

Fixed Priorities or EDF for Distributed Real-Time Systems?

Juan M. Rivas
Computers and Real-Time Group
Universidad de Cantabria
39005-Santander, SPAIN
rivasjm@unican.es

J. Javier Gutiérrez
Computers and Real-Time Group
Universidad de Cantabria
39005-Santander, SPAIN
gutierjj@unican.es

Michael González Harbour
Computers and Real-Time Group
Universidad de Cantabria
39005-Santander, SPAIN
mgh@unican.es

EXTENDED ABSTRACT¹

Although fixed priority (FP) scheduling is the most popular on-line scheduling policy in industrial environments, the earliest deadline first (EDF) policy is starting to get more attention, given its benefits in terms of increased resource usage. A comparison between FP and EDF under several aspects is presented in [1] for single-processor systems, showing that FP does not always have the commonly attributed advantages over EDF.

The starting point of this work is MAST², a suite of tools developed in our group that offers different schedulability analysis and scheduling parameter assignment techniques for event-driven systems. MAST is based on the end-to-end flow model proposed in the MARTE standard [2] and integrates techniques for FP, EDF, or heterogeneous FP/EDF [3] (e.g., EDF processors and a CAN bus), which can be applied to distributed systems. For EDF it supports both local schedulers, where task scheduling deadlines are referenced to their release times in their own processor, and global schedulers, where scheduling deadlines are referenced to the arrival of the event that releases the end-to-end flow, possibly in a different processor, and thus requiring clock synchronization. MAST is under active development [4], and new techniques can be easily integrated.

Our goal is to compare how FP and EDF can influence the schedulability of a distributed real-time system under a variety of conditions: different system sizes, deadline/period ratios, different lengths of end-to-end flows, etc. To this end, we propose an exhaustive study of current schedulability analysis and scheduling parameter assignment techniques over a wide range of examples, with the purpose of evaluating each algorithm and quantify the differences among them, as well as comparing the ability of the different scheduling algorithms to meet the deadlines.

We need the capability to execute the analysis and assignment techniques over an extensive set of examples. Thus, we are developing a benchmark generator tool for MAST, GEN4MAST, which can (1) generate a pool of MAST examples according to some specifications and evolution rules, (2) run a set of tests with different schedulability analysis or scheduling parameter assignment techniques, and finally (3) store and process the results. The first version of the tool deals only with tasks and processors, but it will be extended to messages and networks, and also to the network switches being included in future versions of MAST.

For the generation of examples, the user can specify the basic characteristics of the systems that will form the pool of examples: the number of end-to-end flows, tasks and processors; the deadline and period ranges and ratios; the type of schedulers (FP and local or global EDF); and the system loads. This allows adapting to the characteristics of a specific application domain. Each example is generated with a certain component of randomness (task allocation or actual deadline or period values), and it is also possible to specify how many examples with a similar architecture can be generated, so statistically relevant results can be obtained. In addition, it is possible to specify other parameters that fine-tune the behavior of the schedulability analysis or the scheduling parameter assignment algorithms (e.g., configuration parameters for the HOSPA algorithm).

For each execution of a test (example-algorithm combination), a set of results is recorded, including the worst-case response times obtained, the execution times required for the calculation and, optionally, the sensitivities given as percentages of system slacks.

Preliminary results using holistic analysis techniques show that global EDF reaches higher average utilizations than FP, which in turn gets higher utilizations than local EDF. The table below shows the maximum average schedulable utilization achieved by each technique for a medium to small size experiment (17 tasks and 3 processors). This experiment took 6 hours of computation to analyze 4200 examples.

	<i>FP</i>	<i>Local EDF</i>	<i>Global EDF</i>
Utilization %	82,6	67,0	95,7

REFERENCES

- [1] G. Buttazzo, "Rate Monotonic vs. EDF: Judgment Day," *Real-Time Systems*, 29(1), pp. 5-26, 2005.
- [2] Object Management Group, "UML Profile for MARTE: Modeling and Analysis of Real-Time Embedded systems," 2009 OMG Document, v1.0 formal/2009-11-02.
- [3] Juan M. Rivas, J.Javier Gutiérrez, J. Carlos Palencia, M. Gonzalez Harbour, "Schedulability Analysis and Optimization of Heterogeneous EDF and FP Distributed Real-Time Systems," Proceedings of the 23th Euromicro Conference on Real-Time Systems, Porto (Portugal), 2011.
- [4] M. González Harbour, J. Javier Gutiérrez, José M. Drake, P. López Martínez and J. C. Palencia, "Modeling distributed real-time systems with MAST 2," In press, *Journal of Systems Architecture*, Elsevier, 2012.

¹ This work has been funded in part by the Spanish Government under grant number TIN2011-28567-C03-02 (HIPARTES).

² MAST is available at <http://mast.unican.es/>