



CISTER

Research Centre in
Real-Time & Embedded
Computing Systems

Conference Paper

Increasing Fixed-Priority Schedulability using Non-Periodic Load Shapers

Mitra Nasri

Geoffrey Nelissen*

*CISTER Research Centre

CISTER-TR-170506

2017/06/27

Increasing Fixed-Priority Schedulability using Non-Periodic Load Shapers

Mitra Nasri, Geoffrey Nelissen*

*CISTER Research Centre

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8321159

E-mail: mitra@mpi-sws.org, grrpn@isep.ipp.pt

<http://www.cister.isep.ipp.pt>

Abstract

Many real-time systems use fixed-priority scheduling (FPS) for three main reasons: (i) it is easy to understand, configure and analyze, (ii) it has low runtime overhead, and (iii) it is widely implemented in operating systems. Moreover, it is enforced by some standards such as Autosar for safety-critical real-time applications since it avoids unbounded delays in the case of overload in the system. However, FPS is not optimal since it may prioritize the execution of a high-priority task over a task with an urgent deadline.

In this work, we introduce load-shaping (LS); a technique that shapes the workload of tasks to limit their impact on the system's schedulability while improving their safety. LS is implemented using a reservation server for each task. Each server has a budget that replenishes with a given pattern, which might not necessarily be periodic. In this work, we discuss the open problems regarding parameter assignment for load-shaping techniques.

Increasing Fixed-Priority Schedulability using Non-Periodic Load Shapers

Mitra Nasri

Max Planck Institute for Software Systems (MPI-SWS),
mitra@mpi-sws.org

Geoffrey Nelissen

CISTER-INESC TEC, ISEP, Portugal.
grrpn@isep.ipp.pt

I. MOTIVATIONS

Many real-time systems use *fixed-priority scheduling* (FPS) for three main reasons: (i) it is easy to understand, configure and analyse, (ii) it has low runtime overhead, and (iii) it is widely implemented in commercial operating systems. Additionally, FPS is enforced by some industrial safety related standards such as Autosar. However, FPS is not optimal from a schedulability viewpoint as it may prioritise the execution of non-urgent high-priority tasks over tasks with urgent deadlines.

As an attempt to increase FPS schedulability, variations of the FPS scheduling scheme such as Preemption Threshold Scheduling (PTS) [1] and Dual-Priority (DP) scheduling [2] have been introduced. PTS raises the priority of a task (possibly several times) based on its remaining execution time, while DP promotes the task's priority after a given time following its release. It was shown that PTS is not optimal [3] and designing a strategy for correctly assigning priorities and promotion times such that all deadlines are met is still an open problem for DP [4].

In this work, we introduce *load-shaping* (LS); a technique that shapes the workload of tasks to limit their impact on the system's schedulability while improving the system's safety. LS is implemented using a reservation server for each task. Each server has a budget that replenishes itself following a given pattern.

Example. Fig. 1-(a) shows an example of a periodic task set which is not schedulable with FPS. However, the system may become schedulable if the workload released by τ_1 is properly shaped. An example of such shaping is shown in Fig. 1-(b), where τ_1 's execution is divided in two segments. τ_1 is allowed to execute 4 time units in the first segment and 8 time units in the second.

Note that load-shaping is more generic than a simple period transformation [5]. Period transformation divides the tasks' periods into smaller segments of equal lengths, while load-shaping may generate segments of different lengths executing unequal workloads (see example above). We claim that period transformation may result in having more segments than actually required for the system to become schedulable, e.g., for Fig. 1's example, period transformation would need to reduce τ_1 's period to 10 (i.e., the greatest-common-divisor of the tasks' periods), hence, dividing τ_1 in 3 segments instead of 2 as with load-shaping.

II. SYSTEM MODEL AND NOTATIONS

In this work, we consider a uniprocessor system executing a periodic task set τ . Each task τ_i in τ is characterised by a period T_i , a worst-case execution time (WCET) C_i , a relative deadline $D_i \leq T_i$, and a priority p_i (a lower value denotes a higher priority). Tasks are indexed such that $p_1 \leq p_2 \leq \dots \leq p_n$. We assume that tasks are independent, do not share resources and do not self-suspend. The task set utilisation is denoted by $U = \sum_{i=1}^n C_i/T_i$.

A *shaper task* (or simply shaper) τ_i^S for an original task $\tau_i \in \tau$ is defined as a sequence of segments $\langle S_{i,1}, S_{i,2}, \dots, S_{i,m_i} \rangle$, where each segment $S_{i,j} = (r_{i,j}, c_{i,j}, d_{i,j})$ is defined by a relative release time $r_{i,j}$ w.r.t. the release time of a job of τ_i , a relative deadline $d_{i,j}$ w.r.t. $r_{i,j}$, and a budget $c_{i,j}$ that limits the time during which τ_i can execute within segment $S_{i,j}$.

III. OPEN PROBLEM

As explained earlier, parameters of the shaper tasks are needed in order to use load-shaping. In addition, a schedulability test is needed to ensure that all timing constraints will be met at run-time. The following problem summarises our main goal.

Problem 1. Given a task set τ and a schedulability test S , find a set of parameters for shapers such that τ becomes schedulable with S .

Note that a solution to Problem 1 always exists. Indeed, we know that by using period transformation and breaking the periods of the original tasks into segments that are equal to the greatest-common-divisor (GCD) of all periods and deadlines, then the shaped task set becomes harmonic. Hence, by using Rate Monotonic (RM) it is possible to schedule the shaped task set with no deadline miss if the total utilisation U is smaller than 1 [5]. However, this solution significantly increases the number of preemptions. Therefore, the actual open problem is the following.

Problem 2. Given an unschedulable task set τ with $U \leq 1$, how can we assign parameters to shapers such that the system becomes schedulable and the number of preemptions is minimised.




Fig. 1. An example of load shaping. In a normal execution (a), task τ_2 is not schedulable. Using a load shaper for τ_1 (b) composed of two segments $S_{1,1} = (0, 4, 16)$ and $S_{1,2} = (16, 8, 14)$, task τ_2 becomes schedulable (assuming implicit deadlines).

IV. SKETCH OF A SOLUTION

Since each segment of a shaper task has a fixed release offset w.r.t. to the release time of its associated jobs, one can model segments as independent tasks that share the same priority and period than their task but have different release offsets and deadlines. Therefore, one can use existing schedulability tests for periodic tasks with release offsets (e.g., based on a feasibility-interval) as the underlying response-time analysis.

An efficient solution to assign parameters to shapers and check the system’s schedulability is to do both at the same time, and hence, find a set of parameters that guarantees schedulability by design. A constructive technique can be sketched as follows: **(i)** Start by assuming that each task τ_i has a default shaper $\tau_i^S = \langle S_{i,1} \rangle$, where $c_{i,1} = C_i$, $d_{i,1} = D_i$, and $r_{i,1} = 0$, **(ii)** Check the system’s schedulability by analysing the response time of each job in a feasibility-interval. Stop at the first unschedulable job J_i . **(iii)** Break the shapers of tasks with higher priorities than J_i into more segments such that the supply-bound-function (SBF) becomes larger than the request-bound-function (RBF) of level i in at least one time instant between the release and deadline of J_i . Go back to step (ii).

The challenge of step (iii) is to choose a time instant t (or maybe a set of such instants) at which the RBF is larger than the supply, and try to reduce $RBF(t)$ by $\delta(t) = \max\{0, RBF(t) - SBF(t)\}$. For example, a potential candidate for t is a time instant at which $\delta(t)$ is minimized, since in that case, a smaller amount of higher-priority workload must be postponed. However, in the general case, one may attempt to postpone higher-priority workload for any t at which $\delta(t)$ changes its value. For each time instant t , there is a set of high priority tasks that are not finished yet. Among these tasks, those that have their deadline later than t are the candidates whose workload can be postponed.

V. DISCUSSION

Beside safety-critical systems for which load-shaping allows higher predictability and dependability, applications of load-shaping can be found in multi-constrained systems, where the timing behaviour of a task might be constrained by other non-functional properties. Tasks might use different shapers depending on runtime requirements or system needs. For example, the shaper of a task τ_i can be different according to the following constraints: reduced response-time or reduced processor temperature. Under the first constraint, the shaper should assign more budget at the release of each job while in the latter case the shaper should force the task to be scheduled with a (potentially lower) rate leading to a better temperature profile for the system. The problem therefore becomes: how can one find shapers’ parameters under additional constraints. Further, if a system has several modes of execution for the same set of tasks (e.g., power saving and high performance mode), one should understand how to switch between modes of one or more shapers without causing any deadline miss in the system.

ACKNOWLEDGMENT

This work is supported by a fellowship from Alexander von Humboldt Foundation. It was also partially supported by National Funds through FCT/MEC (Portuguese Foundation for Science and Technology) and co-financed by ERDF (European Regional Development Fund) under the Portugal2020 Program, within the CISTER Research Unit (CEC/04234).

REFERENCES

- [1] Y. Wang and M. Saksena, “Scheduling fixed-priority tasks with preemption threshold,” in *International Conference on Real-Time Computing Systems and Applications (RTCSA)*, 1999, pp. 328–335.
- [2] A. Burns and A. J. Wellings, “Dual Priority Assignment: A Practical Method for Increasing Processor Utilisation,” in *Euromicro Workshop on Real-Time Systems (EWRTS)*, 1993, pp. 48–55.
- [3] R. J. Bril, M. M. van den Heuvel, and J. J. Lukkien, “Improved feasibility of fixed-priority scheduling with deferred preemption using preemption thresholds for preemption points,” in *International conference on Real-Time Networks and Systems (RTNS)*. ACM, 2013, pp. 255–264.
- [4] A. Burns, “Dual Priority Scheduling: Is the Processor Utilisation bound 100%,” in *International Real-Time Scheduling Open Problems Seminar (RTSOPS)*, 2010, pp. 3–5.
- [5] L. Sha, J. P. Lehoczky, and R. Rajkumar, “Solutions for some practical problems in prioritized preemptive scheduling,” in *IEEE Real-Time Systems Symposium (RTSS)*, 1986, pp. 181–191.