



SPEC ResearchSM Group Newsletter

CONTENTS

- 2 SPEC RG Officers and Working Groups
- 3 Welcome to the SPEC RG Newsletter
- 3 Kaivalya Dixit Distinguished Dissertation Award
- 3 SPEC launches new blog section
- 3 ICPE Conference included in CORE 2021 Ranking
- 4 ICPE 2022: Statistics
- 4 ICPE 2023 in Coimbra
- 5 Reports of the Working Groups
- 10 Dagstuhl Seminar 21201 "Serverless Computing"
- 11 HotCloudPerf Workshop 2022
- 12 Selected Abstracts

SPEC KAIVALYA DIXIT DISTINGUISHED DISSERTATION AWARD 2021

The award selection committee for 2021 chose the dissertation of André Bauer of University of Würzburg. The selection committee was impressed by the comprehensive coverage of this thesis including design, benchmark platform, and use cases for hybrid forecasting of time series.

Read more on page 3

ICPE CONFERENCE INCLUDED IN CORE 2021 RANKING

The ACM/SPEC International Conference of Performance Engineering has been included in the CORE 2021 conference ranking with rank B. This decision was made by an academic committee based on objective data requested as part of the submission process.

Read more on page 3

ICPE 2023 WILL BE HELD IN COIMBRA

The ICPE organizing committee and general chair Marco Vieira invite interesting high-quality submissions for the next ACM/SPEC International Conference on Performance Engineering (ICPE 2023). The conference will take place on April 17-21, 2023 in Coimbra, Portugal.

Read more on page 4



CONTACT

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

SPEC Research Group

Chair: Samuel Kounev (rgchair@spec.org)

Web: <https://research.spec.org>

SPEC RESEARCH WORKING GROUPS REPORT ON THEIR PROGRESS

The SPEC Research Working Groups Security, Cloud, DevOps Performance, Predictive Data Analytics, and Power report on their progress, articles, benchmarks, and technical reports published in 2021. The Working Groups are always open for new members. Feel invited to join us!

Read more on pages 5-10

SPEC RESEARCH GROUP OFFICERS

Chair:

Samuel Kounev, University of Würzburg, Germany

Vice-Chair:

André van Hoorn, University of Hamburg, Germany

Secretary:

Cristina Abad, ESPOL University, Ecuador

Steering Committee:

Cristina Abad, ESPOL University, Ecuador

Cor-Paul Bezemer, University of Alberta, Canada

Nikolas Herbst, University of Würzburg, Germany

Alexandru Iosup, VU Amsterdam, NL

Samuel Kounev, University of Würzburg, Germany

Klaus-Dieter Lange, HPE, USA

Alex Podelko, MongoDB, USA

Weiyi Shang, Concordia University, Canada

Petr Tůma, Charles Univ. of Prague, Czech Republic

André van Hoorn, University of Hamburg, Germany

Publicity Officer:

Weiyi Shang, Concordia University, Canada

Release Manager:

Michele Tucci, Charles Univ. Prague, Czech Republic

Newsletter Editors:

Martin Straesser, University of Würzburg, Germany

André Bauer, University of Würzburg, Germany

<https://research.spec.org>

SPEC RESEARCH WORKING GROUPS

DevOps Performance Working Group

Chair:

André van Hoorn, University of Hamburg, Germany

Vice-Chair:

Cor-Paul Bezemer, University of Alberta, Canada

Weiyi Shang, Concordia University, Canada

Secretary:

Simon Eismann, University of Würzburg, Germany

<https://research.spec.org/en/working-groups/rg-devops-performance.html>

Cloud Working Group

Chair:

Alexandru Iosup, VU Amsterdam, The Netherlands

Vice-Chair:

Nikolas Herbst, University of Würzburg, Germany

Secretary:

Sacheendra Talluri, VU Amsterdam, The Netherlands

Release Manager:

André Bauer, University of Würzburg, Germany

<https://research.spec.org/en/working-groups/rg-cloud.html>

Power Working Group

Chair:

Maximilian Meissner, University of Würzburg, Germany

Vice-Chair:

Klaus-Dieter Lange, HPE, USA

Secretary:

Aaron Cragin, Microsoft, USA

<https://research.spec.org/working-groups/rg-power.html>

Security Working Group

Chair:

Aleksandar Milenkoski, Cybereason/Univ. Würzburg, Germany

Vice-Chair:

Nuno Antunes, University of Coimbra, Portugal

Secretary and Release Manager:

José Flora, University of Coimbra, Portugal

<https://research.spec.org/working-groups/rg-ids-benchmarking.html>

Predictive Data Analytics Working Group

Chair:

André Bauer, University of Würzburg, Germany

Vice-Chair:

Pooyan Jamshidi, University of South Carolina, USA

Secretary:

Mark Leznik, University Ulm, Germany

Release Manager:

Daniel Seybold, University Ulm, Germany

<https://research.spec.org/working-groups/rg-predictive-data-analytics/>

WELCOME TO THE SPEC RESEARCH GROUP NEWSLETTER

With over 120 members in 30 countries, and nearly two dozen benchmarks spanning highly diverse aspects of computing performance and energy efficiency, SPEC has become known as a beacon of truth for computing researchers, vendors, users and analysts worldwide. These professionals rely on SPEC to ensure that the marketplace has a fair and useful set of metrics to differentiate computing systems. Founded in 2011, the SPEC Research Group is proud of being part of this remarkable history.

We are delighted to present to you the next 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.

Some highlights from the last year include:

- 12th (and second virtual) ACM/SPEC **ICPE 2021** in Rennes, France
- 2nd IEEE International Conference on Autonomic Computing and Self-Organizing Systems **ACSOS 2021** in Washington D.C., US
- 4th Workshop on Hot Topics in Cloud Computing Performance **HotCloudPerf 2021** at ICPE 2021
- A new working group within the SPEC RG: The *Predictive Data Analytics* Working Group

We have been actively working on preparation, planning, and organization of ICPE 2022, this year again (and hopefully the last time) as a virtual conference. We hope that a vivid exchange of ideas will be a great motivation for the next year of scientific and engineering work.

We hope that you will enjoy reading the newsletter. We welcome and encourage your contributions for articles and suggestions for future coverage.

Samuel Kounev
(SPEC Research Chair, University of Würzburg).
Martin Straesser
(Newsletter Editor, University of Würzburg).
André Bauer
(Newsletter Editor, University of Würzburg).

SPEC, the SPEC logo and the names SERT, SPEC SFS, SPECjbb, SPECvirt, Chauffeur, and PTDaemon are registered 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 ©1988-2022 Standard Performance Evaluation Corporation (SPEC). Reprinted with permission. All rights reserved.

SPEC KAIVALYA DIXIT DISTINGUISHED DISSERTATION AWARD 2021 WINNER

The SPEC Kaivalya Dixit Distinguished Dissertation Award is an annual award that 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 presentation. The award selection committee for 2021 was chaired by Prof. John Murphy of University College Dublin, Ireland and consisted of seven members from academia.

The **winning dissertation** “Automated Hybrid Time Series Forecasting: Design, Benchmarking, and Use Cases” was authored by André Bauer from University of Würzburg, under the supervision of Professor Samuel Kounev. The selection committee was impressed with the importance of hybrid forecasting of time series and the application to autoscaling along with the comprehensive coverage including design, benchmark platform, and use cases.

The award is to be presented at the International Conference on Performance Engineering (ICPE 2022).

Given the high quality of dissertations nominated for this award, the committee decided to publicly recognize another dissertation as Runner-Up, “Enabling High-Performance Large-Scale Irregular Computations” authored by Dr. Maciej Besta, of ETH Zürich, under the supervision of Dr. Torsten Hoefler.

<https://research.spec.org/news/2022-03-15-00-00-winner-of-spec-kaivalya-dixit-distinguished-dissertation-award-2021/>

SPEC LAUNCHES NEW BLOG SECTION

SPEC (Standard Performance Evaluation Corporation) has launched a new blog on its website [1]. The blog features the latest developments and updates from all SPEC groups and members. Members of the Research Group are invited and encouraged to submit content proposals. The RG steering committee is in talks about a formal process outlining how contents can be proposed, and will inform RG members about the outcome in the near future.

Samuel Kounev, SPEC Research Chair
André van Hoorn, SPEC Research Vice-Chair

[1] <https://www.spec.org/blog/>

ICPE CONFERENCE INCLUDED IN CORE 2021 RANKING

The ACM/SPEC International Conference on Performance Engineering was included in the CORE 2021 conference ranking with rank B. The CORE Conference Ranking provides assessments of major conferences in the comput-

ing disciplines. The rankings are managed by the CORE Executive Committee, with periodic rounds for submission of requests for addition or reranking of conferences. Decisions are made by academic committees based on objective data requested as part of the submission process. Conference rankings are determined by a mix of indicators, including citation rates, paper submission and acceptance rates, and the visibility and research track record of the key people hosting the conference and managing its technical program.

The ICPE Steering Committee

ICPE 2022: STATISTICS

The 13th ACM/SPEC International Conference on Performance Engineering (ICPE 2022) will be held virtually from April 9 to April 13. This year the research track of ICPE attracted 38 submissions, 9 of which were selected as full articles after a rigorous review process, yielding an acceptance ratio of 23%. There were also 4 articles accepted as short articles. Of the 20 submissions to the Industry Track, 9 were accepted, of which 5 as full articles. 4 papers of the research and industry tracks were able to achieve ACM artifact badges. The Work-in-Progress and Vision Track received 6 submissions of which 2 were accepted. For the first time, ICPE features a Data Challenge with four accepted short papers. 6 tutorials have been proposed out of which 3 were selected to complement the program. Furthermore, 7 demonstrations and posters are scheduled for an interactive session.

The following workshops are planned for ICPE 2022:

- The Tenth International Workshop on Load Testing and Benchmarking of Software Systems (LTB 2022):
<https://ltb2022.eecs.yorku.ca/>
- The Fifth Workshop on Hot Topics in Cloud Computing Performance (HotCloudPerf-2022):
<https://hotcloudperf.spec.org/>
- The Second Workshop on Performance and Energy Efficiency in Concurrent Systems (PECS 2022):
<https://pecs-workshop.github.io/2022/>
- The First International Workshop on Performance-Data Analytics and Data-Management (PANDA 2022):
<https://panda-workshop.github.io/>
- The First International Workshop on Information Technology Storage Management Specification (ITSMS 2022):
<https://itsms-22.github.io/>

- The Seventh Workshop on Challenges in Performance Methods for Software Development (WOSP-C 2022):

<https://wosp-c.github.io/wosp-c-22/>

The following distinguished speakers will give keynotes at ICPE 2022:

- Ivona Brandic, Professor for High Performance Computing Systems, TU Wien
Data Science Driven Methods for Sustainable and Failure Tolerant Edge Systems
- John Wilkes, Google
Building warehouse-scale computers
- Longxiang Li, Inspur
Performance Optimization of HPC Applications in Large-Scale Cluster Systems

With the pandemic, the original planned face-to-face ICPE in Beijing is not possible. All authors were offered the opportunity to provide video presentations and slides to the community complemented by focused discussion threads. Virtual sessions via teleconference are scheduled for the days of the conference for the invited speakers' presentations and for live paper presentations and discussions. The detailed program can be found at [the ICPE 2022 website](#).

Nikolas Herbst (University of Würzburg, Germany),
Philipp Leitner (Chalmers — University of Gothenburg, Sweden),
Steffen Becker (University of Stuttgart, Germany),
Dan Feng (Huazhong University of Science and Technology, China)

ICPE 2023 IN COIMBRA, PORTUGAL - PRELIMINARY ANNOUNCEMENT

The ACM/SPEC International Conference on Performance Engineering (ICPE) provides a forum for the integration of theory and practice in the field of performance engineering. It brings together researchers and industry practitioners to share ideas, discuss challenges, and present results of both work-in-progress and state-of-the-art research on performance engineering of software and systems.

ICPE 2023 will be held in Coimbra, Portugal, from April 17 to April 21. Coimbra is a vibrant city in the center of Portugal built around the University of Coimbra, a higher education institution created in 1290, that embraces research and education units in a broad set of domains. Coimbra is the home of several high-tech companies that ship software products worldwide. The historical campus has been classified as UNESCO World Heritage in 2013 under the name "*Universidade de Coimbra – Alta e Sofia*". Coimbra is a fascinating destination that you should not miss! From Coimbra you can easily reach many other

fantastic places in Portugal. By train, bus, or car, one can get to Porto, Lisbon, and many small villages in the interior of the country, where the food, the people, and the landscape will surprise you and create memories for life.

The contact person for ICPE 2023 is Marco Vieira, from the University of Coimbra (Portugal), who will be General Co-Chair along with Valeria Cardellini, from the University of Rome Tor Vergata (Italy). The PC Co-Chairs will be announced in the near future. The industrial track chair will be David Daly from MongoDB.

REPORT: DEVOPS PERFORMANCE WORKING GROUP

DevOps is an emerging principle for engineering and operating software systems. It aims to increase the rate and velocity of releasing new software versions, which is, for instance, achieved by a high degree of automation and by integrating development and operations responsibilities. DevOps imposes immense challenges for quality assurance, e.g., concerning performance and related attributes. Key reasons are that respective activities are constrained by time and that the environment in which a software system is running is ever-changing. On the other hand, DevOps provides great opportunities because the integration between development and operations allows for a high degree of automation as well as a streamlined collection and analytics of performance data.

The RG DevOps Performance Working Group is a forum for individuals and organizations interested in the interplay of DevOps and performance engineering. The mission of the working group is to consolidate concepts and tools to better integrate these activities. Its membership body currently includes representatives of Charles University, Concordia University, Imperial College London, Karlsruhe Institute of Technology, MongoDB, University of Alberta, University of L'Aquila, University of Hamburg, University of South Carolina, University of Stuttgart, and University of Würzburg.

The group as a whole meets in online meetings that are held on a monthly basis. In addition to the discussion of organizational topics, these general meetings include a technical presentation by group members or by invited guests. In total, 8 group meetings were held in 2021, including the following talks as part of our monthly lecture series:

- **“A model-driven approach for continuous performance engineering in microservice-based systems”** by Michele Tucci, University of L'Aquila, Italy

- **“On the impact of performance antipatterns in multi-objective software model refactoring optimization”** by Daniele di Pompeo, University of L'Aquila, Italy
- **“Continuous integration of architectural performance models (CIPM) with parametric dependencies”** by Manar Mazkatli, KIT, Germany
- **“White-box analysis over machine learning: Modeling performance of configurable systems”** by Miguel Velez, CMU, US
- **“CADET: Debugging misconfigurations using counterfactual reasoning”** by Shahriar Iqbal, USC, US
- **“Performance regression testing for SAP HANA: Overview and challenges”** by Thomas Bach, SAP, Germany
- **“Measuring and modeling the performance configurations of distributed DBMS”** by Johannes Grohmann, University of Würzburg, Germany
- **“Performance testing driven by reinforcement learning”** by Mahshid Helali Moghadam, RISE Research Institute, Sweden

In addition to the monthly meetings, the group operates in subgroups consisting of 6-8 participants who collaborate closely on concrete topics. Collaborations include joint research papers and jointly supervised student projects. The subgroups meet biweekly and report to the whole group once a month in the regular meeting. The current subgroups are:

1. **Performance testing of next-generation cloud applications:** This subgroup focuses on the challenges of performance testing next-generation cloud applications. In 2021, the group conducted an empirical study on load testing of serverless applications and submitted a paper on the results to the Journal of Systems and Software, which has recently been accepted [1].
2. **Performance change point detection:** This newly established group focuses on the challenges of using performance detection methods in real large-systems performance data. In collaboration with three major tech companies, SAP, MongoDB, and Oracle, this subgroup is establishing the first curated dataset of real performance changes and assessing the quality of current change detection methods.
3. **Search-based software performance engineering:** The focus of this new subgroup (started in Jan 2022) is on developing novel approaches and research directions on the multi-criteria optimization

of performance-related quality attributes in software architectures.

4. **Performance of continuous delivery infrastructures:** This subgroup focuses on the evaluation and improvement of continuous delivery (CD) infrastructures. Building on previously analyzed performance data of a CD system, the group has a current focus on simulating and optimizing CD pipelines.
5. **Resilience engineering for cloud-native applications:** The group investigates novel approaches, tools, and data sets for resilience engineering including interactive visualization, runtime monitoring, scenario-based architecture evaluation, risk analysis, chaos engineering, formal methods, and other techniques. Current joint efforts include the evolution of a resilience simulator and its application to scenario-based analysis of resilience properties and container orchestration strategies.
6. **Cost-effective performance regression detection:** This subgroup focuses on the challenges of predicting system performance and detecting performance regressions in a cost-effective manner. Current efforts include developing novel approaches to constructing and analyzing system architectural performance models incorporating the insights of local component performance models.

Two members of the group were awarded at the SPEC Annual Meeting 2022:

- **Simon Eismann won a SPECtacular award:** “For several years, this individual has made significant contributions to the SPEC Research Group Cloud Working Group and to the RG DevOps Performance working group. He led the study of the performance of microservices and stability for serverless applications. He also led the effort to create a SPEC RG Technical Report reviewing serverless use cases which led to the publication of related peer-reviewed papers. Simon Eismann has been an active participant in the organization of events that benefit the SPEC research community, served as an evaluator of artifacts, and contributed the TeaStore microservice reference application to the SPEC Research Tool Repository. SPEC has significantly benefitted from Simon’s continuous and dedicated work which has resulted in tangible results that advance the SPEC Research Group.”
- **Cor-Paul Bezemer won a SPEC contributor award:** “This individual has contributed a tremendous amount of time and effort to SPEC RG in his positions of RG Secretary, RG DevOps Vice-Chair, Co-program chair of ICPE 2021, workshop and financial chair of ICPE 2020, and program committee

member of ICPE since 2018. Cor-Paul Bezemer is one of the main drivers of the ICPE 2021 event and was one of the local organizers of ICPE 2020. He has contributed to the success of SPEC RG DevOps group with a contribution of significant research.”

Several members of the group are active in the organization of international events. In 2021, the group co-organized:

- Journal special issue on software performance [2]. Group members have initiated a special issue on “Software Performance” in the Empirical Software Engineering Journal, which is a top-ranked software engineering journal.

For more information about the DevOps Performance Working Group (including our mission, activities, meetings, presentations, and projects), please visit our web page [3].

If you are interested in following the discussions or contributing actively, please contact the working group chairs.

Cor-Paul Bezemer (University of Alberta), André v. Hoorn (University of Hamburg), Simon Eismann (University of Würzburg)

<https://research.spec.org/en/working-groups/rg-devops-performance.html>

[1] Simon Eismann, Diego Elias Costa, Lizhi Liao, Cor-Paul Bezemer, Weiyi Shang, André van Hoorn, Samuel Kounev: A Case Study on the Stability of Performance Tests for Serverless Applications. Journal of Systems and Software (accepted and in print). Pre-print available at: <https://arxiv.org/abs/2107.13320>

[2] https://emsejournal.github.io/special_issues/2021_Software_Performance.html

[3] <http://research.spec.org/devopswg/>

REPORT: CLOUD WORKING GROUP

In 2021, the SPEC RG Cloud Group has driven several activities aligned with its long-term mission of furthering cloud benchmarking, quantitative evaluation, and experimental analysis, in directions relevant for both academia and industry. We have focused this year on novel cloud paradigms such as Functions-as-a-Service, Serverless Computing, the Cloud Continuum extending clouds with fog and edge devices, Convergence of HPC and Big Data as cloud services.

The scope of the group is to develop new methodological elements for gaining deeper understanding not only of cloud performance, but also of cloud operation and behavior, through diverse quantitative evaluation tools, including benchmarks, metrics, and workload generators’. We consider properties such as elasticity, performance isolation, dependability, and other non-functional system properties, in addition to classical performance-related metrics such as response time, throughput, scalability, and efficiency. Our work towards benchmark prototypes

includes designing reference architectures, standardizing use cases, observing patterns, and methods for reproducibility.

In 2021, through monthly online meetings facilitated by SPEC's Zoom and meetings focusing on furthering specific activities, and through continuous discussion via a Slack workspace, we have advanced work on the following main topics:

1. **Serverless Use Cases:** Serverless architectures are rapidly adopted by practitioners, but many characteristics of serverless applications are either unknown or contested. Therefore, the SPEC RG Cloud Group started collecting and analyzing serverless use cases in late 2019. This activity culminated in a collection of 89 serverless applications from academic papers, open-source repositories, industrial articles, and scientific applications analyzed along 24 characteristics. To investigate community consensus on this topic, we conducted a meta-analysis of related studies on the characteristics of serverless applications. The results of this activity have been published as a IEEE Transactions on Software Engineering (TSE) article [1].
2. **Serverless Function Sizing:** Selecting the optimal size of serverless functions is quite challenging, so developers often neglect it despite its significant cost and performance benefits. Therefore, members of the SPEC RG Cloud Group worked on an automated approach to determine the optimal size of serverless functions without active performance measurements. This work was published at MIDDLEWARE 2021 and has been recognized by the Best Student Paper Award [2].
3. **Storage for Data Intensive Serverless Computing:** Data analytics is a socially and industrially important workload that benefits from the scale and pay-per-use nature of serverless computing. But, serverless data analytics applications are bottlenecked by serverless storage systems (such as S3) with high latency and variable bandwidth. To overcome this bottleneck, data analytics applications use a caching system between the compute and storage systems. To fulfill the pay-per-use promise, the cache needs to be auto-scaled dynamically along with other resources. We are investigating the impact of auto-scaling policies on the performance of serverless data analytics applications with caches. We designed and implemented a workload generator to evaluate such systems. We intend to compare and analyze the interplay of scheduling, caching, and autoscaling policies for different workloads. The result will be a performance model to pick the appropriate auto-scaling parameters for a

given workload.

4. In the **Edge activity** of the SPEC Cloud Group, we conduct systematic studies of computing models focusing on edge resources. Now, we have two papers under review:
 - (1) The first paper (Titled: Edge Workload Trace Gathering and Analysis for Benchmarking) defines workload classes, and collects traces for three real-world edge applications that represent those classes. This paper has been submitted to the 6th IEEE International Conference on Fog and Edge Computing 2022.
 - (2) The second paper (Titled: The SPEC-RG Reference Architecture for the Edge Continuum) presents a reference architecture for the edge continuum, and maps to it a diverse set of state-of-the-art resource managers from different computing models, providing strong evidence of the generality of this architecture. This paper has been submitted to the 42nd IEEE International Conference on Distributed Computing Systems 2022.

Besides these focused activities, the Cloud WG has been acting successfully in the following directions: (1) Organization of the yearly – meanwhile 5th – workshop Hot-CloudPerf [3], and (2) the Dagstuhl seminar 21201 on “Serverless Computing” [4]. The vice-chair of the group, Nikolas Herbst, serves as ICPE 2022 PC co-chair and Cristina Abad as ICPE 2022 WIP/Vision Track chair.

To conclude, 2021 was a full and successful year for the RG Cloud Group. We are looking forward to an even more successful 2022. For this, we are actively seeking new participants and activities. You can also join ongoing activities.

Alexandru Iosup (Vrije Universiteit Amsterdam, Chair) and Nikolas Herbst (University of Würzburg, Vice-Chair)

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

[1] The State of Serverless Applications: Collection, Characterization, and Community Consensus. Eismann, Simon; Scheuner, Joel; van Eyk, Erwin; Schwinger, Maximilian; Grohmann, Johannes; Herbst, Nikolas; Abad, Cristina; Iosup, Alexandru; in Transactions on Software Engineering (2021)

[2] Sizeless: Predicting the Optimal Size of Serverless Functions. Eismann, Simon; Bui, Long; Grohmann, Johannes; Abad, Cristina; Herbst, Nikolas; Kounev, Samuel; in Proceedings of the 22nd International MIDDLEWARE Conference (2021). 248–259. Best Student Paper Award, ACM Artifacts Evaluated — Functional

[3] 5th Workshop on Hot Topics in Cloud Computing Performance (Hot-CloudPerf 2022) <https://hotcloudperf.spec.org/>

[4] Serverless Computing (Dagstuhl Seminar 21201). Abad, Cristina; Foster, Ian T.; Herbst, Nikolas; Iosup, Alexandru; in Dagstuhl Reports, (C. Abad; I. T. Foster; N. Herbst; A. Iosup, eds.) (2021). 11(4) 34–93. <https://drops.dagstuhl.de/opus/volltexte/2021/14798/>

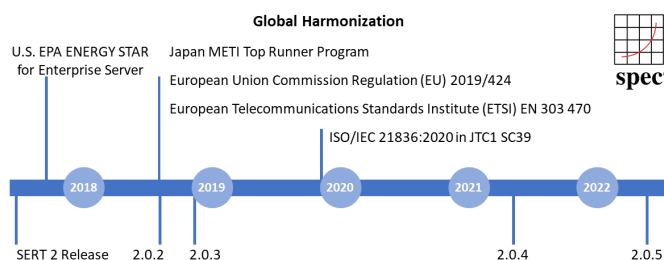
REPORT: POWER WORKING GROUP

The SPEC Research Power Working Group (AMD, Dell, HPE, Intel, IBM, Microsoft, University of Wuerzburg; Chair: Maximilian Meissner, Vice-Chair: Klaus-Dieter Lange) has operated tightly coupled with the SPECpower Committee since its inception in 2017.

Our mission is to research the energy and resource efficiency of computing devices and software. Our close collaboration fosters the interaction between industry and academia by contributing research that enhances and promotes methods and tools for energy and resource efficiency evaluation to address this essential concern for the industry, academia, and regulatory institutions.

Closely connected to the SPEC RG Power WG and the SPECpower Committee is the International Standards Group (ISG), which we introduced to the public at ICPE 2021 [1]. The ISG oversees the establishment of standardized benchmarks primarily developed for the use in government regulations and programs, and collaborates with national and international standard development organizations to enhance global standards. Within the ISG, the Server Efficiency Committee (Chair: Sanjay Sharma) and the Client Efficiency Committee (Chair: Roger Tiple) have been established. Their goal is to oversee the establishment and development of standardized benchmarks for their respective fields. Current members and associates of the ISG are AMD, Ampere, Apple, Dell, HPE, IBM, Inspur, Intel, Microsoft, NVIDIA, Quanta, and the University of Würzburg.

With the release of ISO/IEC 21836:2020 in August 2020 [3], the founding members of ISG already have achieved the standardization of the server energy effectiveness metric (SEEM). ISG is striving towards global harmonization, where one internationally standardized benchmark is used in every national regulation for server energy efficiency. The Server Efficiency Rating Tool (SERT) suite,



which the U.S. EPA already has adopted with Version 3.0 of their Server Energy Star program, will be enhanced further under the responsibility of the ISG starting with release of the SERT 2.0.5 suite earlier this year. Among a series of enhancements and bug fixes, the SERT 2.0.5 release supports Java 17 and its new features.

Apart from our contributions to government regulations,

one of our goals is to raise practitioners' awareness with respect to the influence of different software designs of the same functionality on energy efficiency. To this end, we investigated the energy efficiency of common sorting algorithms in different implementation variants and for different problem sizes [2]. The results demonstrate that simple design choices can have significant impact on energy efficiency, a fact often overseen by developers.

At ICPE 2022, RG Power, SPECpower and ISG jointly offer a tutorial titled "SPEC Server Efficiency Benchmark Development – How to Contribute to the Future of Energy Conservation". To keep pace with the development of the fast moving IT landscape, SPECpower plans to introduce a workload bounty program to encourage researchers to develop novel workloads. Submitted workloads will be considered for inclusion in a future SPECpower benchmark and awarded. The goal of the tutorial is to equip the participants with the necessary knowledge and tools to conduct energy efficiency experiments and to use the SERT 2 suite. The tutorial will show-case how to use the Chauffer Worklet Development Kit to develop next-generation workloads to enhance the real-world relevance of the future SPECpower benchmarks, a critical element for the benchmarks to contribute to future energy conservation.

In the coming years, the group plans to expand its activities to include investigating energy efficiency in the context of Serverless Computing. The SPEC RG Power WG looks forward to new and exciting challenges in power, resource, and energy efficiency benchmarking. The group is happy to accept new members and visions for additional research directions in the general area of energy and resource efficiency benchmarking.

Maximilian Meissner (University of Würzburg),
Klaus-Dieter Lange (SPEC International Standards Group Chair)

<https://research.spec.org/working-groups/rg-power/>

[1] Norbert Schmitt, Klaus-Dieter Lange, Sanjay Sharma, Aaron Cragin, David Reiner, and Samuel Kounev. 2021. SPEC — Spotlight on the International Standards Group (ISG). Companion of the ACM/SPEC International Conference on Performance Engineering. Association for Computing Machinery, New York, NY, USA, 167–168.

[2] Norbert Schmitt, Supriya Kamthania, Nishant Rawtani, Luis Mendoza, Klaus-Dieter Lange, Samuel Kounev. 2021. Energy-Efficiency Comparison of Common Sorting Algorithms. 29th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS).

[3] <https://www.iso.org/standard/71926.html>

[4] [ENERGY STAR Computer Servers SERT V2.0.4 Clarification Memo](#)

REPORT: SECURITY WORKING GROUP

The SPEC RG Security Benchmarking Working Group, after establishing its long-term agenda, continues to work towards devising and conducting impactful results while achieving the defined goals. In 2021, the Working Group had papers published at the 13th International Workshop on Software Aging and Rejuvenation (WoSAR'21) [1, 2], and at the 12th Symposium on Software Performance (SSP'21) [3].

The paper entitled “*Software Testing Strategies for Detecting Hypercall Handlers’ Aging-related Bugs*” shares the insights collected during the experiments conducted using hypercalls and their parameters and to construct test cases for hypervisor aging of Microsoft Hyper-V. The bug detected, which was reported and acknowledged by Microsoft, is presented. The work further proposes the possibility to automate the analysis process for valid parameter detection and execution conditions.

The paper with the title “*My Services Got Old! Can Kubernetes Handle the Aging of Microservices?*” presents a study on the effects of aging in microservices and the utilization of faults to accelerate aging effects. The work is based on three operation scenarios for a representative microservice-based system through the utilization of stress testing and fault injection as techniques to potentiate aging in the services.

The paper published as “*Measuring the Performance Impact of Branching Instructions*” assesses the performance difference between using and avoiding branching instructions to check for values to log for every injection, measured by the hypercall throughput of the injector tool.

The group is also involved in a special section for the IEEE Transactions on Emerging Topics in Computing on Applied Software Aging and Rejuvenation in Computing Systems [4], with Alberto Avitser being one of the guest editors.

The SPEC RG Security Benchmarking Working Group continues to push forward in several research directions, including:

- extending the previously proposed approach and the framework for testing the robustness and performance of hypercall interfaces
- identifying challenges in the area of evaluating robustness and performance aspects of security-relevant system components and security mechanisms
- proposing new methodologies to evaluate security of systems based on injecting the effects of intrusions
- evaluating the performance of different approaches in the context of detecting security intrusions based

on performance and behavior profiles

- constructing a methodology for the evaluation of intrusion detecting systems capable of monitoring over multiple containers’ activities
- design and develop a methodology to benchmark ransomware at multiple levels of its operation

Aleksandar Milenkoski (Cybereason/Univ. Würzburg, Germany),
Nuno Antunes (University of Coimbra),
José Flora (University of Coimbra)

<https://research.spec.org/working-groups/rg-ids-benchmarking.html>

[1] L. Beierlieb, L. Iffländer, A. Milenkoski, A. Avitser, N. Antunes, and S. Kounev, ‘Software Testing Strategies for Detecting Hypercall Handlers’ Aging-related Bugs’, in 2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), Oct. 2021, pp. 48–55. doi: 10.1109/ISSREW53611.2021.00043.

[2] J. Flora, P. Gonçalves, M. Teixeira, and N. Antunes, ‘My Services Got Old! Can Kubernetes Handle the Aging of Microservices?’, in 2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), 2021, pp. 40–47. doi: 10.1109/ISSREW53611.2021.00042.

[3] L. Beierlieb, L. Iffländer, A. Milenkoski, T. Prantl, and S. Kounev, ‘Measuring the Performance Impact of Branching Instructions’, presented at the 2021 Symposium on Software Performance (SSP), Nov. 2021.

[4] IEEE TETC, Special Section on Applied Software Aging and Rejuvenation in Computing Systems <https://www.computer.org/digital-library/journals/ec/cfp-software-aging-rejuvenation>

REPORT: PREDICTIVE DATA ANALYTICS WORKING GROUP

The research field of data analytics/science has grown significantly in recent years as a means to make sense of the vast amount of available data. It has permeated every aspect of computer science and engineering and is heavily involved in business decision-making. For example, in the field of performance engineering, performance prediction is an instrument for controlling and improving the behavior of a system. Analogously, data analytics is playing an essential role for companies and science. In both examples, streamlining data analytics processes (DataOps) enables multi-tenant access to data and models.

To bridge the missing links between the facets involved in data analytics, namely big data storage and provisioning, data versioning, and performance evaluation, the SPEC predictive data analytics working group was established in June 2021. The group’s ambition is to standardize and benchmark the entire data lifecycle, i.e., the analytics/prediction methods and especially pipelines for data analytics ranging from big data storage and preprocessing to analytics and assessment, as well as to provide heuristics for the selection of tools, patterns, and infrastructure.

Furthermore, the group members are interested in promoting the interaction between industry and academia by contributing research towards standardization and benchmarking of the different aspects of data analytics. For this, the group members investigate data analytics-related methodologies, systems, and metrics. Another important goal is to support open data and promote the reproducibility of experiments and benchmarking of data analytics methods. The interests of the group lie in but are not limited to:

1. Performance modeling, analysis, testing, and prediction
2. Performance analysis of ML Systems
3. Resource autoscaling and reconfigurable systems
4. Performance behavior in resource-constrained environments
5. Change point and anomaly detection
6. Time series analysis and forecasting
7. Streamlining the data science process (DataOps)
8. Benchmarking of big data infrastructure

At ICPE 2022, a poster paper [1] and a data challenge paper [2] will be presented by the group. We will introduce the group during the poster session and highlight our research interests and possible collaborations. The poster session will be followed by the group's annual meeting (April 11, 7 pm CEST), which you are welcome to attend if you are interested.

In general, the group is meeting online every second Wednesday at 4 pm CEST; you are welcome to attend. If you are interested in joining the group, please contact André Bauer (andre.bauer@uni-wuerzburg.de).

André Bauer, University of Würzburg, Germany

<https://research.spec.org/working-groups/rg-predictive-data-analytics/>

[1] A. Bauer, M. Leznik, M. Iqbal, I. Trubin, P. Jamshidi 'SPEC Research — Introducing the Predictive Data Analytics Working Group', in Companion of the 13th ACM/SPEC International Conference on Performance Engineering 2022.

[2] M. Leznik, M. Iqbal, I. Trubin, A. Lochner, P. Jamshidi, A. Bauer 'Change Point Detection for MongoDB Time Series Performance Regression', in Companion of the 13th ACM/SPEC International Conference on Performance Engineering 2022.

DAGSTUHL SEMINAR 21201 "SERVERLESS COMPUTING"

In the backbone of our digital society, cloud computing enables an efficient, utility-like ecosystem of developing, composing, and providing software services. Responding

to a trend to make cloud computing services more accessible, fine-grained, and affordable, serverless computing has gained rapid adoption in practice, and garnered much interest from both industry and academia. However successful, serverless computing manifests today the opportunities and challenges of emerging technology: a rapidly growing field but scattered vision, plenty of new technologies but no coherent approach to design solutions from them, many simple applications but no impressive advanced solution, the emergence of a cloud continuum (resources from datacenters to the edge) but no clear path to leverage it efficiently, and overall much need but also much technical complexity. Several related but disjoint fields, notably software and systems engineering, parallel and distributed systems, and system and performance analysis and modeling, aim to address these opportunities and challenges. Excellent collaboration between these fields in the next decade will be critical in establishing serverless computing as a viable technology. We organized this Dagstuhl seminar to bring together researchers, developers, and practitioners across disciplines in serverless computing, to develop a vision and detailed answers to the timely and relevant, open challenges related to the following topics:

1. Topic 1: design decisions for serverless systems, platforms, and ecosystems,
2. Topic 2: software engineering of serverless applications, but also systems, platforms, and ecosystems
3. Topic 3: applications and domain requirements for serverless computing,
4. Topic 4: evaluation of serverless solutions, and beyond (privacy, cyber-physical systems, etc.).

The seminar gathered a diverse group of about 50 high-quality researchers from the large-scale systems, software engineering, and computer performance engineering communities. Following preparations and preliminary work started in 2020, the seminar took place in May 2021 in a special online format.

Intense discussions were kick-started by four distinguished keynotes:

1. "Serverless Predictions: 2021-2030" given jointly by Pedro García López (Universitat Rovira i Virgili – Tarragona, ES) and Bernard Metzler (IBM Research-Zurich, CH)
2. "Developer Experience for Serverless: Challenges and Opportunities" given by Robert Chatley (Imperial College London, GB)
3. "Federated Function as a Service" given jointly by Kyle Chard (University of Chicago, US) and Ian T. Foster (Argonne National Laboratory – Lemont, US)

4. “Characterizing Serverless Systems” given by Mohammad Shahrad (University of British Columbia – Vancouver, CA)

Nikolas Herbst (University of Würzburg)

<https://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=21201>

[1] Serverless Computing (Dagstuhl Seminar 21201). Abad, Cristina; Foster, Ian T.; Herbst, Nikolas; Iosup, Alexandru; in Dagstuhl Reports, (C. Abad; I. T. Foster; N. Herbst; A. Iosup, eds.) (2021). 11(4) 34–93. <https://drops.dagstuhl.de/opus/volltexte/2021/14798/>

FIFTH WORKSHOP ON HOT TOPICS IN CLOUD COMPUTING PERFORMANCE (HOTCLOUDPERF 2022)

Since 2018, the Cloud working group is organizing a yearly workshop in connection with their face-to-face meeting: The fourth, virtual edition of HotCloudPerf was well-attended in conjunction with the ACM/SPEC ICPE 2021. The workshop featured 2 keynotes and 7 papers. The fifth edition is to be held in conjunction with ICPE 2022 and features the following three keynotes:

1. Lydia Chen: Dependability Management for Data-centers: A Machine Learning Perspective
2. Evgenia Smirni: Serverless Machine Learning Serving for Scalable Workflows
3. Paul Brebner: Scaling Open Source Big Data Cloud Applications is Easy/Hard

and the following 6 papers:

1. Floriment Klinaku, Martina Rapp, Jörg Henss and Stephan Rhode. Beauty and the beast: A case study on performance prototyping of data-intensive containerized cloud applications.
2. Thrivikraman V, Vishnu R Dixit, Nikhil Ram S, Vikas K Gowda, Santhosh Kumar Vasudevan and Subramaniam Kalambur. MiSeRTrace: Kernel-level Request Tracing for Microservice Visibility.
3. Laurens Versluis and Alexandru Iosup. TaskFlow: An Energy- and Makespan-Aware Task Placement Policy for Workflow Scheduling through Delay Management.
4. Robert Cordingly and Wes Lloyd. FaaSSET: A Jupyter notebook to streamline every facet of serverless development.
5. Robert Schmitz, Danielle Lambion, Robert Cordingly, Navid Heydari and Wes Lloyd. Characterizing X86 and ARM Serverless Performance Variation: A natural language processing case study.
6. George Kousiouris, Chris Giannakos, Konstantinos Tserpes and Teta Stamatii. Measuring Baseline

Overheads in Different Orchestration Mechanisms for Large FaaS Workflows.

Simon Eismann (University of Würzburg)

<https://hotcloudperf.spec.org/>

COMBINING NODE-RED AND OPENWHISK FOR PATTERN-BASED VISUAL DEVELOPMENT, ORCHESTRATION AND EXECUTION OF COMPLEX FAAS WORKFLOWS

Harokopio University of Athens, member of the SPEC RG and Cloud WG, has been involved in the last year in the H2020 PHYSICS project [1], a Research and Innovation Action funded by the EC. The aim of PHYSICS is to provide a framework for the use of the Function as a Service model, including a design and development environment for application adaptation and migration. The goal of the environment, based on the widely used in the IoT domain Node-RED tool, is to enable a more user friendly and abstract function and workflow creation for complex FaaS applications. To this end, it bypasses workflow description and function reuse limitations of the current FaaS platforms, by providing an extendable, pattern-enriched palette of ready-made, reusable functionality. The latter can be combined in arbitrary ways, including also application logic, through event driven workflows in the Node-RED runtime. Indicative patterns include different state management capabilities, parallelization and optimization, workflow primitives, data acquisition and anonymization as well as orchestration. The environment embeds seamless DevOps processes for generating the deployable artefacts (i.e., functions and images) of the target FaaS platform (Openwhisk). Annotation mechanisms are also foreseen for the developer to dictate diverse execution options or management guidelines towards the deployment and operation stacks. Combination of performance modelling approaches with the patterns is also conducted in order to dynamically decide on pattern parameter settings and configurations for an enhanced runtime behavior.

George Kousiouris, Harokopio University of Athens

[1] <https://physics-faas.eu/>

SELECTED ABSTRACTS

A Case Study on the Stability of Performance Tests for Serverless Applications

Context. While in serverless computing, application resource management and operational concerns are generally delegated to the cloud provider, ensuring that serverless applications meet their performance requirements is still a responsibility of the developers. Performance testing is a commonly used performance assessment practice; however, it traditionally requires visibility of the resource environment. **Objective.** In this study, we investigate whether performance tests of serverless applications are stable, that is, if their results are reproducible, and what implications the serverless paradigm has for performance tests. **Method.** We conduct a case study where we collect two datasets of performance test results: (a) repetitions of performance tests for varying memory size and load intensities and (b) three repetitions of the same performance test every day for ten months. **Results.** We find that performance tests of serverless applications are comparatively stable if conducted on the same day. However, we also observe short-term performance variations and frequent long-term performance changes. **Conclusion.** Performance tests for serverless applications can be stable; however, the serverless model impacts the planning, execution, and analysis of performance tests. In this paper, we discuss the benefits and challenges of microservices from a performance tester's point of view. Through a series of experiments on the TeaStore application, we demonstrate how microservices affect the performance testing process, and we demonstrate that it is not straightforward to achieve reliable performance testing results for a microservice application.

Simon Eismann, Diego Elias Costa, Lizhi Liao, Cor-Paul Bezemer, Weiyi Shang, André van Hoorn, Samuel Kounev: A Case Study on the Stability of Performance Tests for Serverless Applications. *JSS* (under review). Pre-print available at: <https://arxiv.org/abs/2107.13320>

Sizeless: Predicting the Optimal Size of Serverless Functions

Serverless functions are an emerging cloud computing paradigm that is being rapidly adopted by both industry and academia. In this cloud computing model, the provider opaquely handles resource management tasks such as resource provisioning, deployment, and auto-scaling. The only resource management task that developers are still in charge of is selecting how much resources are allocated to each worker instance. However, selecting the optimal size of serverless functions is quite challenging, so developers often neglect it despite its significant cost and performance benefits. Existing approaches aiming to automate serverless functions re-

source sizing require dedicated performance tests, which are time-consuming to implement and maintain. In this paper, we introduce an approach to predict the optimal resource size of a serverless function using monitoring data from a single resource size. As our approach does not require dedicated performance tests, it enables cloud providers to implement resource sizing on a platform level and automate the last resource management task associated with serverless functions. We evaluate our approach on four different serverless applications on AWS, where it predicts the execution time of the other memory sizes based on monitoring data for a single memory size with an average prediction error of 15.3%. Based on these predictions, it selects the optimal memory size for 79.0% of the serverless functions and the secondbest memory size for 12.3% of the serverless functions, which results in an average speedup of 39.7% while also decreasing average costs by 2.6%.

Sizeless: Predicting the Optimal Size of Serverless Functions. Eismann, Simon; Bui, Long; Grohmann, Johannes; Abad, Cristina; Herbst, Nikolas; Kounev, Samuel; in *Proceedings of the 22nd International MIDDLEWARE Conference* (2021). 248–259. DOI: <https://doi.org/10.1145/3464298.3493398>

The State of Serverless Applications: Collection, Characterization, and Community Consensus

Over the last five years, all major cloud platform providers have increased their serverless offerings. Many early adopters report significant benefits for serverless-based over traditional applications, and many companies are considering moving to serverless themselves. However, currently there exist only few, scattered, and sometimes even conflicting reports on when serverless applications are well suited and what the best practices for their implementation are. We address this problem in the present study about the state of serverless applications. We collect descriptions of 89 serverless applications from open-source projects, academic literature, industrial literature, and domain-specific feedback. We analyze 16 characteristics that describe why and when successful adopters are using serverless applications, and how they are building them. We further compare the results of our characterization study to 10 existing, mostly industrial, studies and datasets; this allows us to identify points of consensus across multiple studies, investigate points of disagreement, and overall confirm the validity of our results. The results of this study can help managers to decide if they should adopt serverless technology, engineers to learn about current practices of building serverless applications, and researchers and platform providers to better understand the current landscape of serverless applications.

The State of Serverless Applications: Collection, Characterization, and Community Consensus. Eismann, Simon; Scheuner, Joel; van Eyk,

Erwin; Schwinger, Maximilian; Grohmann, Johannes; Herbst, Nikolas; Abad, Cristina; Iosup, Alexandru; in *Transactions on Software Engineering* (2021). DOI: <https://doi.org/10.1109/TSE.2021.3113940>

An Analysis of Distributed Systems Syllabi With a Focus on Performance-Related Topics

We analyze a dataset of 51 current (2019-2020) Distributed Systems syllabi from top Computer Science programs, focusing on finding the prevalence and context in which topics related to performance are being taught in these courses. We also study the scale of the infrastructure mentioned in DS courses, from small client-server systems to cloud-scale, peer-to-peer, global-scale systems. We make eight main findings, covering goals such as performance, and scalability and its variant elasticity; activities such as performance bench-marking and monitoring; eight selected performance-enhancing techniques (replication, caching, sharding, load balancing, scheduling, streaming, migrating, and offloading); and control issues such as trade-offs that include performance and performance variability.

Cristina L. Abad, Alexandru Iosup, Edwin F. Boza, and Eduardo Ortiz Holguin. 2021. An Analysis of Distributed Systems Syllabi With a Focus on Performance-Related Topics. *Companion of the ACM/SPEC International Conference on Performance Engineering*. Association for Computing Machinery, New York, NY, USA, 121–126. DOI:<https://doi.org/10.1145/3447545.3451197>

Software Testing Strategies for Detecting Hypercall Handlers' Aging-related Bugs

With the continuing rise of cloud technology hypervisors play a vital role in the performance and reliability of current services. As long-running applications, they are susceptible to software aging. Hypervisors offer so-called hypercall interfaces for communication with the hosted virtual machines. These interfaces require thorough testing to ensure their long-term reliability. Existing research deals with the aging properties of hypervisors in general without considering the hypercalls. In this work, we share our experience that we collected during trying to understand hypercalls and their parameters and use them to construct test cases for hypervisor aging of Microsoft Hyper-V. We present a bug that we detected, which was reported and acknowledged by Microsoft. Further, based on our manual binary code analysis, we propose the idea of automating the analysis process to detect valid parameter ranges and execution conditions of hypercalls without manual effort.

L. Beierlieb, L. Iffländer, A. Milenkoski, A. Avitzer, N. Antunes, and S. Kounev, 'Software Testing Strategies for Detecting Hypercall Handlers' Aging-related Bugs', in 2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), Oct. 2021, pp. 48–55. doi: [10.1109/ISSREW53611.2021.00043](https://doi.org/10.1109/ISSREW53611.2021.00043).

My Services Got Old! Can Kubernetes Handle the Aging of Microservices?

The exploding popularity of microservice based applications is taking companies to adopt them along with cloud services to support them. Containers are the common deployment infrastructures that currently serve millions of customers daily, being managed using orchestration platforms that monitor, manage, and automate most of the work. However, there are multiple concerns with the claims put forward by the developers of such tools. In this paper, we study the effects of aging in microservices and the utilization of faults to accelerate aging effects while evaluating the capacity of Kubernetes to detect microservice aging. We consider three operation scenarios for a representative microservice-based system through the utilization of stress testing and fault injection as a manner to potentiate aging in the services composing the system to evaluate the capacity of Kubernetes mechanisms to detect it. The results demonstrate that even though some services tend to accumulate aging effects, with increasing resource consumption, Kubernetes does not detect them nor acts on them, which indicates that the probe mechanisms may be insufficient for aging scenarios. This factor may indicate the necessity for more effective mechanisms, capable of detecting aging early on and act on it in a more proactive manner without requiring the services to become unresponsive.

J. Flora, P. Gonçalves, M. Teixeira, and N. Antunes, 'My Services Got Old! Can Kubernetes Handle the Aging of Microservices?', in 2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), 2021, pp. 40–47. doi: [10.1109/ISSREW53611.2021.00042](https://doi.org/10.1109/ISSREW53611.2021.00042).

Measuring the Performance Impact of Branching Instructions

With the continuing rise of cloud technology, hypervisors play a vital role in the performance and reliability of current services. Hypervisors implement interfaces providing call-based connectivity to hosted virtualization-aware virtual machines. One of them is the hypercall interface, allowing hypervisor service requests from virtual machines. A hypercall injection tool measuring hypercall execution times must minimize internal overhead. Among other things, limiting logging to strictly required information is crucial. However, checks for values to log for every injection requires executing many branch instructions. We assess the performance difference between using and avoiding such branching, measured by the hypercall throughput of the injector tool.

L. Beierlieb, L. Iffländer, A. Milenkoski, T. Prantl, and S. Kounev, 'Measuring the Performance Impact of Branching Instructions', presented at the 2021 Symposium on Software Performance (SSP), Nov. 2021.

Why Is It Not Solved Yet? Challenges for Production-Ready Autoscaling

Autoscaling is a task of major importance in the cloud computing domain as it directly affects both operating costs and customer experience. Although there has been active research in this area for over ten years now, there is still a significant gap between the proposed methods in the literature and the deployed autoscalers in practice. Hence, many research autoscalers do not find their way into production deployments. This paper describes six core challenges that arise in production systems that are still not solved by most research autoscalers. We illustrate these problems through experiments in a realistic cloud environment with a real-world multi-service business application and show that commonly used autoscalers have various shortcomings. In addition, we analyze the behavior of overloaded services and show that these can be problematic for existing autoscalers. Generally, we analyze that these challenges are only insufficiently addressed in the literature and conclude that future scaling approaches should focus on the needs of production systems.

Martin Straesser, Johannes Grohmann, Jóakim von Kistowski, Simon Eismann, André Bauer, and Samuel Kounev. 2022. Why Is It Not Solved Yet? Challenges for Production-Ready Autoscaling. In *Proceedings of the 2022 ACM/SPEC International Conference on Performance Engineering (ICPE'22)*, April 9–13, 2022, Beijing, China. ACM, New York, NY, USA

OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters

Cloud datacenters are important for the digital society, serving stakeholders across industry, government, and academia. Simulation is a critical part of exploring datacenter technologies, enabling scalable experimentation with millions of jobs and hundreds of thousands of machines, and what-if analysis in a matter of minutes to hours. Although the community has already developed powerful simulators, emerging technologies and applications in modern datacenters require new approaches. Addressing this requirement, in this work we propose OpenDC, a new platform for datacenter simulation.

OpenDC includes novel models for emerging cloud datacenter technologies and applications, such as serverless computing with FaaS deployment and TensorFlow-based machine learning. Our design also focuses on convenience, with a web-based interface for interactive experimentation, support for experiment automation, a library of prefabs for constructing and sharing datacenter designs, and support for diverse input formats and output metrics. We implement, validate, and open-source OpenDC 2.0, a significant redesign and release after a multi-year research and development process. We demonstrate the

benefits of OpenDC for the field through a set of representative use-cases: serverless, machine learning, procurement of HPC-as-a-Service infrastructure, educational practices, and reproducibility studies. Overall, OpenDC helps understand how datacenters work, design datacenter infrastructure, and train the next generation of experts.

F. Mastenbroek, G. Andreadis, S. Jounaid, W. Lai, J. Burley, J. Bosch, E. van Eyk, L. Versluis, V. van Beek, A. Iosup, "OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters," *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, 2021

Automated Triage of Performance Change Points Using Time Series Analysis and Machine Learning

Performance regression testing is a foundation of modern DevOps processes and pipelines. Thus, the detection of change points, i.e., updates or commits that cause a significant change in the performance of the software, is of special importance. Typically, validating potential change points relies on humans, which is a considerable bottleneck and costs time and effort. This work proposes a solution to classify and detect change points automatically. On the performance test data set provided by MongoDB, our approach classifies potential change points with an AUC of 95.8% and accuracy of 94.3%, whereas the detection and classification of change points based on previous and the current commits exhibits an AUC of 92.0% and accuracy of 84.3%. In both cases, our approach can save time-consuming and costly human work.

André Bauer, Martin Straesser, Lukas Beierlieb, Maximilian Meissner, and Samuel Kounev. 2022. Automated Triage of Performance Change Points Using Time Series Analysis and Machine Learning: Data Challenge Paper. In *ICPE '22: 13th ACM/SPEC International Conference on Performance Engineering*, April 09–13, 2022, Beijing, China

SPEC — Spotlight on the International Standards Group (ISG)

The driving philosophy for the Standard Performance Evaluation Corporation (SPEC) is to ensure that the marketplace has a fair and useful set of metrics to differentiate systems, by providing standardized benchmark suites and international standards. This poster-paper gives an overview of SPEC with a focus on the newly founded International Standards Group (ISG).

Norbert Schmitt, Klaus-Dieter Lange, Sanjay Sharma, Aaron Cragin, David Reiner, and Samuel Kounev. 2021. SPEC — Spotlight on the International Standards Group (ISG). In *ICPE '21: Companion of the ACM/SPEC International Conference on Performance Engineering*