

Keynote Address

Proprietary Code to Non-Proprietary Benchmarks: Synthesis Techniques for Scalable Benchmarks

Lizy Kurian John

Electrical and Computer Engineering Department
The University of Texas at Austin
ljohn@ece.utexas.edu

ABSTRACT

Real world applications constitute intellectual property and simultaneous design of hardware and software is made very difficult due to the need for disclosing proprietary software to hardware designers. Consider a smart phone for which applications are developed by various third parties or a military system, where the classified applications are developed in-house at the military while hardware is procured from standard vendors. Design of hardware that gives good performance and low power can be done if hardware designers had access to the software, so they can understand the features of the software and tune various hardware features to the software characteristics. While non-disclosure agreements and legal arrangements can be used to partly solve the problem, it will be much more convenient to have a mechanism to create proxies of proprietary benchmarks that have the performance (and power) characteristics of the source, but not the functionality.

In our past research, we created a benchmark synthesis process for early design exploration. The benchmark synthesis process consists of constructing a proxy workload that possesses approximately the same performance and power characteristics as the original workload [1-3]. The synthesis comprises of two steps: (1) profiling the real-world proprietary workload to measure its inherent behavior characteristics, and (2) modeling the measured workload attributes into a synthetic benchmark program. The set of workload characteristics can be thought of as a signature that uniquely describes the workload's inherent behavior, independent of the microarchitecture. The cloned code in fact has no functionality and cannot be reverse engineered to create the original code or algorithms. The cloned software can be freely released to hardware developers so that they can optimize the hardware they deliver to their clients to yield improved performance.

This talk will describe the benchmark synthesis process as well as the status of the research. The process of

constructing synthetic proxies that approximately resemble the original proprietary application will be explained. Synthesis for multicore and multithreaded processors will be described.

Categories and Subject Descriptors: B.8.2
[Hardware]: Performance Analysis and Design Aids

General Terms: Performance, Measurement

Keywords: Workload Synthesis, Synthetic Benchmarks, Computer Design.

REFERENCES:

- [1] Robert H. Bell, Jr. and Lizy K. John, "Improved Automatic Test case Synthesis for Performance Model Validation", *Proceedings of 19th ACM International Conference on Supercomputing (ICS)*, June 2005, pp. 111-120.
- [2] Robert H. Bell, Rajiv R. Bhatia, Lizy John, Jeff Stuecheli, Ravel Thai, John Griswell, Paul Tu, Louis Capps, Anton Blanchard, "Automatic Test-case Synthesis and Performance Model Validation for High Performance PowerPC Processors", *Proceedings of the International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2006, pp. 154-165.
- [3] Ajay Joshi, Lieven Eeckhout, Robert H. Bell Jr., and Lizy K. John, "Performance Cloning: A Technique for Disseminating Proprietary Applications as Benchmarks", *Proceedings of the International Symposium on Workload Characterization, October 2006*, pp. 105-115
- [4] Karthik Ganesan, Jungho Jo, and Lizy K. John, "Synthesizing Memory-Level Parallelism Aware Miniature Clones for SPEC CPU2006 and ImplantBench Workloads", *2010 International Symposium on Performance Analysis of Systems and Software*, March 2010