Influence of the Task Model on the Precision of Scheduling Analysis for Preemptive Systems

Sebastian Altmeyer, Claire Maiza
Saarland University, Germany
{altmeyer,maiza}@cs.uni-saarland.de

In real-time systems, tasks must obey stringent timing constraints. A verification process that checks if these constraints are met consists of a timing analysis of each task and schedulabiltiy analysis of the set of tasks. The interface between these two analyses is the task model constituting an abstraction of the task's timing properties in the system. A very basic task model was presented by Liu and Layland [4]: the execution demand of a task i, often denoted as C_i, abstracts all possible execution times to a single value. The aim of the timing analysis is to compute this abstraction of the timing behavior of the tasks by safely bounding their worst-case execution time (C_i).

By the abstraction step from timing analysis to task model some precision is lost. Especially in preemptive systems or systems with interrupts, timing analysis computes, in addition to the pure time bound for uninterrupted execution, the additional delay due to interrupts or preemptions. As research on this topic has shown [1], preemption costs strongly depend on the specific preemption points and on the preempting task; preemption costs may vary from nearly zero to large fractions of the task's execution time. Thus, timing analysis may compute not only an upper bound on the preemption costs for a task i but also additional bounds for preemption of task i by task j [2, 6], or for the n^{th} preemption of task i, or for preemption occurring at point p [1, 3]. If a schedulability analysis is able to take into account such precise information about the preemption costs, the results may exhibit a higher precision.

However, schedulability analyses are often based either on the basic task model by Liu and Layland with a unified bound on the execution time including preemption costs or on a model using only one separated value for the preemption costs per task [5]. The second task model improves over Liu and Layland's model by distinguishing preemption costs depending on the actual number of preemptions instead of considering an upper bound. Nevertheless, both models exhibit an inherent pessimism. The bound on the additional preemption delay—no matter if part of the execution time bound or considered separately—must comprise all possible preemption scenarios regarding preemptions points, preempting task etc., even if they do not occur in the actual schedule. However, schedulability analysis generally uses a simple model like the one of Liu and Layland to reduce the complexity of the schedulability test.

So, on the one hand, the abstraction of the timing behavior in the task model comes at the cost of inherent pessimism and on the other hand, schedulability analysis may rely on a simplified task model to reduce complexity. The tradeoff between precision and complexity of the schedulability analysis is determined by the task model and its abstraction of the timing of tasks. This tradeoff raises some questions:

- How high is the inherent pessimism and imprecision of a specific task model due to the precision of the abstraction of the timing properties?
- What is a good tradeoff between precision of the task model and complexity of the schedulability analysis?
- How to integrate such detailed information about the timing of tasks in the schedulability analysis?

Furthermore, incorporating more precise information, such as preemption points or preempting task, may be infeasible in general; thus, the following question arises in this context:

- Is it possible to adapt the schedule or the system in order to better benefit from the precision provided by timing analysis?

A typical example in which the system is adapted to achieve higher precision is the use of deferred preemption, i.e., preemption limited to predefined program points. Although flexibility of the schedule is lost to some degree, schedulability may be achieved only due to a strongly reduced bound on the preemption costs.

Answers to the questions listed above enable a better understanding of the influence of preemptions costs on the schedule and provide guidelines for the design and schedulability analysis of real-time systems. Note that the tradeoff between precision of the task model and complexity of the schedulability analysis becomes especially apparent in case of preemptive systems—but is not limited to such. Other issues where more precision from timing analysis could be taken into account are for instance the cache contention for multicore systems or the difference between first and all further executions of one task, which often strongly varies due different initial cache states.

References

[1] S. Altmeyer and C. Burguière. A new notion of useful cache block to improve the bounds of cache-related preemption delay. In *ECRTS 2009*.

[2] S. Altmeyer, C. Burguière, and J. Reineke. Resilience analysis: Refinement of the CRPD bound for set associative caches. In *LCTES 2010*.

[3] C.-G. Lee, J. Hahn, S. L. Min, R. Ha, S. Hong, C. Y. Park, M. Lee, and C. S. Kim. Analysis of cache-related preemption delay in fixed-priority preemptive scheduling. In *RTSS 1996*.

[4] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1), 1973.

[5] P. Meumeu Yomsi and Y. Sorel. Extending rate monotonic analysis with exact cost of preemptions for hard real-time systems. In *ECRTS 2007*.

[6] J. Staschulat, S. Schliecker, and R. Ernst. Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In *ECRTS 2005*.