Dynamic Voltage and Frequency Scaling in NoCs With Supervised and Reinforcement Learning Techniques

Quintin Fettes, Ohio University
Mark Clark, Belcan
Razvan Bunescu and Avinash Karanth, Ohio University
Ahmed Louri, The George Washington University

This installment of Computer's *series highlighting the work published in IEEE Computer Society journals comes from* IEEE Transactions on Computers.

Digital Object Identifier 10.1109/MC.2019.2923827
Date of publication: 27 August 2019

Computer architects often face the challenging task of balancing various design considerations, such as performance, power, cost, and reliability. With its unprecedented success in numerous domains and disciplines, machine learning could be a promising approach to solving complicated architecture design and optimization problems. *IEEE Transactions on Computers* continues to lead research in this area, recently publishing more than 10 papers that propose innovative, viable, and promising ways to take the advantage of machine learning in computer architecture.

In the recent study "Dynamic Voltage and Frequency Scaling in NoCs With Supervised and Reinforcement Learning Techniques" (*IEEE Transactions on Computers*, vol. 68, no. 3, pp. 375–389, 2019), Quintin

0018-9162/19©2019IEEE

Fettes, Mark Clark, Razvan Bunescu, Avinash Karanth, and Ahmed Louri propose learning-enabled, energy-aware dynamic (LEAD) voltage/frequency (VF) scaling for multicore architectures, using supervised learning and reinforcement learning (RL) approaches.

Applying dynamic voltage/frequency scaling (DVFS) to a network-on-chip (NoC) can be an effective technique to reduce dynamic energy, but the decrease occurs at the expense of NoC performance as well as the execution of applications running on multicores. To tackle this problem more smartly, the LEAD model groups an NoC router and its outgoing links into the same VF domain and implements proactive DVFS mode management strategies that rely on offline, trained machine-learning models to provide optimal VF mode selection. The authors present three supervised learning versions of LEAD systems that are based on buffer utilization, a change in buffer utilization, and a change in energy/throughput to allow proactive mode selection based on accurate predictions of future network parameters. This produces an average dynamic energy savings of 15.4% when evaluated using the Princeton Application Repository for Shared Memory Computers and SPLASH-2 benchmarks on a four-by-four concentrated mesh architecture, with merely an 0.8% loss in throughput and no significant impact on latency.

To increase the efficacy, the authors explore the use of modern RL techniques, such as deep Q-networks, noisy networks, replay buffers, and prioritized replay, to select DVFS modes more effectively. This approach, diagrammed in Figure 1, is more scalable than other supervised learning models and does not require expensive threshold tuning, thus, allowing for an easier adjustment of dynamic-energy-versus-throughput tradeoffs. All models are trained offline to minimize their energy consumption and area footprint at runtime. LEAD-RL increases the average dynamic energy savings to 20.3% at the cost of a 1.5–1.7% decrease in throughput and latency.

Ultimately, the more flexible RL approach enables the learning of an optimal behavior for a wider range of load environments, under any desired energy-versus-throughput tradeoffs, and without much human engineering in the automatic training process.

IEEE Transactions on Computers will closely follow these exciting developments in applying machine learning to computer architectures and provide the latest academic and industry research. Please stay tuned for upcoming issues to keep current on this topic as well as computing architectures, memory technologies, real-time systems, and much more. C

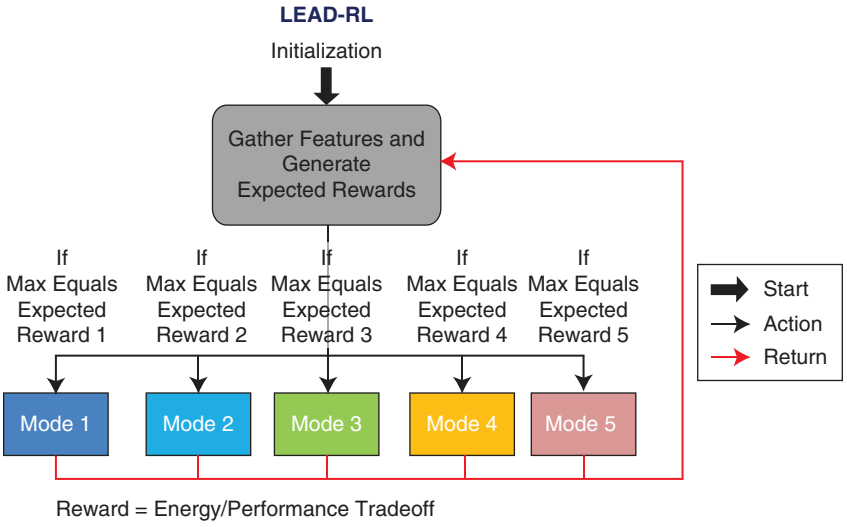

FIGURE 1. LEAD-RL generates long-term, expected rewards for each mode and selects the mode that is predicted to lead to the greatest return during the next time window.

QUINTIN FETTES is a Ph.D. student in computer science at Ohio University, Athens. Contact him at qf731413@ohio.edu.

MARK CLARK is a software engineer at Belcan, Cincinnati, Ohio. Contact him at mc591611@ohio.edu.

RAZVAN BUNESCU is a professor in the School of Electrical Engineering and Computer Science at Ohio University, Athens. Contact him at bunescu@ohio.edu.

AVINASH KARANTH is the Joseph K. Jachinowski Professor in the School of Electrical Engineering and Computer Science at Ohio University, Athens. Contact him at karanth@ohio.edu.

AHMED LOURI is the David and Marilyn Karlgaard Endowed Chair Professor of Electrical and Computer Engineering at The George Washington University, Washington, D.C. Contact him at louri@gwu.edu.