

The Eighth International Workshop on Load Testing and Benchmarking of Software Systems (LTB 2020)

Alexander Podelko
Oracle
Stamford, CT, USA
alex.podelko@oracle.com

Tse-Hsun (Peter) Chen
Concordia University
Montreal, Quebec, Canada
peterc@encs.concordia.ca

Hamzeh Khazaei
York University
Toronto, Ontario, Canada
hkh@eecs.yorku.ca



LTB 2020 Load Testing & Benchmarking

ABSTRACT

The Eighth International Workshop on Load Testing and Benchmarking of Software Systems (LTB 2020) is a full day event bringing together software testing researchers, practitioners and tool developers to discuss the challenges and opportunities of conducting research on load testing and benchmarking software systems. The workshop, co-located with the 11th International Conference on Performance Engineering (ICPE 2020), is held on April 20th, 2020 in Edmonton, Alberta, Canada.

CCS CONCEPTS

• **General and reference** → **Performance**; • **Software and its engineering** → **Software performance**.

KEYWORDS

load testing; benchmarking; software systems; performance

ACM Reference Format:

Alexander Podelko, Tse-Hsun (Peter) Chen, and Hamzeh Khazaei. 2020. The Eighth International Workshop on Load Testing and Benchmarking of Software Systems (LTB 2020). In *Proceedings of the 2020 ACM/SPEC International Conference on Performance Engineering (ICPE '20)*, April 20–24, 2020, Edmonton, AB, Canada. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3358960.3383770>

1 INTRODUCTION

Software systems (e.g., smartphone apps, desktop applications, telecommunication infrastructures, cloud and enterprise systems,

etc.) have strict requirements on software performance. Failure to meet these requirements may cause business losses, customer defection, brand damage and other serious consequences. Hence, in addition to conventional functional testing, the performance of these systems must be verified through load testing or benchmarking to ensure quality service.

Load testing and benchmarking software systems are difficult tasks, which require a great understanding of the system under test and customer behavior. Practitioners face many challenges such as tooling (choosing and implementing the testing tools), environments (software and hardware setup) and time (limited time to design, test, and analyze). This one-day workshop brings together software testing researchers, practitioners and tool developers to discuss the challenges and opportunities of conducting research on load testing and benchmarking software systems.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPE '20, April 20–24, 2020, Edmonton, AB, Canada

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6991-6/20/04.

<https://doi.org/10.1145/3358960.3383770>

etc.) have strict requirements on software performance. Failure to meet these requirements may cause business losses, customer defection, brand damage and other serious consequences. Hence, in addition to conventional functional testing, the performance of these systems must be verified through load testing or benchmarking to ensure quality service.

Load testing and benchmarking software systems are difficult tasks, which require a great understanding of the system under test and customer behavior. Practitioners face many challenges such as tooling (choosing and implementing the testing tools), environments (software and hardware setup) and time (limited time to design, test, and analyze). This one-day workshop brings together software testing researchers, practitioners and tool developers to discuss the challenges and opportunities of conducting research on load testing and benchmarking software systems.

2 THE WORKSHOP

The workshop is held on April 20th, 2020 in Edmonton, Alberta, Canada, and consists of three keynotes, one research paper, and the presentation track for industry or experience talks.

Alberto Avritzer in his keynote, “Integrating Automated Scalability Assessment into DevOps”, discusses about ensuring system performance in DevOps (development and operation). DevOps is an emerging software engineering paradigm that aims for fast feedback cycles between software changes in development and bringing these changes into production. The increased development speed and complexity impose various challenges to scalability assessment. The goal of the keynote is to provide an overview of challenges and approaches for scalability assessment in the context of DevOps and microservices. Specifically, he presents scalability assessment approaches that employ operational data obtained from production-level application performance management (APM) tools, giving access to operational workload profiles and architectural information. In his talk, Alberto discusses the use of such data to automatically create and configure scalability assessments based on models

and load tests. The focus of this keynote is on approaches that employ production usage, because these approaches provide more accurate recommendations for microservice architecture scalability assessment than approaches that do not consider production usage.

Carey Williamson in his keynote, “Things That Go Bump in the Ne”, discusses that we live in a world in which we are very dependent on the Internet for our personal and professional activities. This dependence on network-based applications, and Internet connectivity, makes us vulnerable when things go wrong, or do not behave as expected. Furthermore, diagnosing network-related performance problems is often a challenge, with some of the reasons being data volume, geo-distributed applications, transient behaviours, and encryption. This talk shares some of his experiences on the benchmarking, performance evaluation, and debugging of network-based applications over the past decade. Examples will include Web, video streaming, online social network, email, and LMS applications.

Alexandru Iosup in his keynote, “Will It Rain Today? Understanding the Weather of Computing Clouds, Before it Happens”, discusses that Cloud computing services play an important role in today’s modern society. Combining technology developed since the 1960s (e.g., modes of resource sharing) with new paradigms that could only have emerged in the 2010s (e.g., FaaS), they promise to enable unprecedented efficiency and seamless access to services for many. However successful, we cannot take the cloud for granted: its core does not yet rely on sound principles of science and design, its engineering is often based on hacking, and there have already been worrying signs of unstable operation. In this talk, he posits that we can address the current challenges by focusing on the relatively large complex of systems, and by increasing and focusing the effort put into load testing and benchmarking. He shows examples of how our approach could work in practice, presenting (i) results related to performance variability, (ii) discovery methods that feed into the engineering of future load testing and benchmarking frameworks, and (iii) processes that could improve the reproducibility and credibility of experimental results in this field.

One research paper was accepted after reviewing by members of the program committee. Dheeraj Chahal and Ravi Ojha in their paper “Migrating a Recommendation System to Cloud Using ML Workflow” discuss on-premise implementation of our recommendation system called iPrescribe. They show a methodology to migrate on-premise implementation of recommendation system to a cloud using ML workflow. They also present their study on performance of recommendation system model when deployed on different types of virtual instances.

The list of industry or experience talks is not finalize at the time of writing yet and could be found on the workshop home page <http://ltb2020.eecs.yorku.ca>.

3 ORGANIZERS

Alexander Podelko is a consulting member of technical staff at Oracle, responsible for performance testing and optimization of Enterprise Performance Management (a.k.a. Hyperion) products. He periodically talks and writes about performance-related topics, advocating tearing down silo walls between different groups of performance professionals. He currently serves as a director

for the Computer Measurement Group (CMG), an organization of performance and capacity professionals. He received his PhD in Computer Science from Gubkin University and his MBA from Bellevue University.

Tse-Hsun (Peter) Chen is an Assistant Professor in the Department of Computer Science and Software Engineering at Concordia University, Montreal, Canada. He leads the Software Performance, Analysis, and Reliability (SPEAR) Lab, which focuses on conducting research on performance engineering, program analysis, log analysis, production debugging, and mining software repositories. His work has been published in flagship conferences and journals such as ICSE, FSE, TSE, EMSE, and MSR. He serves regularly as a program committee member of international conferences in the field of software engineering, such as ASE, ICSME, SANER, and ICPC, and he is a regular reviewer for software engineering journals such as JSS, EMSE, and TSE. Dr. Chen obtained his BSc from the University of British Columbia, and MSc and PhD from Queen’s University. Besides his academic career, Dr. Chen also worked as a software performance engineer at BlackBerry for over four years. Early tools developed by Dr. Chen were integrated into industrial practice for ensuring the quality of large-scale enterprise systems. More information at: <http://petertsehsun.github.io>.

Hamzeh Khazaei is an assistant professor of Computer Science in the Department of Electrical Engineering and and Computer Science at York University. Previously he was an assistant professor at the University of Alberta, research associate at the University of Toronto and a research scientist at IBM, respectively. He received his PhD degree in Computer Science from the University of Manitoba where he extended queuing theory and stochastic processes to accurately model the performance and availability of cloud computing systems. More information at <https://www.eecs.yorku.ca/~hkh>.

4 PROGRAM COMMITTEE

Eduardo Cunha deAlmeida, UFPR, Brazil; Nasim Beigi, University of Toronto, Canada; Cor-Paul Bezemer, University of Alberta, Canada; Andreas Brunnert, RETIT GmbH, Germany; Ahmed E. Hassan, Queens University, Canada; Zhen Ming (Jack) Jiang, York University, Canada; Klaus-Dieter Lange, HPE, USA; Marin Litoiu, York University, Canada; Diwakar Krishnamurthy, University of Calgary, Canada; Johannes Kroß, Fortiss, Germany; Gerson Sunyá, University of Nantes, France; Anthony Ventresque, University College Dublin (UCD), Ireland.

5 STEERING COMMITTEE MEMBERS

Ahmed E. Hassan, Queen’s University, Canada; Marin Litoiu, York University, Canada; Zhen Ming (Jack) Jiang, York University, Canada.

ACKNOWLEDGMENTS

We would like to acknowledge the authors who submitted their research to the workshop and the keynote speakers. We would also like to thank the organizers of the ICPE 2020 conference, the steering committee, and the program committee members who contributed to putting together the exciting program.