CNS-0931035: Collaborative Research: Abstraction of Cyber-Physical Interplays and Its Application to CPS Design

K. G. Shin and E. M. Atkins, University of Michigan, Ann Arbor
C. M. Krishna and I. Koren, University of Massachusetts, Amherst

Abstract

The objective of this project is to develop mechanisms to better bridge the gap between the computer and plant sides of a cyber-physical system. Some key contributions are the following:

Cost Function for Automobiles: The execution time delay in a control computer translates to feedback delay from the point of view of the application, and affects the quality of the provided control. For zero order hold systems, this degradation in the quality of control can be quantified as a cost function of the feedback delay and the rate at which the control inputs are updated. Based on a detailed automobile model, we have developed cost functions for wheel torque and steering. Furthermore, we have constructed an algorithm to schedule tasks that aims not only to meet all hard deadlines but also to minimize the total cost, as expressed by the sum of all the task cost function values.

Thermal-aware Scheduling of Real-Time Tasks: A real-time controller must meet hard deadlines to keep the feedback delay sufficiently low. It must also ensure that the processor is not subjected to excessive heating as thermal stress contributes to processor failure. We use a per-core dynamic voltage scheduler as well as job migration to balance these competing goals by evening out the temperature variance across multiple cores.

Managing Battery-Supercapacitor Resources: Batteries have high energy densities while supercapacitors are more amenable to satisfying impulse energy draws. This suggests the use of supercapacitors as an intermediary between the backup energy store (the battery) and the consuming processor. The supercapacitor charging schedule significantly affects the usable energy that can be drawn from the battery. We have developed a model that allows one to schedule appropriate levels of battery energy draws.

Architecture Adaptation for Cyber-Physical Applications: We have shown that architecture adaptation when combined with dynamic voltage scaling provides significant advantages over dynamic voltage scaling alone, in the context of a real-time workload. In particular, in many cases, we achieve 25% additional savings of energy over voltage scaling alone without missing any deadlines. With the increasing prevalence of complex multicore processors, architecture adaptation has been demonstrated to be a viable avenue to further energy savings.

Real-Time Task Scheduling With Tolerance of Wear-out Faults on Multi-Core Platforms: As CMOS shrinks down into nanoscale, wear-out faults become more likely to occur in multi-core chips, making their lifetime shorter and less predictable than ever before. For safety-critical applications running on multi-core chips (e.g., automotive powertrain controls) to operate correctly, both functional and timing constraints must be met, even in the event of core wear-out faults. We have developed a run-time scheduling algorithm that utilizes the computing resources of a multi-core processor efficiently to meet both fault-tolerance and real-time constraints of applications. The algorithm attempts to use job migration whenever possible, along with a proactive fault-detection mechanism, to tolerate failures caused by core wear-out faults with minimum redundancy. At run-time, the scheduler chooses one of the pre-defined fault-tolerance policies with consideration of all task timing constraints and the job-migration cost.