WATERS'15

6th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems
Tuesday, July 7, 2015

Keynote speaker: James H. Anderson, University of North Carolina at Chapel Hill
Having Fun Experimenting with Hardware Management and Mixed Criticality on Multicore

Abstract

In work on real-time multicore systems, the most important unresolved issue today is a problem termed here "*the one-out-of-m problem*": when using an *m*-core platform in a safety-critical domain, analysis pessimism can be so great, the capacity of the "additional" *m* - 1 cores is entirely negated. Both mixed-criticality analysis techniques (which can be applied to provision less critical system components less pessimistically) and hardware-management techniques (which can be applied to reduce pessimism due to shared-hardware interference) can be seen as attempts to address this problem. In the past, these approaches have been largely considered separately. In this talk, experimental results will be presented that show that applying both approaches together may be a more effective way forward.

These results were obtained through work involving UNC's MC^2 (mixed-criticality on multicore) framework, which has been the subject of experimental research for over six years. Its implemented features have continued to evolve during this time. The current version supports hardware-management techniques that allow sharing and isolation tradeoffs to be investigated in a mixed-criticality setting with respect to both last-level caches and DRAM memory banks. These techniques can be applied to isolate highly critical tasks from both less critical tasks and the operating system.

The presented experiments set the stage for an obvious next step in MC^2-related research: conducting a large-scale schedulability study to fully understand relevant resource-allocation tradeoffs pertaining to the deployment of mixed-criticality task systems. However, the scale of such a study would be massive, unless some study parameters (such as the distribution of tasks across criticality levels) can be tightly constrained based on practical use-case scenarios. The talk will end with a discussion of how such a study might be conducted and a plea for advice concerning how to reasonably constrain its scale.