Efficient Schedulability Test for Dynamic-Priority Scheduling of Mixed-Criticality Real-Time Systems

Gu, Xiaozhe; Easwaran, Arvind

2017

Gu, X., & Easwaran, A. (2017). Efficient Schedulability Test for Dynamic-Priority Scheduling of Mixed-Criticality Real-Time Systems. *ACM Transactions on Embedded Computing Systems*, 17(1), 24-.

<https://hdl.handle.net/10356/85848>

<https://doi.org/10.1145/3105922>

© 2017 Association for Computing Machinery (ACM). This is the author created version of a work that has been peer reviewed and accepted for publication by *ACM Transactions on Embedded Computing Systems*, Association for Computing Machinery (ACM). It incorporates referee's comments but changes resulting from the publishing process, such as copyediting, structural formatting, may not be reflected in this document. The published version is available at: [<https://doi.org/10.1145/3105922>].

Downloaded on 26 Apr 2023 20:57:41 SGT

Efficient Schedulability Test for Dynamic-Priority Scheduling of Mixed-Criticality Real-Time Systems

Xiaozhe Gu and Arvind Easwaran, Nanyang Technological University

Systems in many safety-critical application domains are subject to certification requirements. In such a system, there are typically different applications providing functionalities that have varying degrees of criticality. Consequently, the certification requirements for functionalities at these different criticality levels are also varying, with very high levels of assurance required for a highly critical functionality, whereas relatively low levels of assurance required for a less critical functionality. Considering the timing assurance given to various applications in the form of guaranteed budgets within deadlines, a theory of real-time scheduling for such multi-criticality systems has been under development in the recent past. In particular, an algorithm called Earliest Deadline First with Virtual Deadlines (EDF-VD) has shown a lot of promise for systems with two criticality levels, especially in terms of practical performance demonstrated through experiment results. In this paper we design a new schedulability test for EDF-VD that extend these performance benefits to multi-criticality systems. We propose a new test based on demand bound functions and also present a novel virtual deadline assignment strategy. Through extensive experiments we show that the proposed technique significantly outperforms existing strategies for a variety of generic real-time systems.

CCS Concepts: • **Computer systems organization** → **Real-time operating systems**;

General Terms: Schedulability Analysis, Design of Real-Time Scheduler

Additional Key Words and Phrases: Demand Bound Function, Mixed-Criticality System

ACM Reference Format:

Xiaozhe Gu and Arvind Easwaran. 2017. Efficient Schedulability Test for Dynamic-Priority Scheduling of Mixed-Criticality Real-Time Systems. *ACM Trans. Embedd. Comput. Syst.* x, x, Article x (July 2017), 25 pages.

DOI: 0000001.0000001

1. INTRODUCTION

Real-time systems are defined as those systems in which the correctness of the system depends not only on the logical result of computation, but also on the time at which the results are produced [Stankovic et al. 1998]. For example a pacemaker is inserted in a person's chest to provide electrical impulses at regular intervals to help the heart beat. Here the pacemaker must provide service with certain timing constraints, and applications with these kinds of timing constraints are considered real time.

Timing constraints in real-time systems are often modeled as deadlines. If a schedulable activity (e.g., a piece of job) executes and completes before its assigned deadline, the deadline is met (and otherwise it is missed). This means, in order to meet the deadline, the scheduler must have apriori knowledge on the amount of execution that the job would request. Further, to achieve high timing predictability in the presence of various sources of variability, these systems must be built under pessimistic assumptions to cope with worst case scenarios. That is, the scheduler typically assumes that the job would execute for a certain worst-case amount of time (denoted as WCET for

This work was funded in part by the Ministry of Education, Singapore, Tier-1 grant #RG21/13 and Tier-2 grant #ARC9/14, and by the Start-Up-Grant from SCSE, NTU, Singapore. Authors addresses: Xiaozhe Gu and Arvind Easwaran, School of Computer Science and Engineering, Nanyang Technological University, Nanyang Avenue, Singapore 639798; emails: guxi0002@e.ntu.edu.sg, arvinde@ntu.edu.sg

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2017 Copyright held by the owner/author(s). 1539-9087/2017/07-ARTx \$15.00

DOI: 0000001.0000001

Worst-Case Execution Time), which encompasses all the possible variations in execution time. However determining an exact WCET for a job is very difficult, and usually a conservative overestimation of the true WCET is used to analyze and schedule a real-time system.

An increasing trend in embedded systems that multiple functionalities with different levels of “criticality” (or importance) are developed independently and integrated together on a common computing platform. This trend is evident in industry-driven initiatives such as ARINC653 [ARI 2008], Integrated Modular Avionics (IMA) in avionics and AUTOSAR in automotive. An important notion behind this trend is the safe isolation of separate functionalities of different criticality, primarily to achieve fault containment. For example in a modern car, devices and software are integrated into the entertainment system and often run on the same platform as the more critical instrument panel information display subsystem. The challenge in such a system is managing the dramatically-different nature of resource requirements between the less critical infotainment functions characterized by best-effort or soft real-time needs and the more critical display functions that require strong reliability.

Here we define criticality as the level of assurance against severe failure needed for a system component. Typically, failure of high critical functionality would cause a more severe consequence to the whole system than the failure of a low critical functionality. Thus high critical functions require higher assurance that their estimated WCETs will not be exceeded. As a result, their WCET estimates tend to involve very conservative assumptions about the system (cache flushing on preemption, over-provisioning for potentially missed execution paths, etc.) that are very unlikely to occur in practice. Consequently, the system resources are in fact severely under-utilized in practice because high critical functions would rarely execute as much as their WCET estimates.

This raises the challenge of how to balance the conflicting requirements of isolation for safety assurance and efficient resource sharing for economical benefits. The concept of *mixed-criticality* (MC) appears to be important in meeting those two goals. In order to close such a gap in resource utilization, Vestal [Vestal 2007] proposed the MC task model that comprises of different WCET estimates. These estimates are determined at different levels of confidence (“criticality”) based on the following principle. A reasonable low-confidence WCET estimate, even if it is based on measurements, may be sufficient for almost all possible execution scenarios in practice. As long as this estimate is not violated, both low critical tasks and high critical tasks are required to meet their timing constraints. In the highly unlikely event that this estimate is violated, as long as the scheduling mechanism can ensure deadline satisfaction for high critical applications, the resulting system design may still be considered as safe. Considering such MC real-time task systems with two criticality levels, several studies have proposed scheduling algorithms and corresponding schedulability tests in the past [Ekberg and Yi 2012; Baruah et al. 2012a; Baruah et al. 2011b]. There have also been some recent studies that extend some of these results to more than two criticality levels [Ekberg and Yi 2014; Fleming 2013; Baruah et al. 2015].

In this paper we focus on the problem of EDF (earliest deadline first) scheduling of mixed-criticality systems on uniprocessors. In particular, we address the problem of scheduling multi-criticality real-time task systems (systems with more than two criticality levels). Baruah and Vestal [Baruah and Vestal 2008] first considered MC scheduling with EDF. Later Park and Kim [Park and Kim 2011] proposed Criticality Based EDF that applies a combination of off- and on-line analyses to run high critical jobs as late as possible, and low critical jobs in the generated slack. Baruah et al. proposed an algorithm called EDF-VD (EDF - with virtual deadlines) [Baruah et al. 2012a] for a dual-criticality system. High critical tasks have their deadlines reduced by the same factor (if necessary) during low-criticality mode execution. They

demonstrate both theoretically and via evaluations that this is an effective scheme. EDF-VD [Baruah et al. 2012a] is constrained to dual-criticality implicit deadline systems, and was extended to multi-criticality systems in [Baruah et al. 2015]. Although EDF-VD analyses the system across multiple criticality levels together, it is still very pessimistic even for two criticality levels because of the following factors: 1) The virtual deadlines for all the high critical tasks are uniformly assigned based on a single common factor, and 2) Using demand density to characterize the demand of a constrained-deadline task (task with a virtual deadline smaller than period) will always be pessimistic.

A more general demand bound function based analysis for EDF mixed-criticality scheduling was proposed by Ekberg and Yi [Ekberg and Yi 2012]. They also introduced a heuristic virtual deadline tuning algorithm called GreedyTuning where deadlines can be reduced by different factors. GreedyTuning can increase the chances that a task system is schedulable by EDF. From extensive experiments, they show that GreedyTuning outperforms all existing works on MC scheduling for a variety of generic real-time systems. GreedyTuning was extended to multi-criticality systems in [Ekberg and Yi 2014]. However it suffers from a drawback that its schedulability performance drops significantly as the number of criticality levels increases (see Figure 17 in [Ekberg and Yi 2014]). In Figure 9 of this paper, we also show that it is not good at scheduling task systems with a larger percentage of high critical tasks. The primary reason for this drop in performance is that they analyze the system in each criticality mode from the time instant when the mode-switch happens, conservatively assuming maximum carry-over interference when the system behavior switches from a lower criticality level to the one being analyzed.

The test proposed in [Easwaran 2013] addresses this problem by considering the system behavior from the start of a busy interval in a dbf-based analysis, but it is restricted to dual-criticality systems. The first challenge in extending this initial result to more than two criticality levels is that the task execution pattern that can result in worst-case demand and hence the dbf for dual-criticality system is no longer valid in multi-criticality systems. Besides, if we consider the demand from the start of a busy interval as [Easwaran 2013] does, there would be multiple mode-switches happening at $S_1, S_2, \dots, S_m$ during the time interval. Then, we have to consider all possible combinations of these mode-switch instants, and as a result, the complexity of the test is exponential in the number of criticality levels. Also, given a set of these mode-switch instants, the task execution pattern that will result in the worst-case demand depends on all the tasks in the system and their remaining execution time at each of those mode-switches. It is therefore non-trivial to determine a worst-case pattern with low pessimism, and there is no known technique for the same.

Contribution: In sum the contributions of this work can be summarized as follows.

- In this paper, we overcome the challenges and extend the dual-criticality dbf-based test [Easwaran 2013] for multi-criticality systems that have the same time complexity as the dual-criticality ones.
- To further improve the performance of the proposed design, we also develop a new virtual deadline assignment strategy, extending the strategy proposed by [Ekberg and Yi 2014]. Finally, through experimental evaluation, we demonstrate that the proposed technique significantly outperforms the existing ones for generic real-time MC task systems.

Other Related Work: Since Vestal [Vestal 2007] first proposed the MC task model and an algorithm based on Audsley's priority assignment strategy [Audsley 1991], there have been a series of publications on the scheduling of MC systems on uniprocessors. A number of proposed studies are restricted to address the problem of scheduling

a finite set of mixed-criticality jobs with criticality dependent execution times [Li and Baruah 2010; Baruah et al. 2010; Baruah et al. 2012b]. The model these studies use is a constrained one, because in many real-time systems each task is able to generate an infinite number of jobs. For example, the engine control unit in a car periodically senses and processes information to efficiently control the fuel injection and emissions. Hence, these studies are superseded by studies that are applicable to the more general sporadic MC task model, which is also the focus of this paper. Vestal's approach [Vestal 2007] is the first work that uses Response-Time Analysis (RTA) to analyze the schedulability of generic MC task systems. This work was later improved by the Static Mixed Criticality Scheme (SMC) [Baruah and Burns 2011]. Adaptive scheme (AMC) [Baruah et al. 2011b] goes further and it outperforms all the previous works on fixed priority MC scheduling in terms of schedulability. Fleming and Burns [Fleming 2013] extended AMC for task systems with an arbitrary number of criticality levels, focusing particularly on five levels as this is the maximum found in automotive and avionics standards. There are various works (e.g. [Su and Zhu 2013; Burns et al. 2015; Bate et al. 2015a]) concerned about addressing other problems about MC scheduling, e.g., how to switch back to low critical mode or support low critical execution, but is not the focus of this work.

The rest of the paper is organized as follows. In Section 2 we first introduce the mixed-criticality task and scheduling model, and in Section 3 we give a brief introduction about GreedyTuning [Ekberg and Yi 2014], which is also the work we aim to improve. We derive our new multi-mode demand bound function (dbf) in Section 4. In Section 5 we present a novel deadline tuning algorithm to improve the performance of the proposed test. Finally in Section 6 we show that our proposed test dominates GreedyTuning [Ekberg and Yi 2014] from experimentation.

2. TASK AND SCHEDULING MODEL

2.1. MC Task Model

The sporadic task model [Mok 1983] is a generic model for capturing the real-time requirements of many event-driven systems including those with MC such as avionics and automotive. A sporadic task [Mok 1983] can be specified as $\tau_i = (T_i, C_i, D_i)$, where T_i denotes minimum separation between successive job releases, D_i denotes its relative deadline, and C_i denotes its worst-case execution time (WCET). No job of τ_i is expected to execute for more time than its WCET, and otherwise the system is regarded as exhibiting erroneous behavior. Any job released by τ_i is required to complete by its deadline, and deadline miss is regarded as system failure.

The task model widely used in most previous studies on scheduling of MC systems (e.g., [Ekberg and Yi 2012; Baruah et al. 2011b; Vestal 2007; Easwaran 2013]) is a straightforward extension of the classic sporadic task model [Mok 1983] to a MC setting; the worst-case execution times of a single task can vary between criticality levels. However in most of these works, the task model is constrained to two levels, i.e., a task can either be a low critical task or a high critical task. Instead, we use a more general multi-criticality model in this paper. Formally, a task τ_i in a MC sporadic task set $\tau = \{\tau_1, \tau_2, \dots, \tau_k\}$ can be represented as a tuple (T_i, C_i, D_i, L_i) , where:

- $T_i \in \mathbb{Z}^+$ is the minimal time separation between the release of two successive jobs¹,
- $L_i \in \mathbb{Z}^+$ is the criticality level of τ_i , and $L_i = 1$ denotes the lowest criticality level,
- $D_i \in \mathbb{Z}^+$ is the relative deadline,
- $C_i = (C_i^1, C_i^2, \dots, C_i^{L_i})$ is a L_i -tuple of estimated execution time budgets respectively for each criticality level.

¹We focus on integer release time model and hence $T_i \in \mathbb{Z}^+$.

Since the worst-case execution times for higher criticality levels are estimated more conservatively, we make the standard assumption that

$$\forall \tau_i \in \tau : C_i^1 \leq C_i^2 \leq \dots \leq C_i^{L_i} \leq D_i \leq T_i$$

2.2. Mode Switch

The system initially starts in L_1 (short for level one) criticality mode (i.e., the lowest criticality mode), each task $\tau_i \in \tau$ releases a potentially infinite sequence of jobs $\langle J_i^1, J_i^2, \dots \rangle$ in the standard manner: if $r(J)$, $d(J) \in \mathbb{R}$ denote the release time and absolute deadline of job J , then

$$\begin{aligned} - r(J_i^{k+1}) &\geq r(J_i^k) + T_i \\ - d(J_i^k) &= r(J_i^k) + D_i \end{aligned}$$

The system stays in L_1 criticality mode as long as all the jobs of every task τ_i with $L_i \geq 2$ do not execute beyond their L_1 execution time estimate C_i^1 . Once a job executes for its entire L_1 execution time estimate C_i^1 without signaling that it has finished, the system immediately switches to L_2 criticality mode. That is, in general as shown in Figure 1, when the first job J_i such that $L_i \geq m$ executes for more than its L_{m-1} execution time estimate C_i^{m-1} but does not signal that it has finished, the system switches to L_m criticality mode. The mode switch time instant when the system switches from L_{m-1} to L_m criticality mode is called *criticality mode switch instant* and is denoted as S_m .

Fig. 1: Criticality Mode Switch Instant S_m

After this mode switch at S_m , jobs with criticality level $L_i < m$, including those that were released before S_m but did not complete until S_m , are no longer required to meet deadlines. For simplicity of analysis, these jobs are assumed to be dropped thereafter. However we must still meet all the deadlines for jobs with $L_i \geq m$, even if they require up to C_i^m budgets. If the system is now in L_m criticality mode, we assume it will switch back to L_1 mode whenever the processor is idle.

Note that the assumptions in the preceding paragraph on mode switch are consistent with the standard literature on MC scheduling (e.g., [Li and Baruah 2010; Baruah et al. 2011b; Vestal 2007; Easwaran 2013]). There are some studies that focus on dealing with those assumptions, such as reducing the penalty on low-critical tasks (e.g., [Huang et al. 2014; Huang et al. 2013; Gu et al. 2015]), and switching back to a lower criticality mode earlier than the processor idle time (e.g., [Bate et al. 2015b]). These studies are orthogonal to the focus of this paper, which is to derive efficient schedulability tests for EDF-scheduled MC task systems.

Definition 2.1 (MC-Schedulable). We define a MC task system to be MC-schedulable if $\forall m : 1 \leq m \leq M$ where $M = \max_{\tau_i \in \tau} \{L_i\}$, all jobs with criticality level $L_i \geq m$ can receive a budget up to C_i^m and signal completion between their release time and deadline while the system stays in L_m criticality mode.

2.3. Demand Bound Function Based Schedulability Analysis

Demand bound function was first proposed to analyze the schedulability of non-MC real-time workloads [Baruah et al. 1990]. The demand bound function captures the maximum execution demand a task can generate for a given time interval length.

Definition 2.2 (Demand bound function). A demand bound function $dbf(\tau_i, e)$ gives an upper bound on the maximum possible execution demand of task τ_i in any time interval of length e , where demand is calculated as the total amount of required execution time of jobs with their whole scheduling windows within the time interval.

For example a task $\tau_i = (T_i = 5, C_i = 2, D_i = 3, L_i = 1)$ can generate as much as $2 \times C_i = 2 \times 2 = 4$ time units execution demand for a time interval length equal to 8.

For non-MC constrained deadline task model, the demand bound function for a given time interval length e can be computed in constant time [Baruah et al. 1990] using the following equation.

$$dbf(\tau_i, e) = \left(\left\lfloor \frac{e - D_i}{T_i} \right\rfloor + 1 \right) \times C_i \quad (1)$$

As long as we can guarantee that the total execution demand of a task set τ is always smaller than or equal to the time interval length e for all values of e , we can claim τ is schedulable by the EDF algorithm on a uniprocessor platform.

THEOREM 2.3 ([BARUAH ET AL. 1990]). *A non-mixed-criticality sporadic task set τ is successfully scheduled by the earliest deadline first (EDF) algorithm on a dedicated unit speed uniprocessor platform if*

$$\forall e \in \{1, 2, \dots, e^{max}\} : \sum_{\tau_i \in \tau} dbf(\tau_i, e) \leq e$$

where e^{max} is a pseudo-polynomial in the size of the input [Baruah et al. 1990] as long as the system utilization is bounded by some constant smaller than 1.

3. BACKGROUND: EXISTING DBF BASED TEST FOR MC TASK SYSTEMS

In this section, we first extend the idea of demand bound function to the mixed-criticality setting. Then we introduce the Single-Mode (SM) demand bound function of mixed-criticality workloads derived in an existing work [Ekberg and Yi 2012]. Test based on this SM dbf has been shown to dominate previous studies (e.g., EDF-VD [Baruah et al. 2012a] and AMC-max [Baruah et al. 2012b]) in terms of the ability to schedule MC sporadic task systems.

Let $M = \max_{\tau_i \in \tau} \{L_i\}$ and $dbf_{SM}(\tau_i, e, m)$ where $m \in \{1, 2, \dots, M\}$ denote the SM demand bound function of τ_i for the time interval $[S_m, S_m + e)$, when the system is currently in L_m criticality mode and was in L_{m-1} criticality mode before that. As we can observe, SM dbf test separately analyzes the system in each criticality mode because it only considers the demand during $[S_m, S_m + e)$ but the system behavior before S_m is totally ignored. Therefore, Theorem 2.3 can be extended in a straightforward way as follows.

THEOREM 3.1 (PROPOSITION 3 [EKBERG AND YI 2014]). *A MC task set τ is schedulable by EDF on a dedicated unit speed uniprocessor platform for all the criticality modes if the following conditions hold:*

$$\forall m \in \{1, 2, \dots, M\} : \forall e \in \{1, 2, \dots, e^{max}\} : dbf_{SM}(\tau, e, m) \leq e$$

where $dbf_{SM}(\tau, e, m) = \sum_{L_i \geq m} dbf_{SM}(\tau_i, e, m)$, and e^{max} is a pseudo-polynomial in the size of the input if the utilization of each criticality mode is bounded by some constant smaller than 1.

We define condition CN_m^S as follows

$$\forall e \in \{1, 2, \dots, e^{max}\} : \sum_{L_i \geq m} dbf_{SM}(\tau_i, e, m) \leq e$$

Condition CN_m^S captures the schedulability of task set τ for L_m criticality mode on the assumption that τ is schedulable in $L_{m'}$ ($m' < m$) criticality modes. To compute $dbf_{SM}(\tau_i, e, m)$ we need to determine the maximum demand that τ_i can generate in the interval $[S_m, S_m + e]$. However, the demand of any task τ_i in L_m criticality mode depends on the release pattern in all the previous criticality modes.

In L_1 criticality mode, each task τ_i will behave like a normal non-mixed-criticality task, and all jobs are guaranteed to execute for at most C_i^1 time units. Therefore the dbf for non-mixed-criticality tasks can be directly applied to capture the demand of τ_i in L_1 criticality mode, i.e.,

$$dbf_{SM}(\tau_i, e, 1) = \left(\left\lfloor \frac{e - D_i}{T_i} \right\rfloor + 1 \right) \times C_i^1 \quad (2)$$

If τ_i has $L_i < m$, then $dbf_{SM}(\tau_i, e, m) = 0$ because tasks with $L_i < m$ will be discarded after S_m . On the other hand, if $L_i \geq m$, then we need to consider the job that is released before the mode switch instant S_m but has its deadline after S_m , because this job can affect τ_i 's execution demand after S_m . We call such jobs as carry-over jobs.

Definition 3.2 (Carry-over job). A job J_i^m from a $L_i \geq m$ criticality task that is active (released before S_m and has deadline after S_m) at the time of the switch to L_m criticality mode is called a carry-over job for L_m criticality mode.

3.1. Characterizing the Demand of Carry-Over Jobs

While we can discard all the active jobs with $L_i < m$, the remaining execution demand of the carry-over jobs must be completed by their respective deadlines, and hence the demand of carry-over jobs must be accounted for in $dbf_{SM}(\tau_i, e, m)$.

Fig. 2: After a switch to higher criticality mode, the remaining execution demand of a carry-over job must be finished in its remaining scheduling window.

At the time of the switch to L_m criticality mode, a carry-over job J_i^m from task τ_i has x ($x \geq 0$) time units left until its deadline as shown in Figure 2. Since this job would have met its deadline in $L_{m'}$ where ($m' < m$) criticality mode if the mode-switch at S_m had not happened, there can be at most x time units left to finish its maximum possible remaining execution demand C_i^{m-1} . That is, the job must have already executed at least $\max(0, C_i^{m-1} - x)$ time units before the mode-switch at S_m .

(otherwise deadline miss could happen in L_{m-1} criticality mode). After mode-switch S_m , the carry-over job may now execute for up to C_i^m time units in total, and therefore the total execution demand remaining for the carry-over job after the switch is at most $C_i^m - \max(0, C_i^{m-1} - x)$. Unfortunately if $x \rightarrow 0$, condition CN_m^S with $(m > 1)$ cannot be satisfied because as long as $C_i^m - C_i^{m-1} > 0$, we can always find a small x so that $dbf_{SM}(\tau_i, e \rightarrow 0, m) = C_i^m - C_i^{m-1} > 0$.

The problem described above stems from the fact that EDF may execute a carry-over job quite late, and hence it can not finish its remaining execution demand after the mode-switch at S_m . To solve this problem, virtual deadlines in different criticality modes have been introduced [Baruah et al. 2012a; Ekberg and Yi 2012; Easwaran 2013]. When the system is in L_m criticality mode, tasks τ_i with $L_i \geq m$ are scheduled by EDF scheduler according to virtual deadline D_i^m ($D_i^m \leq D_i$). In L_m criticality mode, any task τ_i with $L_i \geq m$ must finish its execution demand C_i^m by its virtual deadline D_i^m . Since a job J_i can now have multiple deadlines, we use $d(J_i, m) = r(J_i) + D_i^m$ to denote its absolute virtual deadline for L_m criticality mode.

Virtual deadline enables the carry-over job to have extra slack time, $D_i^m - D_i^{m-1}$, to finish its additional demand $C_i^m - C_i^{m-1}$ at the cost of a higher load of execution demand in lower criticality mode. However we should note that virtual deadlines are not actual deadlines, and can be determined by deadline tuning algorithms [Ekberg and Yi 2012; Easwaran 2013] to improve EDF schedulability. The remaining execution demand for a carry-over job can then be bounded with the following lemma.

LEMMA 3.3 (DEMAND OF CARRY-OVER JOBS [EKBERG AND YI 2012]). *Assume that EDF uses virtual relative deadline D_i^m in L_m criticality mode for tasks with $L_i \geq m$, and that we can guarantee that the demand is met in all lower criticality modes (i.e., $L_{m'}$ with $(m' < m)$) with respective virtual deadlines. If the carry-over job of τ_i has a remaining scheduling window of x time units until its deadline D_i^m , as illustrated in Figure 3, then the following hold:*

- (1) *If $x < D_i^m - D_i^{m-1}$, then the job must have already finished before S_m .*
- (2) *If $x \geq D_i^m - D_i^{m-1}$, then the job's remaining execution demand after S_m is bounded by $C_i^m - \max(0, C_i^{m-1} - x + D_i^m - D_i^{m-1})$.*

As we can observe, to maximize the total demand in L_m criticality mode, SM dbf conservatively assumes maximum possible carry-over demand $C_i^m - \max(0, C_i^{m-1} - x + D_i^m - D_i^{m-1})$ from L_{m-1} criticality mode.

Fig. 3: A carry-over job of τ_i has a remaining scheduling window of length x after the switch to L_m mode. Here the switch happens before $r(J_i^m) + D_i^{m-1}$.

3.2. Formulating the SM DBF

The execution demand of τ_i for time interval $[S_m, S_m + e)$ is equal to sum of the unfinished execution demand of the carry-over job, and the demand of jobs released after the carry-over job in this interval.

LEMMA 3.4 (MAXIMUM DEMAND PATTERN [EKBERG AND YI 2014]). *Task τ_i with $L_i \geq m$ can generate maximum execution demand in L_m ($m > 1$) criticality mode for a time interval length e when the corresponding virtual deadline of some job $r(J_i) + D_i^m$ is at the end of this time interval e and all preceding jobs are released as late as possible, as shown in Figure 4.*

Therefore the demand bound function of τ_i in L_m ($m > 1$) criticality mode can be summarized as follows:

$$dbf_{SM}(\tau_i, e, m) = \max(0, (1 + \lfloor \frac{e - (D_i^m - D_i^{m-1})}{T_i} \rfloor) \times C_i^m) - done_m(\tau_i, e) \quad (3)$$

Here $done(\tau_i, e)$ captures the execution demand of the carry-over job that must finish before S_m and is equal to:

$$done_m(\tau_i, e) = \begin{cases} \max(0, C_i^{m-1} - \text{mod}(e, T_i) + D_i^m - D_i^{m-1}), & \text{if } D_i^m - D_i^{m-1} \leq \text{mod}(e, T_i) \\ & \text{and } \text{mod}(e, T_i) < D_i^m \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $\text{mod}(e, T_i) = e \bmod T_i$ is equal to the length of the remaining scheduling window of the carry-over job after S_m .

Fig. 4: Demand of τ_i in L_m ($m > 1$) Criticality Mode

4. MULTI-MODE DBF BASED TEST FOR MC TASK SYSTEMS

The SM demand bound function considers the behavior of each criticality mode separately. That is, it does not use the execution demand in the previous criticality mode (e.g., L_{m-1}) to determine the remaining execution for carry-over jobs when mode switch happens at S_m . If the execution load of task set τ in L_{m-1} criticality mode is low, then carry-over jobs of many tasks would finish well before their deadlines, and would not generate any carry-over demand. This property leads to some interesting results like the fact that SM dbf based test cannot even schedule some task sets that are schedulable by reservation based approaches (i.e., all tasks are allocated $C_i^{L_i}$ execution budgets).

Example 4.1. Suppose task set $\tau = \{\tau_1, \tau_2\}$ has two tasks, where τ_1 and τ_2 are given in the following table. Obviously τ is schedulable by reservation based approaches because the utilization $\frac{C_1^2}{T_1} + \frac{C_2^2}{T_2} < 1$. However according to GreedyTuning, τ is not schedulable for any possible combination of virtual deadlines.

Task	T_i	C_i	D_i	L_i
τ_1	15	$\{C_1^1 = 3, C_1^2 = 7\}$	15	2
τ_2	2	$\{C_2^1 = 1, C_2^2 = 1\}$	2	2

To address this drawback, in this section we propose a Multi-Mode (MM) demand bound function that collectively bounds the demand of τ in L_{m-1} and L_m criticality modes. Suppose the time interval of interest is $[S_{m-1}, S_{m-1} + e)$ where $S_{m-1} \leq S_m \leq S_{m-1} + e$. For simplicity we assume $S_{m-1} = 0$ because what determines the total demand is the time interval length $S_m - S_{m-1}$ and e , i.e., the dbf is independent of the exact value of S_{m-1} . As a result, the interval becomes $[0, e)$ where $0 \leq S_m \leq e$. Let $dbf_{MM}(\tau, S_m, e, m)$ denote the total execution demand of τ for the time interval $[0, e)$.

The test proposed in [Easwaran 2013] considers the system behavior from the start of a busy interval in a dbf-based analysis, but is limited to a dual-criticality task system. In a dual-criticality task system, there is at most one mode-switch. However, in a task system with more than two criticality levels, we have to consider the demand from jobs released before S_{m-1} but have deadline after S_{m-1} . As a result, the dbf analysis in [Easwaran 2013] is no longer valid in multi-criticality systems. In this section we extend the dbf analysis in [Easwaran 2013] to multi-criticality systems and present a MM dbf based test for multi-criticality systems.

The following theorem is a straightforward extension of Theorem 3.1, and can be used to determine whether a task set τ is schedulable by EDF on a dedicated unit speed uniprocessor platform.

THEOREM 4.2. *A MC task set τ is schedulable by EDF on a dedicated unit speed uniprocessor platform for all the criticality modes if the following conditions hold:*

$$\forall m \in \{1, 2, \dots, M\} : \forall e \in \{1, 2, \dots, e^{max}\} : \forall S_m \in \{1, 2, \dots, e\} : dbf_{MM}(\tau, S_m, e, m) \leq e$$

where e^{max} is pseudo-polynomial in the size of the input [Gu et al. 2015] if the system utilization of each criticality mode is bounded by some constant smaller than 1, and is also derived in Appendix A.3.

In the time interval of interest $[S_{m-1} = 0, e)$, a job can experience at most two mode-switches at $S_{m-1} = 0$ and S_m , respectively. We categorize such carry-over jobs into four types: J_i^X where $X \in \{A, B, C, D\}$, and use J_i^{X+1} to denote the next job released after J_i^X . The patterns of these jobs are shown in Figure 5.

- J_i^A . $L_i = m - 1$, $r(J_i^A) < S_{m-1}$ and $r(J_i^{A+1}) \geq S_{m-1}$.
- J_i^B . $L_i = m - 1$, $r(J_i^B) \geq S_{m-1}$, $r(J_i^B) \leq S_m$ and $r(J_i^{B+1}) > S_m$.
- J_i^C . $L_i \geq m$, $r(J_i^C) < S_{m-1}$ and $r(J_i^{C+1}) \geq S_{m-1}$.
- J_i^D . $L_i \geq m$, $r(J_i^D) \geq S_{m-1}$, $r(J_i^D) \leq S_m$ and $r(J_i^{D+1}) > S_m$.

Fig. 5: Four types of jobs that experience mode switch.

Let $dbf_{MM}(J_i^X, S_m, e, m)$, where $X \in \{A, B, C, D\}$, to denote the demand of job J_i^X during $[0, e)$. Given the value of $r(J_i^X)$, we can simply calculate its demand. Thus the detailed equations of $dbf_{MM}(J_i^X, S_m, e, m)$ will be presented in Appendix A.1 because it is very intuitive.

4.1. MM dbf for a single task ($dbf_{MM}(\tau_i, S_m, e, m)$)

In the time interval $[S_{m-1} = 0, e)$, tasks with $L_i = m - 1$ could execute during $[0, S_m)$, but would be dropped after S_m . Tasks with $L_i \geq m$ could execute during the whole time interval $[0, e)$. For L_1 criticality mode (the initial mode), $dbf_{MM}(\tau_i, S_0, e, 1)$ can be obtained from demand bound function for non-mixed sporadic tasks (see Equation 2).

4.1.1. Case when $L_i = m - 1$. The maximum demand generated by task τ_i with $L_i = m - 1$ during the interval $[0, S_m)$ is equal to the sum of demand of all jobs released during $[0, S_m)$ and the execution demand of J_i^A during $[0, S_m)$.

Let $DEM_i(r_A)$ denote the demand that τ_i generates during $[0, S_m)$ when $r(J_i^A) = r_A$ and all jobs are released as soon as possible with period T_i . Given $r(J_i^A)$ there will be at most $n_{m-1} = \lfloor (S_m - (r(J_i^A) + T_i)) / T_i \rfloor$ jobs released during $[0, S_m)$. Here n_{m-1} denote the number of jobs released between J_i^A and J_i^B . When $r(J_i^A) + T_i > S_m$, i.e., J_i^A is the only job from τ_i that generates demand during $[0, e)$, and in this case $DEM_i(r_A)$ is equal to $dbf_{MM}(J_i^A, S_m, e, m)$. Then we have

$$DEM_i(r_A) = \begin{cases} dbf_{MM}(J_i^A, S_m, e, m) & \text{if } r(J_i^A) + T_i > S_m \\ dbf_{MM}(J_i^A, S_m, e, m) + n_{m-1}C_i^{m-1} + dbf_{MM}(J_i^B, S_m, e, m) & \text{otherwise} \end{cases} \quad (5)$$

If all tasks have integer release times, we can simply get

$$dbf_{MM}(\tau_i, S_m, e, m) = \max_{r_A \in \{0, -1, -2, \dots, -T_i\}} \{DEM_i(r_A)\}$$

For more generic cases, the lemma defines the release pattern when τ_i ($L_i = m - 1$) generates maximum demand during $[0, S_m)$.

Fig. 6: Job release pattern for tasks with $L_i = m - 1$

LEMMA 4.3. *Task τ_i with $L_i = m - 1$ generates maximum demand during $[0, S_m]$ if $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$ or $r(J_i^A) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$ (i.e., deadline $d(J_i^A, m - 1) = e \vee d(J_i^B, m - 1) = e$), and all the jobs are released as soon as possible with period T_i . These two patterns are shown in Figure 6(a) and Figure 6(b), respectively.*

PROOF. The proof for Lemma 4.3 can be found in Appendix A.2. $\square$

Thus using Lemma 4.3, we define $dbf_{MM}(\tau_i, S_m, e, m)$ for task τ_i with $L_i = m - 1$ as follows.

$$dbf_{MM}(\tau_i, S_m, e, m) = \max \{ DEM_i(-D_i^{m-2} + C_i^{m-2}), DEM_i(\text{mod}(e - D_i^{m-1}, T_i) - T_i) \} \quad (6)$$

4.1.2. Case when $L_i \geq m$. The demand generated by task τ_i with $L_i \geq m$ during the interval $[0, e)$ is equal to the sum of demand of all jobs released during $[S_{m-1} = 0, e)$ and the demand of J_i^C as shown in Figure 7.

Let $DEM_i(r_C)$ denote the demand that τ_i generates during $[0, S_m)$ when $r(J_i^C) = r_C$ and all jobs are released as soon as possible with period T_i . Let n_{m-1} denote the number of jobs released during $[S_{m-1}, r(J_i^D))$ and n_m denote the number of jobs released after S_m with deadline D_i^m before e . Given the value of $r(J_i^C)$, $n_{m-1} = \lfloor (S_m - r(J_i^C) - T_i) / T_i \rfloor$. The first job released after J_i^D, J_i^{D+1} , is released at $r(J_i^C) + (n_{m-1} + 2) \times T_i$, and hence $n_m = \lfloor (e - r(J_i^{D+1}) - D_i^m) / T_i \rfloor + 1$. Note that if $n_{m-1} < 0$, i.e., $S_m < r(J_i^C) + T_i$, it implies there does not exist J_i^D . Therefore we have

$$DEM_i(r_C) = \begin{cases} dbf_{MM}(J_i^C, S_m, e, m) + n_{m-1}C_i^{m-1} + dbf_{MM}(J_i^D, S_m, e, m) + n_m C_i^m & \text{if } S_m \geq r(J_i^C) + T_i \\ dbf_{MM}(J_i^C, S_m, e, m) + n_m C_i^m & \text{if } S_m < r(J_i^C) + T_i \end{cases} \quad (7)$$

Fig. 7: Job release pattern for tasks with $L_i \geq m$

If all tasks have integer release times, we can simply get

$$dbf_{MM}(\tau_i, S_m, e, m) = \max_{r_C \in \{0, -1, -2, \dots, -T_i\}} \{ DEM_i(r_C) \}$$

For more generic cases, the lemma defines the release pattern so that τ_i with $L_i \geq m$ generates maximum demand during $[0, e)$.

LEMMA 4.4. *Task τ_i with $L_i \geq m$ generates maximum demand during $[0, e)$ if $r(J_i^C) = \text{mod}(e - D_i^m, T_i) - T_i$ (i.e., the last job released before e has $d(J_i^L, m) = e$) or $r(J_i^C) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$ (i.e., $d(J_i^L, m - 1) = e$) or $r(J_i^C) = -D_i^{m-2} + C_i^{m-2}$ and all the jobs are released as soon as possible with period T_i .*

PROOF. The proof for Lemma 4.4 can be found in Appendix A.2. $\square$

Using Lemma 4.3, we define the $dbf_{MM}(\tau_i, S_m, e, m)$ for task τ_i with $L_i \geq m$ as follows.

$$dbf_{MM}(\tau_i, S_m, e, m) = \max \{ DEM_i(\text{mod}(e - D_i^m, T_i) - T_i), DEM_i(-D_i^{m-2} + C_i^{m-2}), DEM_i(\text{mod}(e - D_i^{m-1}, T_i) - T_i) \} \quad (8)$$

4.2. Demand bound function for task set τ ($dbf_{MM}(\tau, S_m, e, m)$)

Now we have presented the MM dbf for each task. We can simply add up the demand of all tasks in the system to get the dbf for the task set τ .

$$dbf_{MM}(\tau, S_m, e, m) = \sum_{L_i \geq m-1} dbf_{MM}(\tau_i, S_m, e, m) \quad (9)$$

Discussion: The MM dbf derived in this section cannot be directly applied to a more generalized MC model [Ekberg and Yi 2014] where each task has different periods in different criticality modes. The main challenge in this extension is that we need to figure out the execution pattern (i.e., the release time of J_i^A and J_i^C) that can result in worst-case demand in the more generalized model. Of course, we can still compute the demand for all possible release patterns and choose the maximum one, but it can be very computational expensive. For space reason, we would leave this as our future work.

5. VIRTUAL DEADLINE ASSIGNMENT

From the previous sections we know that virtual deadline for each criticality mode plays a key role in shaping the demand of carry-over jobs. The choice of virtual deadline for each task therefore has a significant impact on the performance of the proposed schedulability test. We can decrease τ_i 's demand in L_m ($m > 1$) criticality mode at the cost of increasing the demand in previous modes. By choosing suitable values for D_i^{m-1} for each τ_i , we can increase the chances that τ is schedulable in L_m criticality mode.

The process of finding suitable values for the virtual deadlines is very challenging, because it is infeasible to try all possible combinations of D_i^m for all the tasks in different criticality modes. In this section, we propose a heuristic virtual deadline tuning algorithm which has pseudo-polynomial time complexity as long as the system utilization of each criticality mode is bounded by some constant smaller than 1.

In Section 3 and Section 4 we introduce the existing single-mode dbf-based schedulability test and our proposed multi-mode dbf-based test, respectively. A MC task system τ is schedulable in L_m mode assuming no deadline miss happens in previous criticality modes if one of the following two condition holds.

- (1) $CN_m^S: \forall e \in \{1, 2, \dots, e^{max}\} : dbf_{SM}(\tau, e, m) \leq e$ (Theorem 3.1) .
- (2) $CN_m^M: \forall e \in \{1, 2, \dots, e^{max}\} : \forall S_m \in \{1, 2, \dots, e\} : dbf_{MM}(\tau, S_m, e, m) \leq e$ (Theorem 4.2).

Though the complexity of both of the above conditions is pseudo-polynomial, there is a quadratic increase in complexity to check condition CN_m^M . For a certain e^{max} , while the SM dbf test [Ekberg and Yi 2014] need to calculate the demand for e^{max} times, the MM dbf test need to calculate the demand for $e^{max} \times e^{max}$ times. Therefore, as shown in line 18-27 in Algorithm 1, we use the condition CN_m^M as a complement, i.e., only when condition CN_m^S fails at a certain e , we will check whether condition CN_m^M is satisfied for that time interval length. Thus, once we find an e_f such that $dbf_{SM}(\tau, e_f, m) > e_f$,

we use the Multi-Mode demand bound function to check whether the following holds.

$$\forall e \in \{e_f, e_f + 1, \dots, e^{max}\} : dbf_{MM}(\tau, S_m = e - e_f, e, m) \leq e$$

If the above inequalities hold, then it means no deadline miss happens e_f time units after S_m . Otherwise a candidate task will be chosen and its virtual deadline D_i^{m-1} is reduced by 1.

In Algorithm TuneMode(m), if both CN_m^M and CN_m^S fail at a certain e , then a candidate task is chosen and its virtual deadline D_i^{m-1} is reduced by one unit. By doing this, we can reduce the demand of these tasks in L_m criticality mode. It does this without considering the schedulability in L_{m-1} criticality mode itself, hoping that τ in L_{m-1} criticality mode can later be made schedulable by decreasing deadlines for L_{m-2} criticality mode, i.e., D_i^{m-2} . Once D_i^{m-1} decreases to C_i^{m-1} , τ_i is eliminated from the candidate sets Ψ_m , and then we need to find another candidate task. We repeat the above steps until TuneMode(m) returns true or the candidate task set Ψ_m becomes empty.

An exceptional scenario is that: if we decrease D_i^1 to reduce the demand in L_2 criticality mode, τ may become unschedulable in L_1 criticality mode. However there does not exist D_i^0 , and hence we cannot tune D_i^0 to reduce the demand in L_1 criticality mode. Therefore in this case, TuneMode(m) will undo the changes to deadlines that would make τ unschedulable in L_1 criticality mode. Here undo the changes to deadlines means $D_i^1 = D_i^1 + 1$, and then the candidate task τ_i is removed from Ψ_m .

Algorithm 2 (TuneSystem(τ)) applies Algorithm 1 (TuneMode(m)) on all the criticality modes starting from L_M criticality mode and proceeding in a reverse order, until it has successfully tuned the deadlines in all the criticality modes or failed to do so in some criticality mode. Therefore, the complexity of Algorithm 2 (TuneSystem(τ)) increases linearly as the number of levels increases because the complexity of Algorithm 1 (TuneMode(m)) is independent of the value of m .

Finally, to select a candidate task for deadline tuning in each iteration, TuneMode(m) uses Algorithm 3 (find.candidate.m(Ψ, e)). Note that, GreedyTuning [Ekberg and Yi 2012] uses a very simple metric to choose a candidate task, i.e., τ_i with the maximum

$$\Delta_i = dbf_{SM}(\tau_i, e, m) - dbf_{SM}(\tau_i, e - 1, m)$$

is always chosen as a candidate task. Here Δ_i denote the demand change of τ_i using Single-Mode dbf if $D_i^{m-1} \leftarrow D_i^{m-1} - 1$.

However, there are many other parameters which are also important in choosing a good candidate task. Therefore, in Algorithm 3, we extend this metric to additionally consider other factors in candidate selection such as the impact of change in virtual deadline on the schedulability for previous mode (L_{m-1}). In Section 6, we show that the new metric for choosing a candidate task outperforms the one in [Ekberg and Yi 2014].

To maximize schedulability in L_m mode by reducing the demand of a candidate task, a task with larger Δ_i is preferable. Meanwhile when a candidate task τ_i 's virtual deadline D_i^{m-1} decreases by one, the impact on the demand of L_{m-1} mode is different.

Suppose task set τ has task τ_1 and τ_2 where $T_1 = D_1^2 = D_1^1 = 10000$ and τ_2 has $T_2 = D_2^2 = D_2^1 = 2$. Also, suppose L_2 mode of τ is currently not schedulable but can be tuned to become schedulable if either D_1^1 or D_2^1 decreases by one. The impact of $D_1^1 \leftarrow D_1^1 - 1$ or $D_2^1 \leftarrow D_2^1 - 1$ is different, i.e., $\frac{\Delta D_1^1}{D_1^1} = 0.0001$ and $\frac{\Delta D_2^1}{D_2^1} = 0.5$. If $D_1^1 \leftarrow D_1^1 - 1$, the system demand of L_1 mode stays the same for any time interval length $e (< 9999)$. As a result, the L_1 mode is easier to be schedulable if D_1^1 decreases to 9999 compared to the case if D_2^1 decreases to 1. From the above discussion, we know a task with larger Δ_i and

ALGORITHM 1: TuneMode(m)

```

1  $\Psi_m \leftarrow \{\tau_i | L_i \geq m\};$ 
2 changed  $\leftarrow$  False;
3 while changed = True do
4   changed  $\leftarrow$  False;
5   for  $e_1 \in \{1, 2, \dots, e^{max}\}$  do
6     if  $m = 2$  and  $dbf_{ISM}(\tau, e_1, 1) > e_1$  then
7       if changed = True then
8          $D_i^1 \leftarrow D_i^1 + 1;$ 
9          $\Psi_m.remove(\tau_i);$ 
10        changed  $\leftarrow$  True;
11        break;
12      end
13      else
14        return False
15      end
16    end
17    if  $dbf_{SM}(\tau, e_1, m) > e_1$  then
18       $CN_m^M \leftarrow$  True;
19      for  $e_2 \in \{e_1, e_1 + 1, \dots, e^{max}\}$  do
20        if  $dbf_{MM}(\tau, S_m = e_2 - e_1, e_2) > e_2$  then
21           $CN_m^M \leftarrow$  False;
22          break;
23        end
24      end
25      if  $CN_m^M =$  True then
26        continue;
27      end
28      if  $\Psi_m = \emptyset$  then
29        return False
30      end
31       $\tau_i \leftarrow$  find.candidate.m( $\Psi_m, e$ ) (Algorithm 3);
32       $D_i^{m-1} \leftarrow D_i^{m-1} - 1;$ 
33      if  $D_i^{m-1} < C_i^{m-1}$  then
34         $D_i^{m-1} \leftarrow D_i^{m-1} + 1;$ 
35         $\Psi_m.remove(\tau_i);$ 
36      end
37      else
38        changed  $\leftarrow$  True;
39        break;
40      end
41    end
42  end
43 end
44 return True

```

D_i^{m-1} is more likely to become a better candidate task. Therefore, among all possible candidate tasks in Ψ_m , we use $\Delta_i \times D_i^{m-1}$ as the main metric to choose a candidate task.

Until now, we did not consider case when all tasks have $\Delta_i = 0 \Rightarrow \Delta_i \times D_i^{m-1} = 0$. This means that, there does not exist a candidate task so that the total demand would decrease if $D_i^{m-1} \leftarrow D_i^{m-1} - 1$. As shown in Figure 4, if $\Delta_i = 0$, it must be

ALGORITHM 2: TuneSystem(τ)

```

1  $M \leftarrow \max_{\tau_i \in \tau} \{L_i\};$ 
2 for  $m \in \{M \text{ to } 2\}$  do
3   if  $TuneMode(m) = \text{False}$  then
4     return False
5   end
6 end
7 return True

```

that $len_i = \text{mod}(e, T_i) - (D_i^m - D_i^{m-1}) - C_i^{m-1} > 0$. The demand of τ_i would start to decrease if $D_i^{m-1} \leftarrow D_i^{m-1} - len_i - 1$. Therefore, we also choose those tasks with the $\min\{\max\{0, len_i\}\}$ among those candidate tasks. Hence in Algorithm 3, we sort tasks in Ψ_m with first key $\Delta_i \times D_i^{m-1}$ in ascending order and then with second key $\max\{0, len_i\}$ in descending order.

ALGORITHM 3: find_candidate_m(Ψ, e)

```

1  $Q \leftarrow \text{empty queue};$ 
2 for  $\tau_i \in \Psi$  do
3    $Q.insert(\Delta_i \times D_i^{m-1}, \max\{0, len_i\}, \tau_i);$ 
4 end
5  $Q.sort(\text{ascending, descending});$  // sort tasks in  $\Psi$  with first key  $\Delta_i \times D_i^{m-1}$  in ascending order
  and then with second key  $\max\{0, len_i\}$  in descending order
6  $\tau_i \leftarrow Q.pop();$  // return the corresponding task
7 return  $\tau_i$ 

```

6. EVALUATION

In this section we evaluate the ability of the proposed virtual deadline assignment strategy and the MM dbf based test (i.e, Algorithm 2) to schedule MC task systems. We use acceptance ratios, i.e., the fraction of schedulable task sets, as the metric to evaluate our proposed approach.

We aim to compare with the existing work named GreedyTuning [Ekberg and Yi 2014] based on SM dbf. From extensive experiments, it has already been shown in [Ekberg and Yi 2014] that this work outperforms existing studies (e.g., [Baruah et al. 2012a; Baruah et al. 2011a; Vestal 2007]) for a variety of generic real-time systems. Therefore, through comparison with this SM dbf work, we aim to show in this section that the techniques proposed in this paper also outperform those studies, including the one based on SM dbf test.

We consider MC sporadic tasks scheduled on a dedicated unit speed uniprocessor platform. We will study the impact of varying parameters of tasks on the acceptance ratios of these approaches: 1): GreedyTuning (GT) [Ekberg and Yi 2014], 2): SM dbf test from [Ekberg and Yi 2014] with the improved deadline assignment strategy (i.e, Algorithm 3) (GTI), 3): our improved dbf-based test (IMPT, i.e., Algorithm 2).

6.1. Task Set Generation

Suppose task set τ is a empty task set ($\tau = \emptyset$) initially. Randomly generated tasks are added to the task set τ repeatedly until certain requirements are met. The parameters of each task is controlled by the following parameters.

— $P(m)$ denotes the probability that τ_i has $L_i = m$.

- C_i^1 is drawn using an uniform distribution over $[1, 10]$.
- RC_m denotes the maximum ratio of C_i^m / C_i^{m-1} .
- C_i^m is drawn using an uniform distribution over $[C_i^{m-1}, RC_m \times C_i^{m-1}]$.
- T_i is drawn using an uniform distribution over $[C_i^{L_i}, 200]$.
- D_i is drawn using an uniform distribution over $[D_i^{MIN}, T_i]$ where $D_i^{MIN} = \lfloor C_i^{L_i} + RD \times (T_i - C_i^{L_i}) \rfloor$, and $RD \in [0, 1]$

Let $U_\tau = \max_{m \in \{1, 2, \dots, M\}} \{ \sum_{L_i \geq m} \frac{C_i^m}{T_i} \}$ denote the utilization bound of a task set τ . For a given utilization bound, our generation procedure requires U_τ to fall within the small interval between $[Ubound - \epsilon, Ubound]$ ($\epsilon = 0.005$). As long as $U_\tau < Ubound - \epsilon$, a new task will be randomly generated and added to τ . Once U_τ of τ falls within the range $[Ubound - \epsilon, Ubound]$, the generation procedure for τ is considered complete. However if U_τ becomes greater than $Ubound$ after a new task τ_i is added to τ , we discard the whole task set and start with a new empty task set.

Fig. 8: Dual-Criticality Task Systems with $RD \in \{0.5, 1\}$ and $P(2) = 0.5$

6.2. Evaluation of Dual-Criticality Systems

Figure 8 shows the acceptance ratio as a function of utilization ($Ubound$) of task sets generated with $RC_2 = 3$, $P(1) = P(2) = 0.5$ and $RD \in \{0.5, 1\}$. Each data is based on 1000 randomly generated task sets. As shown in Figure 8, our improved dbf-based schedulability test (IMPT) strictly outperforms GreedyTuning (GT). However the performance gap between GT and IMPT is very small when $RD = 1$, which implies GT already does quite well in scheduling dual-criticality implicit deadline task systems.

Figure 9 shows the acceptance ratios of task sets generated with $RD = 0.5$, $RC_2 = 3$ and $P(2) \in \{0.25, 0.75\}$. The acceptance ratios of task sets with $RD = 0.5$, $RC_2 = 3$ and $P(2) = 0.5$ can be found in Figure 8. We can observe that as $P(2)$ increases from 0.25 to 0.75, the acceptance ratio of GT drops quickly. Even though the SM dbf test with the improved deadline assignment strategy (GTI) always has a higher acceptance ratio than GT, its acceptance ratio is much lower than IMPT when $P(2) = 0.75$. One

Fig. 9: Dual-Criticality Task Systems with $RD = 0.5$ and $P(2) \in \{0.25, 0.75\}$

interpretation for this trend is that SM dbf test is not good at scheduling systems with a larger percentage of high critical tasks. On the other hand, the acceptance ratio of IMPT drops much slower, and its acceptance ratio becomes almost two times as much as GT when $P(2) = 0.75$ and $U_{bound} = 0.95$. In fact, the acceptance ratio of IMPT when $P(2) = 0.75$ is closer to its acceptance ratio when $P(2) = 0.5$.

6.3. Evaluation of Multi-Level Systems

In this section we compare the acceptance ratios of different tests for multi-criticality task systems. Figure 10 shows the acceptance ratios of three level systems with $P(1) = P(2) = P(3) = 1/3$, $RC_2 = RC_3 = 2$ and $RD \in \{0.5, 1\}$. As we can observe, the performance of all the three approaches drops, but IMPT very well outperforms GT even for small utilization bounds.

To study how the acceptance ratio changes as the number of criticality levels increases, we here present the weighted acceptance ratios of different approaches for each criticality level. Suppose $A(U)$ is the acceptance ratio of a certain approach for utilization $U \in [0.55, 0.6, 0.65, \dots, 0.95]$, then the weighted acceptance ratio is defined as

$$\frac{\sum_{U \in \{0.55, 0.6, 0.65, \dots, 0.95\}} A(U) \times U}{\sum_{U \in \{0.55, 0.6, 0.65, \dots, 0.95\}} U}.$$

As we can observe, the gap in weighted acceptance ratio between GT and IMPT is in fact quite small when the criticality level is equal to 2. However, this gap become much larger when the criticality level becomes greater than 2. The acceptance ratio of IMPT also becomes lower (especially for the case when $RD = 0.5$) as the number of criticality level increases. One interpretation for this trend may be that, there is not enough space to tune the virtual deadlines since we need to set multiple virtual deadlines for each task. However, note that, it does not mean that the proposed approach IMPT has very poor performance, and can schedule only a small portion of generated task sets. Since there is no known exact feasibility test for MC task systems, we are unable to eliminate all the infeasible from our experiments. Nevertheless, we can still conclude that the proposed approach in this paper outperforms GT for a variety of generic systems.

Fig. 10: Three Level Task Systems

(a) Weighted Acceptance Ratio $RD = 1$ (b) Weighted Acceptance Ratio $RD = 0.5$

Fig. 11: Weighted Acceptance Ratio

7. CONCLUSIONS

We first introduced the existing single-mode demand bound functions [Ekberg and Yi 2014], which characterize the demand of mixed-criticality sporadic tasks. They use a pessimistic upper bound to characterize demand of carry-over jobs by assuming that the previous criticality mode is schedulable. As a result, the single-mode dbfs over estimate the demand of carry-over jobs. Due to the drawback of the single-mode dbf based test, it has a severe problem that its performance decreases significantly as the number of criticality levels increases.

To avoid the problem of single-mode dbf based schedulability test, we propose multi-mode dbfs which consider the execution demand in the previous criticality mode to determine the remaining execution for carry-over jobs. The proposed multi-mode dbf based test can avoid the problem of the single-mode dbf at the cost of a quadratic increase in the complexity. In practice it could be computationally expensive if we use the multi-mode dbf based test directly. Therefore we propose a novel heuristic deadline tuning algorithm which uses the multi-mode dbf as a complement to reduce the off-line

generalize $dbf_{MM}(J_i^A, S_m, e, m)$ as follows.

$$dbf_{MM}(J_i^A, S_m, e, m) = \begin{cases} \min \{ \min \{ r(J_i^A) + D_i^{m-2}, C_i^{m-2} \} + C_i^{m-1} - C_i^{m-2}, S_m \} & \text{If } \neg(1) \wedge (4) \\ 0 & \text{Otherwise} \end{cases} \quad (10)$$

J_i^B is similar to J_i^A except that we can ignore condition (1) because it can never be true. To maximize J_i^B 's demand, we assume it executes continuously from $r(J_i^B)$. Again if its deadline $d(J_i^B, m-1) > e$, it would generate 0 demand. Hence, we have

$$dbf_{MM}(J_i^B, S_m, e, m) = \begin{cases} \min \{ C_i^{m-1}, S_m - r(J_i^B) \} & \text{If } (4) \\ 0 & \text{Otherwise} \end{cases} \quad (11)$$

Similar to J_i^A , if (1) : $d(J_i^C, m-2) < S_{m-1} = 0 \vee \neg(4)$: $d(J_i^C, m-1) > e$, J_i^C generates 0 demand during $[0, e)$. If $\neg(1) \wedge (2)$, J_i^C would already finish by S_m , and hence would generate demand equal to $C_i^{m-1} - C_i^{m-2} + \min \{ d(J_i^C, m-2) - S_{m-1}, C_i^{m-2} \}$. If $\neg(1) \wedge \neg(2) \wedge (3)$, J_i^C would generate $C_i^m - C_i^{m-2} + \min \{ d(J_i^C, m-2) - S_{m-1}, C_i^{m-2} \}$ demand as shown in Figure 13. If $\neg(1) \wedge \neg(2) \wedge \neg(3) \wedge (4)$, J_i^C would not generate demand after S_m because its deadline $d(J_i^C, m)$ is out of the interval of interest. In this case its demand is bounded by S_m . Hence, we have

Fig. 13: $J_i^C : \neg(1) \wedge \neg(2) \wedge (3)$

$$dbf_{MM}(J_i^C, S_m, e, m) = \begin{cases} 0 & \text{If } (1) \vee \neg(4) \\ C_i^{m-1} - C_i^{m-2} + \min \{ d(J_i^C, m-2), C_i^{m-2} \} & \text{If } \neg(1) \wedge (2) \\ C_i^m - C_i^{m-2} + \min \{ d(J_i^C, m-2), C_i^{m-2} \} & \text{If } \neg(1) \wedge \neg(2) \\ \min \{ S_m, C_i^{m-1} - C_i^{m-2} + \min \{ d(J_i^C, m-2), C_i^{m-2} \} \} & \text{If } \neg(1) \wedge \neg(2) \wedge \neg(3) \wedge (4) \end{cases} \quad (12)$$

J_i^D behaves similar to J_i^C except we can ignore condition (1) because it can never be true. Hence, we have

$$dbf_{MM}(J_i^D, S_m, e, m) = \begin{cases} 0 & \text{If } \neg(4) \\ C_i^{m-1} & \text{If } (2) \\ C_i^m & \text{If } \neg(2) \wedge (3) \\ \min \{ C_i^{m-1}, S_m - r(J_i^D) \} & \text{If } \neg(2) \wedge \neg(3) \wedge (4) \end{cases} \quad (13)$$

Fig. 14: Case when $e \leq D_i^{m-1}$

A.2. Proofs

PROOF FOR LEMMA 4.3. (C1): If $e \leq D_i^{m-1}$, obviously the demand of τ_i maximizes when $r(J_i^A) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$, i.e., $d(J_i^A, m-1) = e$ as shown in Figure 14.

(C2): If $(S_m \leq C_i^{m-1} \wedge e > D_i^{m-1})$, we can find the demand of τ_i maximizes when $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$, which is bounded by S_m as shown in Figure 15. Now we

Fig. 15: Case when $(S_m \leq C_i^{m-1} \wedge e > D_i^{m-1})$

can exclude the above two cases, and we only need to consider the case when $e > D_i^{m-1} \wedge S_m > C_i^{m-1}$.

(C3): If $r(J_i^B) + D_i^{m-1} \leq e$ when $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$ as shown in Figure 16. As

Fig. 16: Case C3

we shift the release pattern left, i.e., $r(J_i^A) = -D_i^{m-2} + C_i^{m-2} + x | x < 0$, the demand change of J_i^A , $\Delta_1(x)$, is of the following form.

$$\Delta_1(x) = \begin{cases} x & \text{If } x \in [-C_i^{m-2}, 0) \\ -C_i^{m-1} & \text{If } x \in [-C_i^{m-2} + D_i^{m-2} - T_i, -C_i^{m-2}) \end{cases}$$

Meanwhile the demand of J_i^B will at most increase linearly, and hence the total demand of τ_i will stay the same or decrease as we shift the pattern left. As we shift the release pattern right, the demand of τ_i will only decrease or stay the same. Therefore in this case the demand of τ_i maximizes when $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$.

(C4): If $r(J_i^B) + D_i^{m-1} > e$ when $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$ as shown in Figure 17. As we shift the release pattern right, then obviously the demand of τ_i would stay the same or decrease. On the other hand as we shift the release pattern left, i.e.,

Fig. 17: Case C4

$r(J_i^A) = -D_i^{m-2} + C_i^{m-2} + x | x < 0$, demand of J_i^B will first increase from 0 to $y_1 = dbf_{MM}(J_i^B, S_m, e, m) | r(J_i^B) + D_i^{m-1} = e$ (assuming at this time $x = x_1$), and then increases linearly to C_i^{m-1} . When $x > x_1$, the demand of J_i^A will decrease but the demand of other jobs (including J_i^B) stay the same. When $x = x_1$, the change demand of J_i^A is equal to $-C_i^{m-1}$ if $x_1 < -C_i^{m-2}$, or x_1 if $x_1 \geq -C_i^{m-2}$. If $x_1 < -C_i^{m-2}$ and $y_1 + x_1 \geq 0$, then the demand of τ_i maximized at this time ($r(J_i^A) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$) because as we further shift the pattern left, the demand of τ_i would only decrease or stay the same. Otherwise if $y_1 + x_1 < 0$, then total demand of τ_i maximizes demand when $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$ because as we further shift the pattern left, the total demand would only stay the same or decrease.

In sum the demand of $\tau_i | L_i = m - 1$ during $[0, e)$ is maximized if $r(J_i^A) = -D_i^{m-2} + C_i^{m-2}$ or $r(J_i^A) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$.

□

PROOF FOR LEMMA 4.4. (C1): If $e - S_m > D_i^m$, the demand of $\tau_i | L_i \geq m$ during $[0, e)$ is maximized when $r(J_i^C) = \text{mod}(e - D_i^m, T_i) - T_i$ (the last job released before e , J_i^L , has $d(J_i^L, m) = e$) as shown in Figure 18. This is because as we shift the release

Fig. 18: Case when $e - S_m > D_i^m$

pattern right, the demand of J_i^L would decrease from C_i^m to 0 while the increase in demand of other jobs including J_i^C and J_i^D is bounded by C_i^m . On the other hand as we shift the release pattern left, the total demand of all jobs would decrease or stay the same.

(C2): If $e - S_m < D_i^m - D_i^{m-1}$, no job of τ_i could execute more than C_i^{m-1} (no job will execute after S_m). Thus τ_i behaves as a task with $L_i = m - 1$, and from Lemma 4.3 we know the demand of τ_i is maximized when $r(J_i^C) \in \{-D_i^{m-2} + C_i^{m-2}, \text{mod}(e - D_i^{m-1}, T_i) - T_i\}$.

(C3): If $D_i^m - D_i^{m-1} \leq e - S_m \leq D_i^m$, at most one job (either J_i^C or J_i^D) could generate execution demand as much as C_i^m as shown in Figure 19. Thus suppose initially $r(J_i^C) = \text{mod}(e - D_i^m, T_i) - T_i$, i.e., the scenario of Figure 19. If we shift the release pattern left, the demand of τ_i would only decrease or stay the same. On the other hand as we shift the release pattern right, no job would execute after S_m , and hence τ_i be-

Fig. 19: Case when $D_i^m - D_i^{m-1} \leq e - S_m \leq D_i^m$

has as a task with $L_i = m - 1$. Therefore from Lemma 4.3 we know the demand of τ_i is maximized when $r(J_i^C) \in \{-D_i^{m-2} + C_i^{m-2}, \text{mod}(e - D_i^{m-1}, T_i) - T_i\}$.

In sum the demand of $\tau_i | L_i \geq m$ during $[0, e]$ is maximized if $r(J_i^C) = -D_i^{m-2} + C_i^{m-2}$ or $r(J_i^C) = \text{mod}(e - D_i^{m-1}, T_i) - T_i$ or $r(J_i^C) = \text{mod}(e - D_i^m, T_i) - T_i$. $\square$

A.3. Upper Bound of time interval length

For task τ_i with $L_i = m - 1$, we can observe that its demand during $[0, S_m)$ is upper bounded by $(\frac{S_m - 2 \times C_i^{m-1}}{T_i} + 2) \times C_i^{m-1}$. For task τ_i with $L_i \geq m$, its demand during $[0, e]$ is upper bounded by

$$\begin{aligned} & \underbrace{C_i^{m-1}}_{\geq dbf_{MM}(J_i^C, S_m, e, m)} + \underbrace{\frac{S_m - C_i^{m-1}}{T_i} \times C_i^{m-1}}_{\geq n_{m-1} \times C_i^{m-1}} + \underbrace{C_i^m}_{\geq dbf_{MM}(J_i^D, S_m, e, m)} + \underbrace{\frac{e - S_m - D_i^m + T_i}{T_i} \times C_i^m}_{\geq n_m \times C_i^m} \\ &= C_i^{m-1} \times \frac{S_m - C_i^{m-1} + T_i}{T_i} + \frac{e - S_m - D_i^m + 2 \times T_i}{T_i} \times C_i^m \end{aligned}$$

Suppose $dbf_{MM}(\tau, S_m, e, m) > e$, then it must be that

$$\begin{aligned} & \sum_{L_i \geq m} \left(C_i^{m-1} \times \frac{S_m - C_i^{m-1} + T_i}{T_i} + \frac{e - S_m - D_i^m + 2 \times T_i}{T_i} \times C_i^m \right) + \sum_{L_i = m-1} \left(\frac{S_m + 2 \times (T_i - C_i^{m-1})}{T_i} \right) \times C_i^{m-1} \\ &= S_m \times \underbrace{\left(\sum_{L_i = m-1} \frac{C_i^{m-1}}{T_i} + \sum_{L_i \geq m} \frac{C_i^{m-1} - C_i^m}{T_i} \right)}_{\text{Exp A}} + \underbrace{\sum_{L_i = m-1} \frac{T_i - C_i^{m-1}}{T_i} \times 2 \times C_i^{m-1}}_{\text{Exp B}} \\ &+ \underbrace{\left(\sum_{L_i \geq m} \frac{2 \times T_i - D_i^m}{T_i} \times C_i^m + \frac{T_i - C_i^{m-1}}{T_i} \times C_i^{m-1} \right)}_{\text{Exp C}} + \sum_{L_i \geq m} \frac{C_i^m}{T_i} \times e > e \end{aligned}$$

If $\text{Exp A} > 0$, the total demand is maximized if $S_m = e$, and else if $\text{Exp A} \leq 0$, the total demand is maximized if $S_m = 0$. Therefore the value of e is bounded by $(\text{Exp B} + \text{Exp C}) / (1 - \sum_{L_i \geq m} \frac{C_i^m}{T_i})$ or $(\text{Exp B} + \text{Exp C}) / (1 - \sum_{L_i \geq m} \frac{C_i^m}{T_i} - \text{Exp A})$

REFERENCES

- 2008. "ARINC653 - An Avionics Standard for Safe, Partitioned Systems". Wind River Systems / IEEE Seminar.
- N Audsley. 1991. Optimal Priority Assignment and Feasibility of Static Priority Tasks with Arbitrary Start Times. In *The University of York Technical Report*.
- S Baruah, V Bonifaci, G D'Angelo, H Li, and A Marchetti-Spaccamela. 2012a. The Preemptive Uniprocessor Scheduling of Mixed-Criticality Implicit-Deadline Sporadic Task Systems. In *ECRTS*.

- S Baruah, V Bonifaci, G D'Angelo, H Li, A Marchetti-Spaccamela, N Megow, and L Stougie. 2012b. Scheduling Real-Time Mixed-Criticality Jobs. *IEEE Transactions on Computers* 61, 8 (2012), 1140–1152.
- Sanjoy Baruah, Vincenzo Bonifaci, Gianlorenzo D'angelo, Haohan Li, Alberto Marchetti-Spaccamela, Suzanne Van Der Ster, and Leen Stougie. 2015. Preemptive Uniprocessor Scheduling of Mixed-Criticality Sporadic Task Systems. *Journal of the ACM (JACM)* 62, 2 (2015), 14.
- S Baruah, V Bonifaci, G D'Angelo, A Marchetti-Spaccamela, S Van Der Ster, and L Stougie. 2011a. Mixed-criticality Scheduling of Sporadic Task Systems. In *Proceedings of the European Symposium on Algorithms (ESA)*.
- S Baruah and A Burns. 2011. Implementing Mixed Criticality Systems in Ada. In *Proceedings of the Ada-Europe International Conference on Reliable Software Technologies*. 174–188.
- S Baruah, A Burns, and R Davis. 2011b. Response-Time Analysis for Mixed Criticality Systems. In *RTSS*. 34–43.
- S Baruah, H Li, and L Stougie. 2010. Towards the Design of Certifiable Mixed-criticality Systems. In *RTAS*. 13–22.
- S Baruah, A Mok, and L Rosier. 1990. Preemptively Scheduling Hard-Real-Time Sporadic Tasks on One Processor. In *RTSS*. 182–190.
- S Baruah and S Vestal. 2008. Schedulability Analysis of Sporadic Tasks with Multiple Criticality Specifications. In *ECRTS*. 147–155.
- Iain Bate, Alan Burns, and Robert I Davis. 2015a. A bailout protocol for mixed criticality systems. In *Real-Time Systems (ECRTS), 2015 27th Euromicro Conference on*. IEEE, 259–268.
- Iain Bate, Alan Burns, and Robert I. Davis. 2015b. A Bailout Protocol for Mixed Criticality Systems. In *27th Euromicro Conference on Real-Time Systems, ECRTS 2015, Lund, Sweden, July 8-10, 2015*. 259–268.
- A. Burns, T. Fleming, and S. Baruah. 2015. Cyclic Executives, Multi-core Platforms and Mixed Criticality Applications. In *Real-Time Systems (ECRTS), 2015 27th Euromicro Conference on*. 3–12.
- Arvind Easwaran. 2013. Demand-based Scheduling of Mixed-Criticality Sporadic Tasks on One Processor. In *RTSS*. 78–87.
- P Ekberg and W Yi. 2012. Bounding and Shaping the Demand of Mixed-Criticality Sporadic Tasks. In *ECRTS*. 135–144.
- Pontus Ekberg and Wang Yi. 2014. Bounding and shaping the demand of generalized mixed-criticality sporadic task systems. *Real-Time Systems* 50, 1 (2014), 48–86.
- Thomas Fleming. 2013. *Extending mixed criticality scheduling*. Ph.D. Dissertation. University of York.
- Xiaozhe Gu, A. Easwaran, Kieu-My Phan, and Insik Shin. 2015. Resource Efficient Isolation Mechanisms in Mixed-Criticality Scheduling. In *Real-Time Systems (ECRTS), 2015 27th Euromicro Conference on*. 13–24.
- Pengcheng Huang, Georgia Giannopoulou, Nikolay Stoimenov, and Lothar Thiele. 2014. Service Adaptions for Mixed-Criticality Systems. In *In Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*.
- Pengcheng Huang, Pratyush Kumar, Nikolay Stoimenov, and Lothar Thiele. 2013. Interference Constraint GraphA new specification for mixed-criticality systems. In *Emerging Technologies & Factory Automation (ETFA), 2013 IEEE 18th Conference on*. IEEE, 1–8.
- H Li and S Baruah. 2010. An Algorithm for Scheduling Certifiable Mixed-Criticality Sporadic Task Systems. In *RTSS*. 183–192.
- Aloysius K Mok. 1983. Fundamental design problems of distributed systems for the hard-real-time environment. (1983).
- Taeju Park and Soontae Kim. 2011. Dynamic scheduling algorithm and its schedulability analysis for certifiable dual-criticality systems. In *Proceedings of the ninth ACM international conference on Embedded software*. ACM, 253–262.
- John A Stankovic, Marco Spuri, Krithi Ramamritham, and Giorgio C Buttazzo. 1998. Introduction. In *Deadline Scheduling for Real-Time Systems*. Springer, 1–11.
- Hang Su and Dakai Zhu. 2013. An Elastic Mixed-criticality Task Model and Its Scheduling Algorithm. In *Proceedings of the Conference on Design, Automation and Test in Europe (DATE '13)*. EDA Consortium, San Jose, CA, USA, 147–152. <http://dl.acm.org/citation.cfm?id=2485288.2485325>
- S Vestal. 2007. Preemptive Scheduling of Multi-criticality Systems with Varying Degrees of Execution Time Assurance. In *RTSS*.