Department of Computer Science
St. Francis Xavier University
presents

Studying Practical Kernel-Level DVFS Techniques for Deadline-Constrained Periodic Applications

by

Zhou Ti
St. Francis Xavier University
M.Sc. Thesis Proposal Presentation
Time: Aug 9, 2022 10:00 AM Halifax
Join Zoom Meeting
https://us02web.zoom.us/j/89851044443

Dynamic voltage/frequency scaling (DVFS) is a widely used technique in modern computers to save energy for tasks by slowing down the speed of the tasks. Many research works have studied how to save energy for real-time systems through DVFS techniques. However, these methods need a detailed specification of the tasks and are not practical due to imprecise or pessimistic WCET estimation (usually coupled with machine configuration) and the complexity of deploying the DVFS to work with the OS scheduler. Thus, the current practice for DVFS-based energy-saving primarily relies on the built-in DVFS governors in Linux devices to save energy for tasks without prior knowledge about the tasks. However, the effectiveness of the Linux-built-in governors is limited for energy saving for real-time systems.
The purpose of this thesis is to study the DVFS for a class of real-time systems with a deadline, commonly seen in many IoT systems or pervasive control systems.

In this proposal, we identify some patterns that cannot be dealt with by Linux-built-in governors, even after configuration, through low-cost kernel-level profiling. To avoid the detailed modeling of the task and machine and overcome the limitation of built-in governors, deep reinforcement learning is proposed for CPU power management for soft-deadline periodic real-time systems. Extensive experiments were performed for various applications and deadlines on Nvidia Jetson Nano Board with measured board-level energy consumption. The results show that the proposed method can learn a DVFS policy that can save substantially more energy consumption than built-in governors. The learned DVFS controller is implemented as a module with a quantization technique in the kernel, and it is lightweight with low overhead in decision making. The effect of the learned DVFS controller can be visualized through our kernel-level profiling tool. This work provides a starting point and tools for studying practical kernel-level DVFS for real-time systems.