

## **Three Aspects of Real-Time Multiprocessor Scheduling: Timeliness, Fault Tolerance, Mixed Criticality**

Risat Mahmud Pathan

*Department of Computer Science and Engineering  
Chalmers University of Technology, Sweden*

### **Abstract**

The design of real-time systems faces two important challenges: incorporating more functions/services on existing hardware to make the system more attractive to the market, and deploying existing software on multiprocessors (e.g., multicore) to utilize more processing power. Adding more services on the same hardware needs efficient resource utilization. In addition, satisfying the real-time constraints, while at the same time efficiently utilizing the multiprocessor platform, is a challenging problem. This thesis deals with *global multiprocessor scheduling* for real-time systems, that is, the *fixed-priority* scheduling of sporadic tasks, where each task is allowed to run on any processor.

More specifically, this thesis considers *three* aspects of the design and analysis of global scheduling algorithms: timeliness, fault tolerance, and mixed criticality. *Timeliness* is about meeting the deadlines of the tasks; *fault tolerance* is about producing the correct output within the deadline even in the presence of faults; and *mixed criticality* is about facilitating the certification of system when tasks having different criticality (or importance) are hosted on a common computing platform.

With respect to the timeliness aspect, global multiprocessor scheduling is analyzed (by assuming no faults and the same criticality for all the tasks) in order to propose new fixed-priority assignment policies and efficient schedulability tests. The proposed schedulability tests are shown to not only dominate (from a theoretical point of view) but also significantly outperform (by using simulation experiments) the state-of-the-art schedulability tests for global fixed-priority scheduling.

To allow for the combination of fault tolerance and timeliness, new scheduling algorithms that use time redundancy (i.e., execution of backup task) to tolerate multiple hardware and software faults are proposed. To account for the potential intrusive effect of time-redundant execution of backup tasks on the capability to meet task deadlines, new efficient schedulability tests for the proposed algorithms are derived. If a task set satisfies the schedulability tests, then all the task deadlines are met even when multiple faults (restricted by the assumed fault model) are to be tolerated using time redundancy.

To allow mixed-criticality tasks to be hosted on the same multiprocessor platform, a new algorithm for fixed-priority scheduling is proposed. The purpose of the algorithm is to facilitate certification, while at the same time efficiently utilizing the processing platform. A schedulability test for the algorithm can determine whether the appropriate level of assurance, according to the requirement of some certification authority/standard for meeting the deadlines of the mixed-criticality tasks, is guaranteed or not.

**Keywords:** Real-Time Systems, Sporadic Tasks, Fixed Priority, Global Multiprocessor Scheduling, Time Redundancy, Fault-Tolerant Scheduling, Mixed-Criticality Systems