

Rigorous Evaluation of Computer Processors with Statistical Model Checking

Filip Mazurek¹, Arya Tschand¹, Yu Wang², Miroslav Pajic¹, and Daniel J. Sorin¹

¹Duke University
²University of Florida

October 28, 2023

1 Abstract

Experiments with computer processors must account for the inherent variability in executions. Prior work has shown that real systems exhibit variability, and random effects must be injected into simulators to account for it. Thus, we can run multiple executions of a given benchmark and generate a distribution of results. Prior work uses standard statistical techniques that are not suitable. While the result distributions may take any forms that are unknown *a priori*, many works naively assume they are Gaussian, which can be far from the truth. To allow rigorous evaluation for arbitrary result distributions, we introduce statistical model checking (SMC) to the world of computer architecture. SMC is a statistical technique that is used in research communities that depend heavily on statistical guarantees. SMC provides a rigorous mathematical methodology that employs experimental sampling for probabilistic evaluation of properties of interest, such that one can determine with a desired confidence whether a property (e.g., System X is 1.1x faster than System Y) is true or not. SMC alone is not enough for computer architects to draw conclusions based on their data. We create an end-to-end framework called SMC for Processor Analysis (SPA) which utilizes SMC techniques to provide insightful conclusions given experimental data.

2 Disclaimers

This work is under preparation. Full text will be released through ACM publishing in the 56th Proceedings of the IEEE/ACM International Symposium on Microarchitecture (MICRO), 2023.

3 Links to Code

<https://github.com/filipmazurek/spa>