

Methodological Principles for Reproducible Performance Evaluation in Cloud Computing

Alessandro V. Papadopoulos,¹ Laurens Versluis,² André Bauer,³ Nikolas Herbst,³ Jóakim von Kistowski,³ Ahmed Ali-Eldin,⁴ Cristina L. Abad,⁵ José Nelson Amaral,⁶ Petr Tůma,⁷ Alexandru Iosup²

Abstract: In this work, we report on our recent article on “Methodological Principles for Reproducible Performance Evaluation in Cloud Computing” published in the IEEE Transactions on Software Engineering (TSE)[Pa19]. This publication is a result of a joint effort⁸ coordinated by the Cloud Research Working Group of SPEC (Standard Performance Evaluation Corporation).

Keywords: IEEE Keywords: Cloud computing, Performance evaluation, Benchmark testing, Systematics, Computer performance, Software engineering

Author Keywords: Experimental evaluation, observation study, experimentation

Summary

The rapid adoption and the diversification of cloud computing technology exacerbate the importance of a sound experimental methodology for this domain. It has been a decade since the first commercial cloud has opened for a general public (Amazon AWS, in 2007). Cloud computing systems are now in much demand, and command significant industrial and scientific interest. Although scientific progress and industry growth depend on using sound principles for measuring and reporting cloud-system performance, this process remains complex.

This work investigates how to measure and report performance in the cloud, and how well the cloud research community is already doing it. There are several difficulties, including controlling the time-accuracy trade-off, deciding which data to collect and to report, and using the right summarization statistics in the report.

¹ Mälardalen University, Sweden, Email: alessandro.papadopoulos@mdh.se

² Vrije Universiteit Amsterdam, The Netherlands

³ University of Würzburg, Germany

⁴ Umeå University, Sweden

⁵ Escuela Superior Politecnica del Litoral, Ecuador

⁶ University of Alberta, Canada

⁷ Charles University, Czech Republic

⁸ <https://research.spec.org/news/single-view/article/technical-report-on-reproducible-performance-evaluation-in-cloud-computing-published.html>

We propose a set of eight important methodological principles that combine best-practices from nearby fields with concepts applicable only to clouds, and with new ideas about the time-accuracy trade-off.

We show how these principles are applicable using a practical use-case experiment. To this end, we analyze the ability of the newly released SPEC Cloud IaaS benchmark to follow the principles, and showcase real-world experimental studies in common cloud environments that meet the principles.

Last, we report on a systematic literature review including top conferences and journals in the field, from 2012 to 2017, analyzing if the practice of reporting cloud performance measurements follows the proposed eight principles. Worryingly, this systematic survey and the subsequent two-round human reviews, reveal that few of the published studies follow the eight experimental principles. We conclude that, although these important principles are simple and basic, the cloud community is yet to adopt them broadly to deliver sound measurement of cloud environments.




Fig. 1: Presence of selected principles in the cloud performance evaluation.

Bibliography

[Pa19] Papadopoulos, A. V.; Versluis, L.; Bauer, A.; Herbst, N.; Von Kistowski, J.; Ali-eldin, A.; Abad, C.; Amaral, J. N.; Tüma, P.; Iosup, A.: Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. *IEEE Transactions on Software Engineering*, pp. 1–1, 2019.