Learning-Based Analytical Power and Performance Modeling

By: Andreas Gerstlauer
http://users.ece.utexas.edu/~gerstl/

Bio

Andreas Gerstlauer is an Associate Professor in Electrical and Computer Engineering at The University of Texas at Austin. He received his Ph.D. in Information and Computer Science from the University of California, Irvine (UCI) in 2004. Prior to joining UT Austin in 2008, he was an Assistant Researcher in the Center for Embedded Computer Systems (CECS) at UC Irvine, leading a research group to develop electronic system-level (ESL) design tools. Commercial derivatives of such tools are in use at the Japanese Aerospace Exploration Agency (JAXA), NEC Toshiba Space Systems and others. Dr. Gerstlauer is co-author on 3 books and more than 70 publications. He has presented in numerous conference and industrial tutorials, is an editor for ACM TECS and TODAES journals, and has served as the topic, track or program chair of major international conferences such as DAC, DATE, ICCAD and CODES+ISSS. His research interests include system-level design automation, system modeling, design languages and methodologies, and embedded hardware and software synthesis.

Abstract

Next to performance, early power and energy estimation is a key challenge in the design of systems-on-chip (SoCs) today. Traditional simulation-based methods are often too slow while existing analytical models are often not sufficiently accurate. In this talk, I will present our recent work on applying advanced machine learning techniques to synthesize analytical models that can accurately predict the power and performance of various hardware and software components using measurements obtained from fast, source- or transaction-level functional-only simulations on a host. We have developed such approaches for both hardware IP components as well as software running on programmable processors of a SoC. On the hardware side, learning based models for white-box and black-box IPs reach simulation speeds of 1 Mcycles/s at 97% accuracy. On the software side, depending on the granularity at which prediction is performed, learning-based models can achieve more than 95% accuracy at more than 1000 MIPS of equivalent simulation throughput. Instead of directly employing learning models for cross-platform prediction at the source, we have also used learning-based methods to calibrate traditional low-level analytical reference models, such as McPAT, against post-silicon measurements. Using target-specific activity data, such as counter measurements, automatically calibrated McPAT models are thereby able to predict transient power to within 5% of a modern ARM SoC.

This work was performed in collaboration with Prof. Lizy K. John, the Semiconductor Research Corporation (SRC) and Intel Labs.