

I Aver: Providing Declarative Experiment Specifications Facilitates the Evaluation of Computer Systems Research

Ivo Jimenez and Carlos Maltzahn
UC Santa Cruz

Jay Lofstead
Sandia National Laboratories

Adam Moody and Kathryn Mohror
Lawrence Livermore National Laboratory

Remzi Arpaci-Dusseau and Andrea Arpaci-Dusseau
University of Wisconsin-Madison

ABSTRACT

Validating experimental results in the field of computer systems is a challenging task, mainly due to the many changes in software and hardware that computational environments go through. Determining if an experiment is reproducible entails two separate tasks: re-executing the experiment and validating the results. Existing reproducibility efforts have focused on the former, envisioning techniques and infrastructures that make it easier to re-execute an experiment. By focusing on the latter and analyzing the validation workflow that an experiment re-executioner goes through, we notice that validating results is done on the basis of experiment design and high-level goals, rather than exact quantitative metrics.

Based on this insight, we introduce a declarative format for describing the high-level components of an experiment, as well as a language for specifying generic, testable statements that serve as the basis for validation [1,2]. Our language allows to express and validate statements on top of metrics gathered at runtime. We demonstrate the feasibility of this approach by taking an experiment from an already published article and obtain the corresponding experiment specification. We show that, if we had this specification in the first place, validating the original findings would be an almost entirely automated task. If we contrast this with the current state of our practice, where it takes days or weeks (if successful) to reproduce results, we see how making experiment specifications available as part of a publication or as addendum to experimental results can significantly aid in the validation of computer systems research.

Acknowledgements: Work performed under auspices of US DOE by LLNL contract DE-AC52-07NA27344 ABS-684863 and by SNL contract DE-AC04-94AL85000.

BODY

Providing declarative statements that describe the outcome of an experiment can significantly improve the task of validating its results.

REFERENCES

- [1] I. Jimenez, C. Maltzahn, J. Lofstead, A. Moody, K. Mohror, R. Arpaci-Dusseau, and A. Arpaci-Dusseau, "Tackling the reproducibility problem in storage systems research with declarative experiment specifications," *Proceedings of the 10th parallel data storage workshop*, New York, NY, USA: ACM, 2015, pp. 25–30. Available at: <http://doi.acm.org/10.1145/2834976.2834979>.
- [2] I. Jimenez, "Aver," 2015. Available at: <https://github.com/ivotron/aver>.

Volume 4 of Tiny Transactions on Computer Science

This content is released under the Creative Commons Attribution-NonCommercial ShareAlike License. Permission to make digital or hard copies of all or part of this work 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. CC BY-NC-SA 3.0: <http://creativecommons.org/licenses/by-nc-sa/3.0/>.