content-related work activities. For the second part of 2012, 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.

Experience sharing was a key activity for the Cloud Working Group in 2012 and will continue to be for the upcoming years. Through online and face-to-face meetings, we have exchanged best-practices related to cloud benchmarking methodologies, performance evaluation tools and benchmarking suites, performance metrics, and performance evaluation results for commercial cloud services (e.g., the IaaS Amazon Web Services, the PaaS Google App Engine, the SaaS Dropbox, etc.) Among the lecturers, we have hosted representatives of the EPFL (CH), KIT (DE), and TU Delft (NL) universities; and SAP (DE and UK) and IBM (US). Our dedicated group of industry and academic experts includes now several of these representatives and more. The Cloud Working Group has also fostered collaboration between its members, with joint academic work, research visits, and company internships.

Over the course of 2012, the Cloud Working Group has been engaged in meaningful dialogue about how the performance of cloud-based systems can be measured. The dialogue has resulted, for example, in an innovative formalism for cloud usage patterns, which

will help academics and practitioners share information about IaaS, PaaS, and SaaS usage. Other topics for which preliminary work was started in 2012 include surveying existing cloud benchmarking metrics and tools, and (depending on future participation) formulating methodologies for benchmarking various cloud application domains (e.g., IaaS cloud services, IaaS/PaaS/SaaS Big Data processing services, and SaaS online gaming services).

In 2012, the Cloud Working Group has been represented by its members at many international academic and industrial venues, including through award-winning publications. 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, 2012 was a successful year for the Cloud Working Group. Our promise is to continue pursuing our mission and strengthen our contribution in 2013 and the years to come.

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

Alexandru Iosup (TU Delft)

NEW TECHNICAL REPORT OF THE RG CLOUD WG: CLOUD USAGE PATTERNS

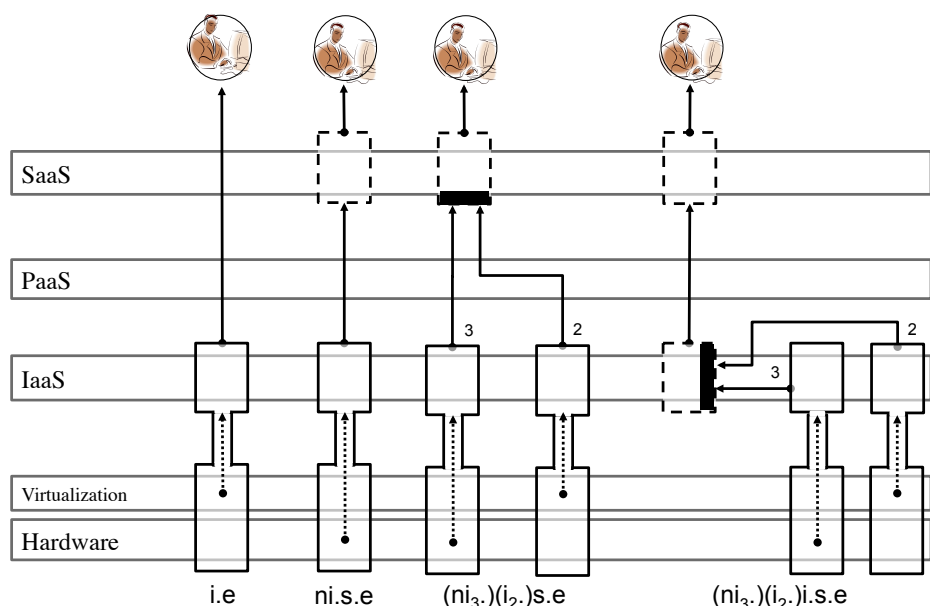
April 4, 2013

The SPEC Research Cloud Working Group announces the publication of the technical report "Cloud Usage Patterns: A Formalism for Description of Cloud Usage Scenarios", which is the result of collaboration between representatives of multiple research and industrial SPEC Research Group member organizations, such as Karlsruhe Institute of Technology (Germany), TU Delft (Netherlands), SAP AG (Germany), and Salesforce.com (USA). Cloud computing is becoming an increasingly lucrative branch of the existing information and communication technologies. Enabling a debate about cloud usage scenarios can help with attracting new customers, sharing best-practices, and designing new cloud services. In contrast to previous approaches, which have attempted mainly to formalize the common service delivery models (i.e., Infrastructure-as-a-Service, Platform-as-a-Service, and Software-as-a-Service), in the technical report, the authors propose a formalism for describing common cloud usage scenarios referred to as cloud usage patterns. The proposed formalism takes a structuralist approach allowing decomposition of a cloud usage scenario into elements corresponding to the

common cloud service delivery models. Furthermore, the formalism considers several cloud usage patterns that have recently emerged, such as hybrid services and value chains in which mediators are involved. We propose a simple yet expressive textual and visual language for our formalism, and we show how it can be used in practice for describing a variety of real world cloud usage scenarios. The scenarios for which we demonstrate our formalism include resource provisioning of global providers of infrastructure and/or platform resources (e.g., Amazon Web Services, Google App Engine), online social networking services (e.g., Facebook), user-data processing services, online customer and ticketing services, online asset management and banking applications, CRM (Customer Relationship Management) applications, and online social gaming applications (e.g., Zynga's applications). For instance, the use of leased infrastructure resources by Go!Animate from the infrastructure provider Amazon Web Services for servicing users processing amateur videos, can be formally described with the string "i₃.s.e₆", where each letter relates to one of the known cloud service delivery models (i.e., Infrastructure-as-a-Service and Software-as-a-Service in the given string), and/or a stakeholder (i.e., an end-user). Also, each number relates to a volume of provisioned resources that may be expressed, for example, as a logarithmic value (e.g., 3 for 1000 in the given string). For further details on the cloud usage patterns, we refer the reader to the technical report.

http://research.spec.org/fileadmin/user_upload/documents/rg_cloud/endorsed_publications/SPEC-RG-2013-001_CloudUsagePatterns.pdf

Aleksandar Milenkoski (KIT), Alexandru Iosup (TU Delft)



Visual form of several cloud usage patterns.

REPORT FROM THE 4TH ACM/SPEC INTERNATIONAL CONFERENCE ON PERFORMANCE ENGINEERING (ICPE 2013) AND COLLOCATED EVENTS

May 7, 2013

Introduction

Over 160 participants from industry and academia coming from 25 different countries attended the 4th ACM/SPEC International Conference on Performance Engineering (ICPE) 2013 in Prague, Czech Republic. The technical program included 20 full research papers, 2 short papers, 8 industrial/experience papers and 10 work-in-progress papers selected from over 80 submissions over all tracks. It also featured three keynote speakers of international repute, six poster/demo presentations, four tutorials and one invited talk on SPECjbb2013.

The conference collocated two workshops, HotTopICS and MultiCloud. HotTopics attracted 54 participants with its focus on new and “hot” topics in the field of Cloud Services. MultiCloud had 47 participants and aimed to investigate concepts, models and tools for running applications on multiple and heterogeneous Cloud systems.

Collocated Workshops

This year’s ICPE was accompanied with two collocated workshops. On Saturday, April 20, the HotTopICS Workshop started with the keynote given by Alexandru Iosup (TU Delft). The topic of the talk concerned approaches, challenges and experiences with IaaS Cloud benchmarking. On Sunday, Samuel Kounev (KIT) gave the invited talk entitled “RELATE: A Research Training Network on Engineering and Provisioning of Service-Based Cloud Applications”. During the workshop’s sessions, the participants had a unique occasion to follow ten presentations about the newest trends and ideas in the Cloud Computing domain. The workshop ended with a joint lunch on Sunday.

On Monday morning, the second collocated workshop started: “MultiCloud”. This one-day event gathered researchers and companies interested in issues, problem and ideas concerning the following topics: Multi-Cloud background and general issues, development of Multi-Cloud applications, monitoring and execution support, Cloud migration. The workshop was concluded with the discussion panel and a joint coffee at the roof of the Dorint Hotel.

The Conference

The main conference, 4th ACM/SPEC International Conference on Performance Engineering (ICPE 2013), started with the Tutorials track on Sunday, April 21. The main tracks of the conference started on Monday morning with the keynote speech “Automated Modeling of Complex Data Center Applications” given by Prashant Shenoy (University of Massachusetts). After the keynote, all participants listened to the talks given by candidates to the Research Track Best Paper Award. The sessions that followed were a perfect occasion

to see the highest ranked papers of the conference. During the coffee breaks between the sessions, the participants were discussing their ideas and problems.

The second day of the conference was opened by the second keynote, entitled “Supporting Operations Personnel through Performance Engineering”, given by Len Bass from National Information and Communications Technology Australia (NICTA). The Tuesday’s research track sessions attracted the attention of listeners until late afternoon, when the Vision and Work in Progress session started. During this session, researchers presented ten interesting topics that were selected for publication, although the research is still in progress. We hope that the interesting ideas presented during this session will be a great motivation to submit follow-up results to the next ICPE in Dublin.

The main conference ended on Wednesday, April 24, when the last keynote talk was given. The talk was entitled “Data Centric Computing for Internet Scale Enterprises” and presented by Yuqing Gao from IBM T.J. Watson Research Center. Before the conference was closed, the participants listened to the presentations of the candidates to the Industry Track Best Paper Award.

SPEC Face-to-Face Meetings

Each ICPE is a perfect occasion for face-to-face meetings of researchers and industry representatives that work in the domain of performance engineering. The members of SPEC also use this opportunity to meet in person and discuss the topics of their interests. During this year’s ICPE, SPEC organized five face-to-face meetings: SPEC OSG CPU, SPEC OSG Java Enterprise, SPEC OSG Power, SPEC Research and the meetings of the Steering Committees. The reports from the face-to-face meeting within SPEC Research are included in this issue of the newsletter (read more on page 11). We hope, they will give the readers an insight of selected discussion topics and be an invitation to join SPEC Research or one of its working groups.

Piotr Rygielski (KIT), Qais Noorshams (KIT), Petr Tuma (Charles University of Prague)

ICPE 2013: BEST PAPER AWARDS

May 14, 2013

At the 2013 International Conference on Performance Engineering (ICPE 2013) in Prague, Czech Republic, two outstanding papers were distinguished by the Program and Awards Chairs.

The Best Research Track Paper Award from research track was presented to Balaji Subramaniam and Wu-Chun Feng for the paper “Towards Energy-Proportional Computing for Enterprise-Class Server Workloads”.

The nominees to the Best Paper Award in the research track were the following two papers: Christoph Heger, Jens Happe, Roozbeh Farahbod for the paper “Automated Root Cause Isolation of Performance Regressions Variaduring Software Development”, and Jean-Michel Fourneau, Katinka Wolter, Philipp

Reinecke, Tilman Krauss for the paper “Multiple Class G-Networks with Restart”.

The distinction—Best Industry Track Paper Award—was presented to Klaus-Dieter Lange, Jeremy A. Arnold, Hansfried Block, Nathan Totura, John Beckett, Mike G. Tricker for the paper “Further Implementation Aspects of the Server Efficiency Rating Tool (SERT)”.

The other nominee to the Best Paper Award in the Industry/Experience Track was the paper by Tilmann Rabl, Meikel Poess, Hans-Arno Jacobsen, Pat O’Neil, Betty O’Neil “Variations of the Star Schema Benchmark to Test Data Skew in Database Management Systems”

Congratulations to the award winners and the authors of distinguished papers!

Piotr Rygielski (KIT)

SHORT REPORTS FROM FACE-TO-FACE MEETINGS

May 14, 2013

The SPEC Research Annual Meeting

The members of the SPEC Research Group met together on May 24th, 2013. The meeting was chaired by Samuel Kounev who presented the group and its working groups highlighting many successes achieved in the previous year. After this introduction, the attendees introduced themselves, presented their scientific interests and involvement in the groups. During the meeting, the chairman highlighted that the SPEC Research Group is open for everyone, has diverse areas of research focus and its main aim is to foster the collaboration between industry and academia. Moreover, the benefits of joining the group have been presented.

Later, the chair of the IDS Working Group, Marco Vieira, shared his experiences about the process of establishing a new working group within SPEC Research. The meeting was concluded by an introduction of the two working groups—Cloud and IDS. After a short break, the meeting of the Cloud Working Group followed.

Piotr Rygielski (KIT), Meikel Poess (Oracle)

The Cloud Working Group

The SPEC RG Cloud Working Group met face-to-face in Prague on May 24th, 2013. The meeting opened with the introduction of the chair Alexandru Iosup. He introduced topics of interest of the group and announced to the attendees that the Cloud Working Group looks for new members (read more on page 3).

After the introductory part, Aleksandar Milenkoski presented the first Technical Report of the Cloud Working Group entitled “Cloud Usage Patterns: A Formalism for Description of Cloud Usage Scenarios” (see page 9). The presentation was concluded by a discussion that showed that the document is important for industrial partners. Moreover, many ideas of extending the formalism were proposed.

In the next part of the meeting, Nikolas Herbst presented material about an upcoming technical report that focuses on cloud performance benchmarking. The

presentation gave insight to the initial progress in the preparation of this document, especially on the elasticity benchmarking. The presentation interested many attendees and they actively discussed on the presented topics. The meeting was closed by Alexandru Iosup and the attendees headed to the Prague castle for a joint dinner where discussions were continued.

Piotr Rygielski(KIT)

The IDS Working Group

A face-to-face meeting of the SPEC RG IDS Benchmarking Working Group took place in Prague, on May 25th, 2013. This was the first face-to-face meeting of the group, which periodically meets online. All members of the working group were represented in the meeting, which was also attended by representatives of other SPEG RG working groups.

During this meeting, we first discussed management aspects of the group and approaches for attracting new members. Afterwards, we focused on technical aspects, discussing requirements and approaches for the generation of malicious workloads, including the use of vulnerability and attack injection techniques.

Marco Vieira (University of Coimbra)

SPECTACULAR AWARDS FOR OUTSTANDING SERVICE TO SPEC IN 2012

May 29, 2013

During SPEC’s 2012 Quarterly Meeting in Prague (Czech Republic), RG Officers Alexandru Iosup, Qais Noorshams and Aleksandar Milenkoski were recognized by the SPEC President and the SPEC RG Chair for their outstanding service to SPEC in 2012.

Alexandru Iosup (TU Delft) received an award for his outstanding service as elected Chair driving the development and establishment of the SPEC RG Cloud Working Group. Qais Noorshams (KIT) received an award for his outstanding service as elected SPEC RG Release Manager. Aleksandar Milenkoski (KIT) received an award for his outstanding service as elected Secretary of the RG Cloud and RG IDS working groups, among other things leading the work on the first official technical report of SPEC RG. In addition, Samuel Kounev (KIT) was recognized by SPEC President for his continued service driving the further establishment and development of the SPEC Research Group and the ICPE conference.

Congratulations to all award recipients!

Piotr Rygielski (KIT)



Samuel Kounev (left) receives the SPECTacular Award from SPEC President Walter Bays.



Aleksandar Milenkoski receives the SPECTacular Award from Walter Bays and Samuel Kounev.



Samuel Kounev and Walter Bays present the Award to Alexandru Iosup.



Walter Bays presents the Award to Qais Noorshams.

ICPE 2013: ABSTRACTS OF SELECTED PAPERS

May 14, 2013

Balaji Subramaniam, Wu-chun Feng: Towards Energy-Proportional Computing for Enterprise-Class Server Workloads (Research Track Best Paper Award)

Massive data centers housing thousands of computing nodes have become commonplace in enterprise computing, and the power consumption of such data centers is growing at an unprecedented rate. Adding to the problem is the inability of the servers to exhibit energy proportionality, i.e., provide energy-efficient execution under all levels of utilization, which diminishes the overall energy efficiency of the data center. It is imperative that we realize effective strategies to control the power consumption of the server and improve the energy efficiency of data centers. With the advent of Intel Sandy Bridge processors, we have the ability to specify a limit on power consumption during runtime, which creates opportunities to design new power-management techniques for enterprise workloads and make the systems that they run on more energy proportional.

In this paper, we investigate whether it is possible to achieve energy proportionality for an enterprise-class server workload, namely SPECpower ssj2008 benchmark, by using Intel's Running Average Power Limit (RAPL) interfaces. First, we analyze the power consumption and characterize the instantaneous power profile of the SPECpower benchmark within different subsystems using the on-chip energy meters exposed via the RAPL interfaces. We then analyze the impact of RAPL power limiting on the performance, per-transaction response time, power consumption, and energy efficiency of the benchmark under different load levels. Our observations and results shed light on the efficacy of the RAPL interfaces and provide guidance for designing power-management techniques for enterprise-class workloads.

Christoph Heger, Jens Happe, Roozbeh Farahbod:

Automated Root Cause Isolation of Performance Regressions During Software Development (Research Track Best Paper Nomination)

Performance is crucial for the success of an application. To build responsive and cost efficient applications, software engineers must be able to detect and fix performance problems early in the development process.

Existing approaches are either relying on a high level of abstraction such that critical problems cannot be detected or require high manual effort. In this paper, we present a novel approach that integrates performance regression root cause analysis into the existing development infrastructure using performance-aware unit tests and the revision history. Our approach is easy to use and provides software engineers immediate insights with automated root cause analysis. In a realistic case study based on the change history of Apache Commons Math, we demonstrate that our approach can automatically detect and identify the root cause of a major performance regression.

**Jean-Michel Fourneau, Katinka Wolter, Philipp Reinecke, Tilman Krauss, Alexandra Danilkina :
Multiple Class G-Networks with Restart
(Research Track Best Paper Nomination)**

Restart is a common technique for improving response-times in complex systems where the causes of delays can either not be discerned, or not be addressed by the user. With restart, the user aborts a running job that exceeds a deadline, and resubmits it to the system immediately. In many common scenarios, this approach can reduce the response-times that the user experiences. Restart has been well-studied for scenarios where only one user applies restart, and typically in cases where queueing effects can be neglected. In this paper we approach the question of restart in a scenario where restart is applied by many users in a system that can be modelled as an open queueing network. We apply the G-Networks formalism to this problem. We use negative customers to model the abortion and retry of a request. The open G-network uses multiple classes with phase-type distributed service times. This allows the approximation of a preemptive repeat different behaviour as it is natural for multiple restarts of a request. We compute the response time of a request and show that an optimal restart interval can be found. The results are compared with simulation.

**Klaus-Dieter Lange, Jeremy A. Arnold, Hansfried Block, Nathan Totura, John Beckett, Mike G. Tricker: Further Implementation Aspects of the Server Efficiency Rating Tool (SERT)
(Industry Track Best Paper Award)**

The Server Efficiency Rating Tool (SERT) has been developed by the Standard Performance Evaluation Corporation (SPEC) at the request of the US Environmental Protection Agency (EPA). Almost 3% of all electricity consumed within the US in 2010 went to running datacenters. With this in mind, the EPA released Version 2.0 of the ENERGY STAR for Computer Servers program in early 2013 to include the mandatory use of the SERT. Other governments world-wide that are also concerned with growing power consumption of servers and datacenters are considering the adoption of the SERT.

**Tilmann Rabl, Meikel Poess, Hans-Arno Jacobsen, Patrick O'Neil, Elizabeth O'Neil: Variations of the Star Schema Benchmark to Test the Effects of Data Skew on Query Performance
(Industry Track Best Paper Nomination)**

The Star Schema Benchmark (SSB), has been widely used to evaluate the performance of database management systems when executing star schema queries. SSB, based on the well known industry standard benchmark TPC-H, shares some of its drawbacks, most notably, its uniform data distributions. Today's systems rely heavily on sophisticated cost-based query optimizers to generate the most efficient query execution plans. A benchmark that evaluates optimizer's capability to generate optimal execution plans under all circumstances must provide the rich data set details on which optimizers rely (uniform and non-uniform distributions, data sparsity, etc.). This is also true for other database system parts, such as indices and operators, and ultimately holds for an end-to-end benchmark as well. SSB's data generator, based on TPC-H's dbgen, is not easy to adapt to different data distributions as its meta data and actual data generation implementations are not separated. In this paper, we motivate the need for a new revision of SSB that includes non-uniform data distributions. We list what specific modifications are required to SSB to implement non-uniform data sets and we demonstrate how to implement these modifications in the Parallel Data Generator Framework to generate both the data and query sets.



5th ACM/SPEC International Conference on Performance Engineering



ICPE 2014

Dublin, Ireland - March 23-26



<http://icpe2014.ipd.kit.edu/>

CALL FOR PAPERS

A Joint Meeting of WOSP/SIPEW sponsored by ACM SIGMETRICS and ACM SIGSOFT in cooperation with SPEC. 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:

- Performance and software development processes
- Performance modeling and prediction
- Performance measurement and experimental analysis
- Benchmarking, configuration, sizing and capacity planning
- System management / optimization
- Performance in cloud, virtualized and multi-core systems
- Performance and Power
- Performance modeling and evaluation in other domains

For more details, please visit: http://icpe2014.ipd.kit.edu/call_for_papers/

RESEARCH PAPERS – PROGRAM COMMITTEE

- Jose Nelson Amaral, University of Alberta, Canada
- Alberto Avritzer, Siemens Corporate Research, USA
- Steffen Becker, University of Paderborn, Germany
- Simona Bernardi, Centro Universitario de la Defensa, Spain
- Steve Blackburn, Australian National University, Australia
- Andre Bondi, Siemens Corporate Research, USA
- Edson Borin, Universidade de Campinas, Brazil
- Thouraya Bouabana-Tebibel, ESI, Algeria
- Radu Calinescu, University of York, UK
- Lydia Chen, IBM Research Zurich Lab, Switzerland
- Vittorio Cortellessa, Università di L'Aquila, Italy
- Tony Field, Imperial College London, UK
- Vincenzo Grassi, University of Roma "Tor Vergata", Italy
- Lars Grunske, TU Kaiserslautern, Germany
- Wilhelm Hasselbring, University of Kiel, Germany
- Matthias Hauswirth, University of Lugano, Switzerland
- Alexandru Iosup, TU Delft, Netherlands
- Carlos Juiz, University of the Balearic Islands, Spain
- Jorge Julvez, Universidad de Zaragoza, Spain
- David Kaeli, Northeastern University, USA
- Samuel Kounev, Karlsruhe Institute of Technology, Germany
- Anne Koziolok, University of Zurich, Switzerland
- Heiko Koziolok, ABB Corporate Research, Germany
- Diwakar Krishnamurthy, University of Calgary, Canada
- Doug Lea, SUNY Oswego, USA
- Yan Liu, Concordia University, Canada
- Stefano Marrone, Seconda Università di Napoli, Italy
- Daniel Menasce, George Mason University, USA
- Raffaella Mirandola, Politecnico di Milano, Italy
- David Pearce, Victoria University of Wellington, New Zealand
- Diego Perez-Palacin, Politecnico di Milano, Italy
- Dorina Petriu, Carleton University, Canada
- Jerry Rolia, HP Labs, Canada
- Giuseppe Serazzi, Politecnico di Milano, Italy
- Catia Trubiani, Università di L'Aquila, Italy
- Petr Tuma, Charles University in Prague, Czech Republic
- Murray Woodside, Carleton University, Canada
- Peng Wu, IBM T. J. Watson Research Center, USA
- Thomas Würthinger, Oracle Labs, Austria
- Xiaoyun Zhu, VMware, USA

SUBMISSION GUIDELINES

Authors are invited to submit original, unpublished papers that are not being 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: <http://icpe2014.ipd.kit.edu/> 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 2014 conference proceedings that will be published by ACM and included in the ACM Digital Library. For submission instructions, please visit: <http://icpe2014.ipd.kit.edu/>

IMPORTANT SUBMISSION DATES

Research Papers	27 Sep 2013
Industrial / Experience Papers	25 Oct 2013
Tutorial Proposals	15 Nov 2013
Poster and Demo Papers	13 Jan 2014
Work-in-Progress and Vision Papers	13 Jan 2014

ORGANIZING COMMITTEE

General Chairs

- Klaus-Dieter Lange, Hewlett-Packard, USA
- John Murphy, University College Dublin, Ireland

Program Chairs

- Walter Binder, University of Lugano, Switzerland
- José Merseguer, Universidad de Zaragoza, Spain

Industrial Chair

- Raghunath Nambiar, Cisco Systems, USA

Tutorial Chair

- Alexandru Iosup, TU Delft, Netherlands

Demos and Posters Chairs

- Kirk W. Cameron, Virginia Tech, USA
- Anthony Ventresque, University College Dublin, Ireland

Publication Chair

- David Gregg, Trinity College Dublin, Ireland

Finance Chair

- Anja Bog, SAP, USA

Publicity Chairs

- Danilo Ansaloni, University of Lugano, Switzerland
- Bob Cramblitt, Cramblitt & Company, USA

Registration Chair

- Nicola Stokes, University College Dublin, Ireland

Awards Chairs

- Samuel Kounev, Karlsruhe Institute of Technology, Germany
- Meikel Poess, Oracle, USA

Local Organization Chair

- Patrick McDonagh, Dublin City University, Ireland

Web Chair

- Cathy Sandifer, SPEC, USA

INDUSTRY PAPERS – PROGRAM COMMITTEE

- Jeremy A. Arnold, IBM, USA
- Prasenjit Biswas, ARM, USA
- Jeffrey J. Blessing, Milwaukee School of Engineering, USA
- Hansfried Block, Fujitsu, Germany
- Andrew Bond, Red Hat, USA
- Wen Chen, University of North Dakota, USA
- Mathew Colgrove, The Portland Group, USA
- Malte Dreyer, Humboldt University, Germany
- Debo Dutta, Cisco Systems, USA
- Diwakar Krishnamurthy, University of Calgary, Canada
- Qais Noorshams, Karlsruhe Institute of Technology, Germany
- Kamesh Raghavendra, NetApp, USA
- Lisa Roderick, VMware, USA
- David Schmidt, Hewlett-Packard, USA
- Sanjay Sharma, Intel, USA
- Cloyce Spradling, Oracle, USA
- Mike Tricker, Microsoft, USA