



Modelling and Analysis Suite for Real Time Applications (MAST 1.5.0)

Analysis Techniques used in MAST

By:	Michael González Harbour	mgh@unican.es
	José Carlos Palencia Gutiérrez	palencij@unican.es
	José Javier Gutiérrez García	gutierjj@unican.es
	Juan María Rivas Concepción	rivasjm@unican.es

Copyright © 2000-2014 Universidad de Cantabria, SPAIN

1. Introduction

The MAST toolset contains several schedulability analysis tools capable of analysing single processor and distributed systems scheduled with fixed priority, EDF, and EDF within priorities scheduling policies. The tools are based on different scheduling analysis techniques published in the literature:

- *Classic RM Analysis*. This analysis implements the classic exact response time analysis for single-processor fixed-priority systems first developed by Harter [3] and Joseph and Pandya [8], and later extended by Lehoczky to handle arbitrary deadlines [10] and by Tindell to handle jitter [29]. It corresponds to Technique 5, “Calculating response time with arbitrary deadlines and blocking”, in [9].
- *Varying Priorities Analysis*. This analysis implements the response time analysis for single processor fixed priority systems in which tasks may explicitly change their priorities, developed by González, Klein and Lehoczky [4]. It corresponds to Technique 6, “Calculating response time when priorities vary”, in [9]. In terms of the MAST model, transactions for this tool are built as linear transactions, with a sequence of activities, each representing the execution of an operation with a possibly different priority level specified through a *Permanent_Overriden_Priority* attribute. However, each transaction is limited to having a single segment. A segment is a continuous sequence of activities executed by the same server. Note that all the other fixed priority tools used in MAST require that the priority of an activity is the same at the start and at the end, and thus they do not allow using *Permanent_Overriden_Priority* attributes.
- *EDF Monoprocessor Analysis*. This analysis implements the exact response time analysis for single-processor EDF systems first developed by Spuri [26]. In the MAST implementation we use the EDF Within Priorities (see below), because there may be interrupt service routines (modelled as fixed priority tasks) in addition to the EDF tasks.
- *EDF Within Priorities Analysis*. This analysis is a mixture of the response time analysis for fixed priority systems [9][10][29] and for EDF [26]. It is capable of analysing systems with hierarchical schedulers, in which the underlying primary scheduler is based on fixed priorities, and there may be other EDF (secondary) schedulers scheduling tasks at a given priority level. It was developed by González and Palencia [5].
- *Holistic Analysis*. This analysis extends the response time analysis to multiprocessor and distributed systems. It is not an exact analysis, because it makes the assumption that tasks of the same transaction are independent. It was first developed for fixed priority



systems by Tindell and Clark [29][30] and refined by Palencia et al [14]. For EDF systems, Spuri [27] and later Palencia [15] extended this technique for EDF systems scheduled with global deadlines (referred to a global synchronized clock), and Palencia [32] for systems scheduled with local deadlines (referred to local clocks in each processor).

- *Offset Based Approximate Analysis*. This is a response time analysis for multiprocessor and distributed systems that greatly improves the pessimism of the holistic analysis by taking into account that tasks of the same transaction are not independent, through the use of offsets. Offset based analysis for fixed priorities was first introduced by Tindell [31] and then extended to distributed systems by Palencia and González [15]. It was later extended to EDF systems by Palencia and González [17]. Although it provides much better results than the holistic analysis, it is not an exact method because the exact analysis is intractable. The method approximates the interference of a transaction with a maximum interference function that is independent on the phase of the transaction.
- *Offset Based Approximate with Precedence Relations Analysis*. This is an enhancement of the offset based approximate analysis for fixed priority systems in which the priorities of the tasks of a given transaction are used together with the precedence relations among those tasks to provide a tighter estimation of the response times. It was developed by Palencia and González [16], and later enhanced by Redell [20].
- *Offset Based Slanted Analysis*. This is another enhancement of the offset based approximate analysis for fixed priority systems in which the maximum interference function is defined with a tighter approximation. This method provides better results than the Offset-Based Approximate Analysis, but it is uncertain if it gets better results than the method with precedence relations. In general both techniques should be applied, and the best results used as an upper bound on the response times, as both methods are pessimistic. It was developed by Mäki-Turja [13].
- *Offset Based Brute Force Analysis*. This is an exact analysis of offset-based transactions, initially developed by Tindell [31]. It analyses all possible combinations of tasks initiating the critical instant for each transaction, which leads to a combinatorial problem that only offers results for very small systems.
- *Analysis for heterogeneous distributed systems*. This technique [21] allows the integration of different response-time analysis techniques so that they can be applied to heterogeneous distributed systems with different scheduling policies in each resource. The Holistic and all the Offset-Based analysis techniques are now integrated in MAST into the distributed analysis for heterogeneous systems.

The MAST model is able to support multi-path transactions in which the event flow can be divided into several paths with fork (multicast) or branch (delivery or query server) elements, and paths may be joined with join (barrier) or merge (concentrator) elements. The analysis of these multipath transactions is described in [7] using the holistic analysis technique, and is implemented in MAST 1.5 with the following restrictions: the transaction has a single input event and only fork, join or merge elements are supported; branch is not supported, as the analysis in [7] only works for transactions with deadlines at or before the end of the period. The holistic analysis technique is the most pessimistic, so in the future we plan to extend offset-based analysis to handle these multipath transactions.

The MAST toolset is able to automatically calculate the blocking times caused by mutual exclusion synchronization. The model includes shared resources with the immediate ceiling



[2], priority inheritance [24], and stack resource [2] synchronization protocols. It also allows mixtures of these protocols and their use in multiprocessor systems [18][19], with some restrictions that are described below. For the basic priority inheritance, Rodríguez and García [22] showed that most implementations do not strictly follow the rules in [24], and that the amount of blocking is usually higher than that predicted by the theory. In MAST we take this into account when calculating the blocking times due to the use of the priority inheritance protocol.

The MAST toolset also contains tools to automatically assign priorities and other scheduling parameters. Priority assignment tools are provided for single-processor and distributed systems. In single-processor systems, if deadlines are within the periods the optimum deadline monotonic priority assignment developed by Leung and Layland is used [11]. The Liu and Layland classic rate monotonic priority assignment for the case of deadlines equal to the periods [12] is known to be a special case of the deadline monotonic assignment. When deadlines are larger than the task periods, the optimum priority assignment developed by Audsley is used [1]. This technique is based on the iterative use of the schedulability analysis tools for different solutions, until a schedulable solution is obtained.

In multiprocessor and distributed systems the problem of assigning scheduling parameters is much harder, as there are strong interrelations between the response times in the different resources. For distributed systems scheduled with fixed priorities, we provide two heuristic solutions based on iteratively applying the schedulability analysis tools. The first one is based on the use of the simulated annealing optimization techniques, first used by Tindell, Burns, and Wellings for assigning priorities [28]. The second heuristic, which usually provides better and faster results is the HOPA algorithm developed by Gutiérrez and González [6]. For distributed EDF systems, the HOSDA [32] algorithm, which is an evolution of HOPA, is provided. This algorithm is capable of assigning and optimizing local and global scheduling deadlines. In addition to HOSDA, other simpler mechanisms algorithms are provided for comparison purposes: the Proportional Deadline assignment (PD) [34], and the Normalized Proportional Deadline assignment (NPD) [34], which takes into account the CPU utilization. These techniques can be applied for the assignment of both local [33] and global [32] scheduling deadlines. All these techniques are integrated into HOSPA [21], the scheduling parameters assignment tool provided for heterogeneous FP/EDF systems.

The MAST toolset is able to determine not only whether the system is schedulable or not, but also how far it is from being schedulable, or how much capacity is available until the system becomes unschedulable. It does so by calculating slacks, which are defined as the percentage by which we can increase the execution times of some operations while keeping the system schedulable (for positive slacks) or the percentage by which we have to decrease the execution times to make the system schedulable (for negative slacks). A slack of zero means that the system is just schedulable, and that even the smallest increment in the execution times would lead to non schedulability. There are different kinds of slacks provided:

- *System slack*: affects all the operations in the system
- *Processing resource slack*: affects only the operations executed in a given processing resource
- *Transaction slack*: affects only the operations used in a given transaction
- *Operation slack*; affects only one single operation



The slacks are calculated by modifying the worst-case execution times and repeating the analysis using a binary search to find the point in which the system becomes unschedulable (or schedulable if it wasn't). The slack is calculated with a 1% precision to limit the amount of times the analysis is repeated to around 20 times.

In addition to these tools, some additional analysis techniques were needed to develop the MAST toolset, because it allows combinations of scheduling policies and synchronization protocols that did not have a global treatment in the published analysis techniques that we know about. This document describes these additional extensions that were specifically developed for MAST.

References

- [1] N.C. Audsley, "Optimal Priority Assignment and Feasibility of Static Priority Tasks with Arbitrary Start Times", Dept. Computer Science, University of York, December 1991.
- [2] T.P. Baker, "Stack-Based Scheduling of Realtime Processes", Journal of Real-Time Systems, Volume 3, Issue 1 (March 1991), pp. 67–99.
- [3] P.K. Harter, "Response times in level-structured systems". ACM Transactions on Computer Systems, vol. 5, no. 3, pp. 232-248, Aug. 1984.
- [4] M. González Harbour, M.H. Klein, and J.P. Lehoczky. "Timing Analysis for Fixed-Priority Scheduling of Hard Real-Time Systems". IEEE Trans. on Software Engineering, Vol. 20, No.1, January 1994.
- [5] M. González Harbour and J.C. Palencia, "Response Time Analysis for Tasks Scheduled under EDF within Fixed Priorities", Proceedings of the 24th IEEE Real-Time Systems Symposium, Cancun, México, December, 2003.
- [6] J.J. Gutiérrez García and M. González Harbour, "Optimized Priority Assignment for Tasks and Messages in Distributed Real-Time Systems". Proceedings of the 3rd Workshop on Parallel and Distributed Real-Time Systems, Santa Barbara, California, pp. 124-132, April 1995.
- [7] J. J. Gutiérrez García, J.C. Palencia Gutiérrez and M. González Harbour, "Schedulability Analysis of Distributed Hard Real-Time Systems with Multiple- Event Synchronization". Proceedings of 12th Euromicro Conference on Real-Time Systems, Stockholm (Sweden), IEEE Computer Society Press, pp. 15-24, June 2000.
- [8] M. Joseph and P. Pandya, "Finding Response Times in a Real-Time System". The Computer Journal (British Computing Society) 29, 5, pp. 390-395, October 1986.
- [9] M.H. Klein, T. Ralya, B. Pollak, R. Obenza, and M. González Harbour. "A practitioner's Handbook for Real-Time Analysis". Kluwer Academic Pub., 1993.
- [10] J.P. Lehoczky. "Fixed Priority Scheduling of Periodic Tasks Sets with Arbitrary Deadlines". IEEE Real-Time Systems Symposium, 1990.
- [11] J. Leung, and J.W. Layland. "On Complexity of Fixed-Priority Scheduling of Periodic Real-Time Tasks". Performance Evaluation 2, 237-50, 1982.
- [12] C.L. Liu and J.W. Layland. "Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment". Journal of the ACM, Vol. 20, No. 1, 1973, pp. 46-61.
- [13] Jukka Mäki-Turja and Mikael Nolin. "Efficient implementation of tight response-times for tasks with offsets". Journal of Real-Time Systems. Volume 40 Issue 1, October 2008, pp. 77 - 116.
- [14] J.C. Palencia Gutiérrez, J.J. Gutiérrez García, M. González Harbour. "On the schedulability analysis for distributed hard real-time systems". 9th Euromicro Workshop on Real-Time Systems. Toledo, 1997.



- [15] J.C. Palencia Gutiérrez and M. González Harbour, "Schedulability Analysis for Tasks with Static and Dynamic Offsets". Proceedings of the 18th. IEEE Real-Time Systems Symposium, Madrid, Spain, December 1998.
- [16] J.C. Palencia, and M. González Harbour, "Exploiting Precedence Relations in the Schedulability Analysis of Distributed Real-Time Systems". Proceedings of the 20th IEEE Real-Time Systems Symposium, 1999.
- [17] J.C. Palencia and M. González Harbour, "Offset-Based Response Time Analysis of Distributed Systems Scheduled under EDF". Euromicro conference on real-time systems, Porto, Portugal, June 2003.
- [18] R. Rajkumar, L. Sha, and J.P. Lehoczky. "Real-Time Synchronization Protocols for Multiprocessors". IEEE Real-Time Systems Symposium, December 1988.
- [19] R. Rajkumar. "Real-Time Synchronization Protocols for Shared Memory Multiprocessors". Proceedings of the 10th International Conference on Distributed Computing, 1990.
- [20] Redell, O. Response Time Analysis for Implementation of Distributed Control Systems, Doctoral Thesis, TRITA-MMK 2003:17, ISSN 1400-1179, ISRN KTH/MMK/R--03/17--SE, 2003
- [21] Juan M. Rivas, J. Javier Gutiérrez, J. Carlos Palencia, and Michael González Harbour. "Schedulability Analysis and Optimization of Heterogeneous EDF and FP Distributed Real-Time Systems". ECRTS-2011, Porto, June 2011.
- [22] P. Rodríguez Hernández and J.J. García Reinoso, "Nota sobre la Implementación del mecanismo de herencia", VI Jornadas de Tiempo Real, Gijón, February 2003 (In Spanish).
- [23] L. Sha, T. Ralya, and J.B. Goodenough, "An analytical approach to real-time software engineering", Software Engineering Institute Annual Technical Review, 1989.
- [24] L. Sha, R. Rajkumar, and J.P. Lehoczky. "Priority Inheritance Protocols: An approach to Real-Time Synchronization". IEEE Trans. on Computers, September 1990.
- [25] B. Sprunt, L. Sha, and J.P. Lehoczky. "Aperiodic Task Scheduling for Hard Real-Time Systems". The Journal of Real-Time Systems, Vol. 1, 1989, pp. 27-60.
- [26] M. Spuri. "Analysis of Deadline Scheduled Real-Time Systems". RR-2772, INRIA, France, 1996.
- [27] M. Spuri. "Holistic Analysis of Deadline Scheduled Real-Time Distributed Systems". RR-2873, INRIA, France, 1996.
- [28] K.W. Tindell, A. Burns, and A.J. Wellings. "Allocating Real-Time Tasks. An NP-Hard Problem Made Easy". Real-Time Systems Journal, Vol. 4, No. 2, May 1992. pp. 145- 166.
- [29] K. Tindell, "An Extendible Approach for Analysing Fixed Priority Hard Real-Time Tasks". Journal of Real-Time Systems, Vol. 6, No. 2, March 1994.
- [30] K. Tindell, and J. Clark, "Holistic Schedulability Analysis for Distributed Hard Real-Time Systems". Microprocessing & Microprogramming, Vol. 50, Nos.2-3, pp. 117-134, April 1994.
- [31] K. Tindell, "Adding Time-Offsets to Schedulability Analysis", Technical Report YCS 221, Dept. of Computer Science, University of York, England, January 1994.
- [32] J.M. Rivas, J.J. Gutiérrez, J.C. Palencia, and M. González Harbour, "Optimized Deadline Assignment and Schedulability Analysis for Distributed Real-Time Systems with Local EDF Scheduling". Proceedings of the 2010 International Conference on Embedded Systems and Applications, Las Vegas, Nevada, USA, ESA'2010. CSREA Press, July, 2010, ISBN:1-60132-141-4, pp. 150,156.
- [33] J. Sun, "Fixed Priority scheduling of end-to-end periodic tasks", Ph.D. thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 1997.
- [34] J. Liu, "Real-Time Systems," Prentice Hall, 2000.