



CISTER

Research Centre in
Real-Time & Embedded
Computing Systems

Technical Report

Technical Report: Techniques and Analysis for Mixed-criticality Scheduling with Mode-dependent Server Execution Budgets

Muhammad Ali Awan

Konstantinos Bletsas

Pedro F. Souto

Benny Akesson

Eduardo Tovar

CISTER-TR-190202

Technical Report: Techniques and Analysis for Mixed-criticality Scheduling with Mode-dependent Server Execution Budgets

Muhammad Ali Awan, Konstantinos Bletsas, Pedro F. Souto, Benny Akesson, Eduardo Tovar

CISTER Research Centre

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8321159

E-mail:

<https://www.cister-labs.pt>

Abstract

In mixed-criticality systems, tasks of different criticality share system resources, mainly to reduce cost. Cost is further reduced by using adaptive mode-based scheduling arrangements, such as Vestal's model, to improve resource efficiency, while guaranteeing schedulability of critical functionality. To simplify safety certification, servers are often used to provide temporal isolation between tasks. However, a server's computational requirements may greatly vary in different modes, but state-of-the-art techniques and schedulability tests do not allow different budgets to be used by a server in different modes. This results in a single conservative execution budget for all modes, increasing system cost.

The goal of this paper is to reduce the cost of mixed-criticality systems through three main contributions: (i) a scheduling arrangement for uniprocessor systems employing fixed-priority scheduling within periodic servers, whose budgets are dynamically adjusted at run-time in event of a mode change, (ii) a new schedulability analysis for such systems, and (iii) heuristic algorithms for assigning budgets to servers in different modes and ordering the execution of the servers. Experiments with synthetic task sets demonstrate considerable improvements (up to 52.8%) in scheduling success ratio when using dynamic server budgets, compared to static "one-size-fits-all-modes" budgets.

1 Technical Report: Techniques and Analysis for 2 Mixed-criticality Scheduling with 3 Mode-dependent Server Execution Budgets

4 **Muhammad Ali Awan** 

5 CISTER Research Centre and ISEP, Porto, Portugal
6 muaan@isep.ipp.pt

7 **Konstantinos Bletsas** 

8 CISTER Research Centre and ISEP, Porto, Portugal
9 ksbs@isep.ipp.pt

10 **Pedro F. Souto** 

11 University of Porto, Faculty of Engineering and CISTER Research Centre, Porto, Portugal
12 pfs@fe.up.pt

13 **Benny Akesson** 

14 ESI (TNO), Eindhoven, the Netherlands
15 benny.akesson@tno.nl

16 **Eduardo Tovar** 

17 CISTER Research Centre and ISEP, Porto, Portugal
18 emt@isep.ipp.pt

19 Abstract

20 In mixed-criticality systems, tasks of different criticality share system resources, mainly to reduce cost. Cost
21 is further reduced by using adaptive mode-based scheduling arrangements, such as Vestal's model, to improve
22 resource efficiency, while guaranteeing schedulability of critical functionality. To simplify safety certification,
23 servers are often used to provide temporal isolation between tasks. However, a server's computational require-
24 ments may greatly vary in different modes, but state-of-the-art techniques and schedulability tests do not allow
25 different budgets to be used by a server in different modes. This results in a single conservative execution budget
26 for all modes, increasing system cost.

27 The goal of this paper is to reduce the cost of mixed-criticality systems through three main contributions:
28 (i) a scheduling arrangement for uniprocessor systems employing fixed-priority scheduling within periodic
29 servers, whose budgets are dynamically adjusted at run-time in event of a mode change, (ii) a new schedulability
30 analysis for such systems, and (iii) heuristic algorithms for assigning budgets to servers in different modes
31 and ordering the execution of the servers. Experiments with synthetic task sets demonstrate considerable
32 improvements (up to 52.8%) in scheduling success ratio when using dynamic server budgets, compared to static
33 "one-size-fits-all-modes" budgets.

34 **2012 ACM Subject Classification** Computer systems organization → Real-time systems; Computer systems
35 organization → Real-time operating systems; Computer systems organization → Real-time system architecture

36 **Keywords and phrases** Mixed-criticality scheduling, Varying execution server, Vestal model

37 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

38 1 Introduction

39 Mixed-criticality systems are an important niche of real-time embedded systems. Their defining
40 characteristic is the fact that computing tasks of different criticalities execute on the same hardware



© Muhammad Ali Awan, Pedro F. Souto, Konstantinos Bletsas, Benny Akesson, and Eduardo Tovar;
licensed under Creative Commons License CC-BY

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:30

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

41 and share system resources, typically in order to reduce system costs¹. A task's criticality is a measure
42 of the severity of the consequences of a task failing which in the context of real-time scheduling
43 means missing its deadline. The higher it is, the more conservative, and costlier, in terms of effort,
44 time and money, the approach employed to upper-bound that task's worst-case execution time.

45 Crucially for the certifiability of a mixed-criticality system, sufficient isolation must exist by
46 design between the timing behavior of different applications. Namely, the timing behavior of one
47 application must not possibly compromise the timeliness of a different application, especially if the
48 former has low criticality and the latter has high criticality. If tasks of different criticalities share
49 the same resources, even on a uniprocessor platform, which is the focus of this work, they must
50 all be engineered to the same strict standard of safety as the highest-criticality task among them.
51 This is evidently grossly resource-inefficient and such over-engineering can have severe real-world
52 costs. Most mixed-criticality systems are embedded, often meaning that profit margins for the end
53 product can be thin and/or that engineering constraints on size, weight and power (SWaP) can be
54 tight. Examples of such domains include automotive, aerospace and avionics.

55 Fortunately, the relevant guidelines (e.g., [1, 2] for the avionics domain), do not insist in zero
56 interference among mixed-criticality applications, but instead expect any intra-application interference
57 to be carefully accounted for and adequately mitigated. Two scheduling arrangements that can be
58 useful to a designer faced with the above considerations, and trying to achieve both safety and
59 efficiency, are *servers* and adaptive *mode-based* scheduling.

60 Servers provide scheduling isolation via *time partitioning*. In the simplest arrangement, a server
61 is a periodically repeating fixed-length contiguous time window on a given core. Only the tasks
62 served by the particular server are allowed to use the core within the confines of that time window,
63 scheduled under, e.g., a fixed-priority or an EDF policy. Conversely, task are not allowed to execute
64 at all outside the confines of their respective server's time window. This arrangement both provides
65 a predictable supply of processing time to the set of tasks served *and* also ensures that they cannot
66 interfere with other applications (tasks served by other servers). Crucially, the implications of any
67 task misbehaving temporally are localised to the respective server.

68 Meanwhile, mode-based scheduling arrangements [6] based on Vestal's model [35] can be applied
69 in order to more efficiently use the available processing capacity. Rather than using extremely
70 pessimistic worst-case execution time (WCET) estimates for both low- and high-criticality tasks that
71 interact, this approach uses less pessimistic estimates, by default – *probably*, but not *provably* safe.
72 In the statistically unlikely case of a task executing for longer than its assumed WCET, a carefully
73 managed *mode change* is triggered. The less critical (and/or less important [9]) tasks, as specified at
74 design time, are dispensed with. The remaining tasks must then be provably schedulable, assuming
75 more pessimistic WCET estimates. In the general case, there can be many such mode changes. A
76 mode switch is *not* a failure – it constitutes system behavior explicitly accounted for at design time
77 (including the set of tasks to drop, and the implications of doing so). This adaptive mode-based model
78 improves resource efficiency without compromising the system requirements.

79 Servers and Vestal's mode-based model can be combined. In this work, we consider a time-
80 partitioned system, with multiple servers that share the same period. To each server, one or more
81 mixed-criticality applications are assigned, in turn consisting of multiple tasks. Every server is
82 scheduled using fixed-priority scheduling (which is known as AMC [5], in the context of Vestal's
83 model). The use of Vestal's model can reduce the processing budget requirements for the different
84 servers (compared to a naive approach that would always assume pessimistic WCETs for all tasks)
85 and the use of servers can provide timing isolation between applications assigned to different servers.

¹ When tasks of different criticalities exist but are completely isolated, such systems are *multiple-criticality*, as opposed to mixed-criticality, and they constitute a different class of systems. See Footnote 1 in [12] and in [11].

86 However, it can still be inefficient if the execution time budgets used for the servers remain the
87 same in different modes, because **a given server can have very different processing needs in one**
88 **mode than in another**. This realisation motivates the present work, which considers a server- and
89 AMC-based scheduling arrangement whereby the server budgets are dynamically adjusted, at mode
90 change, for greater resource efficiency, at no detriment to the predictability of the system and the
91 provision of safety guarantees. Such an arrangement requires new analysis, because, even if the
92 original analysis for AMC can be applied to periodic server-based scheduling with minor changes²,
93 this is no longer the case when the server budgets change, in response to a mode change. This is a
94 short-coming of the state-of-the-art. By proposing varying server budgets and providing analysis for
95 this arrangement, our work hence allows for greater resource efficiency and cost savings.

96 The main contributions of our work are the following: First, we formulate new schedulability
97 analysis for uniprocessor systems using periodic servers with AMC as their scheduling policy and
98 whose execution time budgets are dynamically adjusted in response to a mode change. Secondly,
99 we discuss the complex interdependencies between the parameters of different servers and propose
100 heuristics for the ordering of servers in the schedule and the assignment of server execution budgets
101 in the different modes, for good schedulability performance. Thirdly, we explore via experiments
102 with synthetic task sets, the schedulability performance of dynamic server budgets under different
103 server orderings and budget assignment heuristics, compared to the baseline of static-budget servers.
104 The results strongly validate our approach by demonstrating up to 52.8% improvement in scheduling
105 success ratio over the baseline heuristic.

106 The rest of this paper is structured as follows. Section 2 discusses related work. Section 3 presents
107 the task model and the server scheduling model assumed. Section 4 contains the formulation of our
108 new scheduling analysis. Section 5 enumerates the different scheduling arrangements and heuristics
109 considered in our experimental evaluation, while Section 6 contains the evaluation itself and discusses
110 our findings. Section 7 concludes.

111 **2 Related work**

112 In the literature, several works try to combine an adaptive mixed-criticality task scheduling model
113 with servers in different ways. There also exist other works that are more loosely related.

114 Papadopoulos et al. [31] consider the adaptive mixed-criticality model of Vestal [6] with two
115 criticality levels and a uniprocessor whose processing time supply is partitioned into two servers, one
116 for the tasks of each criticality level, with corresponding target execution time budgets. At run-time,
117 the execution times of the tasks are monitored, and in a control feedback loop, the execution budgets
118 of each successive server instance are adjusted, according to the principles of control theory [32].
119 The budget adjustments may introduce some deviation from strict server periodicity. The aim of this
120 approach is to preserve as much service to lower-criticality tasks as can be afforded by the state of the
121 system, instead of dropping them altogether when a mode change occurs.

122 Ren et al. [33] consider the same task model [6] with similar aims, but with a different approach.
123 In a multicore, each core is assigned different high- and low-criticality tasks (i.e., task to core
124 partitioning). On each core, disjoint groups of one high-criticality task and one or more low-criticality
125 tasks are formed. For each group, different servers for the high- and low-criticality tasks are employed,
126 such that the provision of scheduling guarantees to the former has the least impact on the latter. At
127 the task group level, groups are scheduled on the corresponding core using EDF.

² We are referring to the use of a “fake interfering task”, exactly as done for EDF servers in [3], that models the periodic unavailability of the server.

128 The on-the-fly fast overrun budgeting mechanism by Hu et al. [23] improves the system’s quality
129 of service for low-criticality tasks by postponing the mode-switch instance. This approach (inspired
130 by procrastination techniques [4, 29]) utilises the collected static and dynamic slack for job overruns.
131 Similarly, Gu et al. [18] compute the sufficient L-mode budget for high-criticality applications
132 collectively at design time. This budget is utilised at run time to schedule the high-criticality
133 applications in L-mode with the objective of postponing the mode switch as much as possible. The
134 budget assignment can be tuned for system-wise objectives of schedulability and service guarantees for
135 low-criticality applications. Similarly, Hu et al. [22] regulate the low-criticality workload considering
136 the online demand of high-criticality applications with the objective of improving the quality of
137 service for low-criticality applications.

138 Lipari and Buttazzo’s [17] reservation-based approach, assigns to each high-criticality task a
139 server with a computation bandwidth equal to its high-criticality-mode utilisation. A single low-
140 criticality server (initially assigned the leftover utilisation) serves all low-criticality tasks. The
141 bandwidth reclaimed from high-criticality servers is assigned to the low-criticality server. Fei et
142 al. [20] adjust server budgets recurrently, based on a predictor of future job execution times. Evripidou
143 and Burns [16] consider a uniprocessor platform with multiple partitions. There exist a periodic server
144 for the periodic tasks and a deferrable higher-priority server for the sporadic tasks in each partition.
145 Missimer et al. [30] similarly employ sporadic servers and priority-inheritance bandwidth-preserving
146 servers integrate I/O- and task-scheduling, albeit under a fixed-priority scheduling policy.

147 Gu et al. [19] focused on component-based systems. Within a component, as long as the number
148 of low-criticality mode execution time overruns does not exceed its predefined tolerance level, other
149 tasks (of any criticality) in other components are unaffected. Collectively, tasks are scheduled with an
150 EDF-based policy. Some recent studies [13, 21, 24, 27] explored the implementation-level details of
151 scheduling approaches (including hierarchical ones) for mixed-criticality systems and how to combine
152 adaptive mixed-criticality scheduling with predictable hardware.

153 In [3], we considered different server-based arrangements for strict isolation between criticalities
154 and good schedulability on multicores, using scaled-deadline EDF [15]. One approach uses one server
155 per core for high-criticality tasks and uses the spare capacity for low-criticality task servers. At mode
156 change, the latter are dropped and the high-criticality servers can use the entire core. Another approach
157 uses fixed-budget servers serving tasks of mixed-criticalities. At mode change, low-criticality tasks in
158 each server are dropped but the servers and their budgets persist, serving the remaining tasks.


159 By comparison with [3], our present work brings fully adjustable server budgets at mode change
160 and targets a fixed-priority-scheduled [5] uniprocessor platform. Multiple servers exist, each serving
161 tasks of mixed criticalities. At mode change, the low-criticality tasks are still discarded, however the
162 server budgets are adjusted (some upwards, others downwards), to account for the different processing
163 needs in the new mode. This improves the efficiency in the use of processing capacity, allowing more
164 demanding task sets to be schedulable, without using a faster processor.

165 **3 Task model and system model**

166 **3.1 Task model**

167 This paper assumes the established adaptive variant of Vestal’s mixed-criticality model, with execution
168 time monitoring and mode changes [5]. In particular, we assume a set $\tau \stackrel{\text{def}}{=} \{\tau_1, \dots, \tau_n\}$ of n mixed-
169 criticality sporadic tasks. Each task has a minimum interarrival time T_i , a relative deadline D_i that is
170 constrained (i.e., $D_i \leq T_i$) and a criticality level κ_i . In the general case, these tasks may be grouped
171 together into disjoint applications, possibly consisting of tasks of different criticalities.

172 At run-time, the system operation is based on different *modes* wherein only tasks of a given
173 criticality or higher execute. For each task, different execution time estimates are assumed with



■ **Figure 1** At mode change, the L-tasks are dropped and the remaining H-tasks must be schedulable as long as they execute for up to their C_i^H estimates (including any jobs thereof caught up in the mode change). However, server execution budgets are only adjusted from X_q^L to X_q^H at the start of the next timeslot.

174 corresponding confidence in their safety. For simplicity, in this paper, we consider only two criticality
 175 levels, high (H) and low (L), hence two modes of operation (L-mode and H-mode). The H-mode
 176 worst-case execution time estimate (H-WCET) of a task τ_i is denoted by C_i^H and it is demonstrably
 177 unexceedable, but typically very pessimistic. The L-mode worst-case execution time estimate (L-
 178 WCET) is denoted as $C_i^L \leq C_i^H$. The system boots in L-mode, wherein all tasks execute and all their
 179 deadlines must be met. If any task attempts to execute for more than its execution time estimate for
 180 that mode, a *mode change* is triggered, whereupon all low-criticality tasks (L-tasks) are dispensed
 181 with. In H-mode, only the H-tasks execute and the deadlines of all jobs by H-tasks must be met,
 182 including any jobs released before the mode change.

183 In this work, we consider scheduling based on fixed priorities. This implies that each task has
 184 a unique static priority assigned to it, which is the basis of scheduling decisions. Fixed-priority
 185 scheduling, in the context of the above mixed-criticality model is known as AMC. However, in our
 186 work, we assume that tasks are partitioned to *servers* that use AMC as their internal scheduling policy.
 187 We next introduce the server model that we assume.

188 3.2 Server-based system model

189 Consider a uniprocessor platform and a set of periodic servers $\{\tilde{P}_q\}$, $q = 1, 2, \dots, Q$ assigned to it.
 190 All servers share the same period S (called the “timeslot length”, to stick to the terminology used in
 191 related work [3]) and they execute one after the other in the same order, in a form of cyclic executive
 192 with a periodicity of S . This fixed order in which the servers execute is specified by the designer, at
 193 design time. Each server \tilde{P}_q is assigned a mixed-criticality set of tasks $\tau[\tilde{P}_q] \subseteq \tau$ which are scheduled
 194 within the server according to their fixed priorities.

195 The system conforms to the task model defined in Section 3.1, meaning that there exist two
 196 modes, L and H. Each server \tilde{P}_q has a fixed time budget X_q^L in the L-mode and a respective fixed
 197 time budget X_q^H for the H-mode. Additionally, $\sum_{q=1}^Q X_q^L \leq S$ and $\sum_{q=1}^Q X_q^H \leq S$ (i.e., in each
 198 mode, the servers are sized such that they fit into the timeslot S). When the system is in L-mode, all
 199 servers execute with their X_q^L budgets and all tasks present in the L-mode must provably meet their
 200 deadlines under those budgets, as long as no job executes for more than its C_i^L . However, if such
 201 an execution overrun occurs, a transition to H-mode is triggered. Then, all L-tasks are *immediately*
 202 dispensed with. The remaining tasks (including any jobs thereof released before the mode change)
 203 must meet their deadlines, assuming they can execute for up to their C_i^H estimates. Additionally,
 204 the server budgets are adjusted to their respective X_q^H values *at the start of the next timeslot*. This
 205 implies that there is a time interval of s' time units ($0 \leq s' < S$) after the mode change, during which
 206 the system is already in H-mode, but the server budgets are not yet adjusted from the values (X_i^L)
 207 used in the L-mode. Figure 1 illustrates this arrangement via an example schedule.

208 **4** Schedulability Analysis

209 **4.1** Schedulability analysis for an individual server

210 In this subsection, we are going to derive a sufficient schedulability test for a server conforming
 211 to the model described earlier. More specifically, given as input the tasks assigned to a server \tilde{P}_i ,
 212 its period S and its execution time budgets (X_i^L and X_i^H) and starting offsets (O_i^L and O_i^H) for
 213 the two modes, our analysis will establish whether the server is mixed-criticality-schedulable. The
 214 questions of how these attributes (X_i^L , X_i^H , O_i^L and O_i^H) are determined for each server and how the
 215 derivations of these attributes of different servers inter-depend is discussed later, in Section 4.2, where
 216 the schedulability test for the entire system is formulated. Our analysis builds upon AMC-max [5]
 217 and tests the schedulability of a task (i) in L-mode and (ii) in H-mode separately.

218 **4.1.1** Steady L-mode analysis

219 In L-mode (i.e., prior to the occurrence of a mode switch), tasks behave as conventional Liu-and-
 220 Layland tasks with a WCET of C_i^L . Therefore, as in AMC-max, to test the schedulability of a task
 221 τ_i in L-mode, we use the standard worst-case response time (WCRT) recurrence [26]. However, as
 222 in [34] and [3], we add a “fake” top-priority periodic interfering task τ_f that equivalently models the
 223 fact that the tasks do not execute directly on the processor, but within a periodic server:

$$224 \quad R_i^L = C_i^L + \sum_{\tau_j \in hp(i)} \left\lceil \frac{R_i^L}{T_j} \right\rceil C_j^L + \underbrace{\left\lceil \frac{R_i^L}{S} \right\rceil (S - X^L)}_{\text{fake task}} \quad (1)$$


225 In (1), the server index is omitted for clarity of presentation and $hp(i)$ is the set of higher-priority
 226 tasks served by the same server as τ_i . The fake task’s WCET is $(S - X^L)$ (equal to the time interval
 227 between the ending of one server instance and the start of the next instance of the same server); its
 228 interarrival time is S . A server \tilde{P}_q is schedulable in L-mode if all of its tasks are schedulable in that
 229 mode (i.e., if $R_i^L \leq D_i \forall i \in \tau[\tilde{P}_q]$).

230 **4.1.2** Schedulability testing in H-mode

231 To test for the schedulability of an H-task in the event of a mode change, which, in the general case,
 232 also entails a potential server budget change (though not necessarily coincident), we have to consider
 233 four mutually exclusive and jointly exhaustive cases (elaborated below). The task under analysis must
 234 meet its deadline in all of those cases. The reason for having to consider four separate cases is the
 235 following: If the server budgets change after a mode change (not necessarily immediately), then the
 236 worst-case scenario, maximising the response time of a job that completes in the H-mode, does not
 237 necessarily involve that job being released before the mode change – unlike what holds for classic
 238 AMC.

239 Let $t_{hp(i)}^{idle}$ denote the first instant after the mode change that the server is active and no task in
 240 $hp(i)$ is executing inside it. Note that, in the general case $t_{hp(i)}^{idle}$ might be located before the server
 241 budget change instant or after it, or may coincide with it. The four cases to consider are:

- 242 ■ **Case 1:** The H-task under consideration is released before the mode change instant s (or even at
 243 the mode change instant, as a corner case), but completes after the mode change.
- 244 ■ **Case 2:** The H-task is released after the mode change, but before $t_{hp(i)}^{idle}$.
- 245 ■ **Case 3:** The H-task is released at or after $t_{hp(i)}^{idle}$ **and** also before the server budget change instant.



■ **Figure 2** All possible relative orderings between the mode change instant (s), the budget change instant (b), the arrival time (a_i) of the H-task under analysis and $t_{hp(i)}^{idle}$, for the four cases.


246 ■ **Case 4:** The H-task is released at or after $t_{hp(i)}^{idle}$ **and** also at or after the budget change instant.

247 Figure 2 depicts the possible relative orderings between the mode change instant (s), the budget
 248 change instant (b), the arrival time (a_i) of the H-task under analysis and $t_{hp(i)}^{idle}$, for the four cases.

249 **Case 1:** This case considers H-tasks that are caught in their execution window by the mode
 250 switch, and hence may suffer the interference both from L-tasks and H-tasks of higher priority. The
 251 interference from such tasks can be upper-bounded according to the existing AMC-max analysis.
 252 However, the task in consideration also suffers interference from the unavailability of the server
 253 (which we model as a fake task). Below, we will upper-bound that separately for in each mode.

254 Assume that the mode change occurs some s' time units after the last timeslot boundary. Then,
 255 $s' \in [0, S)$. Let $W_f^{H|tr.}(s', \Delta t)$ denote the worst-case workload function of the fake task modelling
 256 the unavailability of the reserve, for the interval $[s, s + \Delta t)$. The first parameter denotes the phasing
 257 of the mode change relative to the timeslot boundary, as explained earlier. Since s' cannot be known
 258 offline (but only established *a posteriori*), to upper-bound that workload function in the general case
 259 would need to upper bound $W_f^{H|tr.}(s', \Delta t)$ for all $s' \in [0, S)$. Fortunately, since the scheduling is by
 260 fixed priorities, we need only consider two values for s' . Namely, $s' = 0$ and $s' = O^L + X^L$. The
 261 first value ($s' = 0$) involves a mode change coincident with a timeslot boundary; the server is denied
 262 the processor for the next O^L time units (i.e., until its starting offset). For any $s' > 0$, up to the value
 263 of O^L , this initial interference would be smaller (i.e., $O^L - s' < O^L$). The other value that we need
 264 to consider ($s' = O^L + X^L$) corresponds to the mode change occurring just as the server has ran out
 265 budget. Smaller values of s' (i.e., $O^L \leq s' < O^L + X^L$) would mean that the server is executing
 266 immediately after the mode change; greater values ($O^L + X^L < s' < S$) would only decrease the
 267 amount of time (i.e., $S - O^L - X^L + O^H$) until the server gets to execute for the first time after the
 268 mode change. Figure 3 illustrates these cases. Accordingly,

269
$$W_f^{H|tr.}(\Delta t) = \max_{s' \in [0, S)} W_f^{H|tr.}(s', \Delta t) = \max \left(W_f^{H|tr.}(0, \Delta t), W_f^{H|tr.}(O^L + X^L, \Delta t) \right) \quad (2)$$



■ **Figure 3** A mode change occurs s' time units after a timeslot boundary ($0 \leq s' < S$). However, the joint consideration of cases ($s' = 0$) and ($s' = O^L + X^L$) dominates all other $s' \in [0, S)$ under our analysis, because (i) shifting s' from $s' = 0$ to the right, by up to O^L time units can only decrease the post-mode-change interference from the fake task and, similarly, (ii) shifting s' from $s' = O^L + X^L$ to the left by up to X^L time units, or by any amount to the right, within the same timeslot, has the same effect. Plotted at the top (blue) and bottom (red) are the corresponding workload curves, $W_f^{H|tr.}(0, \Delta t)$ and $W_f^{H|tr.}(O^L + X^L, \Delta t)$ for the fake task after the mode change.

270 Equation (2) upper-bounds the fake task's "execution" (i.e., unavailability of the server) over any
 271 interval of length Δt starting at s , the mode change instant. By inspection of Figure 3 (blue plot):

$$272 \quad W_f^{H|tr.}(0, \Delta t) = \begin{cases} \Delta t, & \Delta t \leq O^L \\ O^L, & O^L < \Delta t \leq O^L + X^L \\ \Delta t - X^L, & O^L + X^L < \Delta t \leq S + O^H \\ S + O^H - X^L + \left\lfloor \frac{\Delta t - (S + O^H)}{S} \right\rfloor (S - X^H) \\ + \max\left(0, \Delta t - (S + O^H) - \left(\left\lfloor \frac{\Delta t - (S + O^H)}{S} \right\rfloor S\right) - X^H\right), & S + O^H < \Delta t \end{cases} \quad (3)$$

273 Likewise, by inspecting Figure 3 (red plot), we obtain for $W_f^{H|tr.}(O^L + X^L, \Delta t)$ the expression:

$$274 \quad W_f^{H|tr.}(O^L + X^L, \Delta t) = W_f^{H|tr.}(0, \Delta t + O^L + X^L) - O^L \quad (4)$$

275 Accordingly, the corresponding equivalent request-bound functions for any interval $[s, s + \Delta t)$ are:

$$276 \quad I_f^{H|tr.}(0, \Delta t) = O^L + \min\left(1, \left\lfloor \frac{\Delta t - O^L - X^L}{S} \right\rfloor_0\right) (S - X^L - O^L + O^H) \\ 277 \quad + \left\lfloor \frac{\Delta t - S - O^H - X^H}{S} \right\rfloor_0 (S - X^H) \quad (5)$$

$$\begin{aligned}
278 \quad I_f^{H|tr.}(O^L + X^L, \Delta t) &= (S - X^L - O^L + O^H) \\
279 \quad &+ \left[\frac{\Delta t - S + O^L - O^H + X^L - X^H}{S} \right]_0 (S - X^H) \quad (6)
\end{aligned}$$

280 Analogously as before,

$$281 \quad I_f^{H|tr.}(\Delta t) = \max \left(I_f^{H|tr.}(0, \Delta t), I_f^{H|tr.}(O^L + X^L, \Delta t) \right) \quad (7)$$

282 Next, we will incorporate $I_f^{H|tr.}(\Delta t)$ into a hybrid AMC-max schedulability test for this case.
283 Namely, we can upper-bound the response time of an H-task τ_i , released at or before the mode change
284 instant s but not yet completed by s , as

$$285 \quad R_i^{H|1} = \max(R_i^{s|1}), \forall s \in [0, R_i^L] \quad (8)$$

286 where

$$287 \quad R_i^{s|1} = C_i^H + \sum_{\tau_j \in hpL(i)} IL_j(s) + \sum_{\tau_k \in hpH(i)} IH_k(s, R_i^{s|1}) + \underbrace{IL_f(s) + I_f^{H|tr.}(R_i^{s|1} - s)}_{fake\ task} \quad (9)$$

288 where $hpL(i)$ and $hpH(i)$ are the sets of higher-priority low- and high-criticality tasks in the
289 same server, respectively, and

$$290 \quad IL_j(s) = \left(\left\lfloor \frac{s}{T_j} \right\rfloor + 1 \right) C_j^L \quad (10)$$

$$291 \quad IL_f(s) = \left(\left\lfloor \frac{s}{T_f} \right\rfloor + 1 \right) (S - X^L) \quad (11)$$

$$292 \quad IH_k(s, t) = M(k, s, t)C_k^H + \left(\left\lceil \frac{t}{T_k} \right\rceil - M(k, s, t) \right) C_k^L \quad (12)$$

293 where, classically from [5], $M(k, s, t) = \min \left(\left\lceil \frac{t-s-T_k-D_k}{T_k} \right\rceil + 1, \left\lceil \frac{t}{T_k} \right\rceil \right)$.

294 By replacing IH_k in (9) with the RHS of (12), we obtain:

$$\begin{aligned}
295 \quad R_i^{s|1} &= C_i^H + \sum_{\tau_j \in hpL(i)} IL_j(s) + \sum_{\tau_k \in hpH(i)} \left(M(k, s, R_i^{s|1})C_k^H + \left(\left\lceil \frac{t}{T_k} \right\rceil - M(k, s, R_i^{s|1}) \right) C_k^L \right) \\
&+ IL_f(s) + I_f^{H|tr.}(R_i^{s|1} - s)
\end{aligned}$$

296 Splitting the second summation and reordering the terms:

$$\begin{aligned}
297 \quad R_i^{s|1} &= C_i^H + \sum_{\tau_j \in hpL(i)} IL_j(s) + IL_f(s) \\
&+ \sum_{\tau_k \in hpH(i)} \left(\left\lceil \frac{t}{T_k} \right\rceil - M(k, s, R_i^{s|1}) \right) C_k^L + \sum_{\tau_k \in hpH(i)} M(k, s, R_i^{s|1})C_k^H + I_f^{H|tr.}(R_i^{s|1} - s)
\end{aligned}$$

298 Finally, in a slight accuracy optimisation of ours,

$$299 \quad R_i^{s|1} = C_i^H + \left\lceil IL(s, R_i^{s|1}) \right\rceil^s + IH(s, R_i^{s|1}) \quad (13)$$

300 where:

$$\begin{aligned}
IL(s, R_i^{s|1}) &= \sum_{\tau_j \in hpL(i)} IL_j(s) + IL_f(s) + \sum_{\tau_k \in hpH(i)} \left(\left\lceil \frac{t}{T_k} \right\rceil - M(k, s, R_i^{s|1}) \right) C_k^L \\
301 \quad IH(s, R_i^{s|1}) &= \sum_{\tau_k \in hpH(i)} M(k, s, R_i^{s|1}) C_k^H + I_f^{H|tr.}(R_i^{s|1} - s)
\end{aligned}$$

302 and the operator $\llbracket \cdot \rrbracket^{\max}$ is defined as $\llbracket x \rrbracket^{\max} \stackrel{\text{def}}{=} \begin{cases} x & \text{if } x \leq \max \\ \max & \text{if } x > \max \end{cases}$.

303 Under the original AMC-max, the operator $\llbracket \cdot \rrbracket^s$ is not used. Our slight improvement acknowledges
304 that the interference from all jobs completed before the mode change cannot exceed s .³

305 **Case 2:** In this case, the H-task τ_i under analysis is released at some instant a_i , after the mode
306 change instant s , but before $t_{hp(i)}^{idle}$.

307 Since τ_i is not released before the mode change, it does not suffer any *direct* interference from
308 jobs of L-tasks. However, in the general case, it may suffer indirect *push-through* interference by
309 such tasks that executed before its release. By this, we mean that any higher-priority H-task jobs
310 released before the mode change instant s and not completed before time a_i (the release of τ_i) may
311 have suffered interference from L-tasks of even higher priority (if any), before the mode change. This
312 would comensurately push their execution to the right, along the time axis. In any case, we do not
313 need to quantify the push-through interference from L-tasks in Case 2 (or even test schedulability in
314 Case 2 at all!), because, as we will prove below, if τ_i is proven to be schedulable in the L-mode and
315 in Case 1, it will always also be schedulable in Case 2.

316 **► Lemma 1.** *If an H-task τ_i is schedulable in L-mode and schedulable in H-mode under the*
317 *assumptions of Case 1, then it is also schedulable in H-mode under the assumptions of Case 2.*

318 **Proof.** Assume that a H-task τ_i is schedulable in L-mode and in H-mode under Case 1. Then,
319 consider some schedule σ wherein a_i is the first time instant strictly after the mode change s that τ_i is
320 released and it also holds that $a_i < t_{hp(i)}^{idle}$. This schedule then fulfills the assumptions of Case 2. Any
321 immediately preceding job by τ_i will have been released no later than s , so it would fall under Case 1
322 or will have completed before the mode change, so in either case, it would have been schedulable;
323 this means that the job by τ_i released at time a_i suffers no interference by previous jobs of the same
324 task; it only suffers interference from higher-priority tasks (including the fake task).

325 Let us then transform this original schedule σ to another schedule σ' where, all other things
326 remaining equal, the release of the job by τ_i under analysis is shifted earlier, to some time instant
327 t'' (with the releases of all other jobs by τ_i also shifted earlier by the same amount), such that the
328 following things hold:

- 329 ■ The instant t'' is located at or before the mode change (i.e., $t'' \leq s$).
- 330 ■ The entire interval $[t'', a_i]$ is occupied by execution of tasks in $hp(i)$ or the fake task.

331 The fact that in the original schedule, the entire interval $[s, a_i]$ was busy by higher-priority tasks
332 (including the fake tasks), means that such an instant t'' exists. It could be time instant s itself, or
333 some even earlier time instant.

334 Then, the response time of the job under analysis cannot decrease as a result of the schedule
335 transformation. Namely, in the transformed schedule σ' , the task τ_i does not execute at all over
336 $[t'', a_i]$ (where a_i refers to its release in the original schedule σ), and from time a_i onwards, its

³ Not enclosing the expression by the operator $\llbracket \cdot \rrbracket^s$, would allow the analysis to hold even for a variant model that permits any L-jobs caught up in the mode change to complete, executing for up to the respective L-WCETs. This follows from the original AMC-max – see [5].

337 execution intervals are the same as in the original schedule. Therefore its absolute completion time
 338 f_i is unchanged, even though its release is shifted earlier, to time $t'' < a_i$. In turn, the transformed
 339 schedule σ' belongs to Case 1, analysed earlier (i.e., τ_i being released no later than s but completing
 340 after s). Therefore, the increased response time of the job is upper-bounded by D_i , from the
 341 assumption that τ_i is schedulable in Case 1. Therefore the original response time of the job in
 342 schedule σ was also upper-bounded by D_i .

343 We will now show by contradiction that, if τ_i is schedulable in L-mode and in H-mode under
 344 Case 1, there cannot be more than one job of τ_i released strictly after s and strictly before $t_{hp(i)}^{idle}$.
 345 Assume that in schedule σ the next job by τ_i , after the one released at a_i , was released at time a'_i and
 346 that $a'_i <_{hp(i)}^{idle}$. Then, the job released at time a_i does not receive any execution time at all during the
 347 interval $[a_i, a'_i)$, therefore it misses its deadline at time $a_i + D_i \leq a'_i$. This contradicts the fact that it
 348 is schedulable, which we proved earlier. ◀

349 This means that the schedulability test for Case 2 is subsumed by the one for Case 1. In other
 350 words, if τ_i provably meets its deadline in L-mode and in H-mode under Case 1, then it also does so
 351 under Case 2. Accordingly, since we have to test for Case 1 anyway, it is redundant to test for Case 2.

352 **Case 3:** In this case, because the H-task τ_i under analysis is released at $t_{hp(i)}^{idle}$ or later, there can
 353 be no push-through interference from L-tasks. Therefore, there is only direct interference, from the
 354 tasks in $hpH(i)$ and from the unavailability of the server (i.e., from the fake task). The worst-case, in
 355 terms of interference from tasks in $hpH(i)$, is when these are released simultaneously as τ_i , at time
 356 a_i . This is the same as the worst-case interference from those tasks when we are in Case 1 and $s = 0$.

357 As for the worst-case interference from the fake task, given that in Case 3 the release time a_i of τ_i
 358 is before the server budget change instant, it is upper-bounded by (7), as in Case 1 (using the exact
 359 same reasoning).

360 Combining our observations, the schedulability test for Case 3 is also subsumed by the test for
 361 Case 1.

362 **Case 4:** Since τ_i is released at $t_{hp(i)}^{idle}$ or later and also at the server budget change or later, it is not
 363 subject to any transitive effects either from the mode change or from the budget change. Therefore, to
 364 compute the WCRT of the task, we can apply classic fixed-priority response time analysis, considering
 365 (i) the tasks present in the H-mode and their respective WCET estimates for that mode, and (ii) a fake
 366 top-priority task, modelling the unavailability of the server, with a WCRT of $S - X^H$ and a period of
 367 S . The corresponding equation is:

$$368 \quad R_i^{H|4} = C_i^H + \sum_{\tau_j \in hpH(i)} \left[\frac{R_j^{H|4}}{T_j} \right] C_j^H + \underbrace{\left[\frac{R_i^{H|4}}{S} \right] (S - X^H)}_{\text{fake task}} \quad (14)$$

369 Note that the worst-case processor request by the fake task in Case 4, during an interval of Δt
 370 time units, which is $\left[\frac{R_i^{H|4}}{S} \right] (S - X^H)$ is not necessarily upper-bounded, in the general case, by the
 371 expression $I_f^{H|tr.}(\Delta t)$ (Equation (7)) that describes the request by the fake task in Cases 1, 2 and 3
 372 (i.e., when τ_i is released before the budget change). Therefore, the schedulability test for Case 4 is
 373 not dominated by the schedulability test for Case 1, so we need to test for Case 4 separately.

374 4.2 Schedulability analysis at the system level

375 Having formulated how to test for the schedulability of a given server, we can now explain how the
 376 schedulability of the entire system is tested and how the assignments of execution budgets and starting
 377 offsets for the different servers in the two modes interdepend.

378 A system is schedulable if all servers are assigned non-overlapping execution windows inside the
 379 timeslot in both modes and if they are all found schedulable by the server schedulability test from
 380 Section 4.1 with the assigned execution budgets and starting offsets. In notation:

$$\begin{aligned}
 & (\forall i : \tilde{P}_i \text{ is schedulable with } (X_i^L, X_i^H, O_i^L, O_i^H)) \\
 \wedge & (\forall i : (O_i^L + X_i^L \leq S) \wedge (O_i^H + X_i^H \leq S)) \\
 \wedge & (\forall i, j, i \neq j : (O_i^L + X_i^L \leq O_j^L) \vee (O_j^L + X_j^L \leq O_i^L)) \\
 \wedge & (\forall i, j, i \neq j : (O_i^H + X_i^H \leq O_j^H) \vee (O_j^H + X_j^H \leq O_i^H))
 \end{aligned}$$

385 If the servers appear in the same order inside the timeslot in both modes (as we already assume)
 386 and their time windows are arranged back-to-back, with the first server aligned with the start of the
 387 timeslot, then (if, without loss of generality, the servers are indexed from left to right), we have

$$388 \quad O_1^L = 0; O_i^L = O_{i-1}^L + X_{i-1}^L, \forall i > 1 \quad (15)$$

$$389 \quad O_1^H = 0; O_i^H = O_{i-1}^H + X_{i-1}^H, \forall i > 1 \quad (16)$$

390 This allows the schedulability condition to be simplified to

$$391 \quad (\forall i : \tilde{P}_i \text{ is schedulable with } (X_i^L, X_i^H, O_i^L, O_i^H)) \wedge \left(\sum X_i^L \leq S \right) \wedge \left(\sum X_i^H \leq S \right)$$


392 4.2.1 Interdependencies between the parameters of different servers

393 From (15) and (16), one can see that for a given ordering (indexing) of the servers, the execution
 394 budgets of preceding servers in one mode determine the starting offset of a given server in that
 395 mode. In turn, these offsets (O_i^L and O_i^H) are inputs (along with the budgets X_i^L and X_i^H) to
 396 the schedulability test for that server \tilde{P}_i . Therefore the execution budgets of all preceding servers
 397 indirectly affect whether or not a server \tilde{P}_i is schedulable with a given budget pair (X_i^L, X_i^H). The
 398 ordering of the servers within the timeslot thus matters a lot for the system schedulability.

399 Intuitively, one would expect that ordering the servers such that they appear in the timeslot by
 400 non-decreasing $X_i^H - X_i^L$ would be a helpful heuristic for achieving good scheduling performance.
 401 The reasoning is that if the servers appear in order of $X_i^H - X_i^L$ in the timeslot, then $O_i^L - O_i^H \geq 0$
 402 for all servers – and, in the timeslot where a mode change occurs, this “benign” jitter would mean
 403 that the interval between a server completing with a budget of X_i^L and the start of the execution of
 404 the next server instance, with budget X_i^H , would be S or (in most cases) smaller than S . This implies
 405 a shorter effective transition time to the new budgets (greater responsiveness to the new processing
 406 requirements), compared to the case of $O_i^L - O_i^H > 0$. However, in the general case, we cannot know
 407 a priori in which order to arrange the servers such that they appear in order of $X_i^H - X_i^L$, because the
 408 budgets can only be computed a posteriori, using the offsets O_i^L and O_i^H as input, which themselves
 409 depend on the server ordering, as we just explained earlier.

410 Additionally, in the general case, and for a given pair of starting offsets (O_i^L, O_i^H) there may
 411 exist multiple budget pairs (X_i^L, X_i^H) for which a server is schedulable. If sensitivity analysis (e.g.,
 412 binary search) is used to determine the least feasible budget X_i^H for a given offset pair (O_i^L, O_i^H) as
 413 a function of X_i^L , then the pair (X_i^L, X_i^H) will exhibit the Pareto property. Namely, a more generous
 414 L-mode budget X_i^L might require a smaller X_i^H budget for the L-mode (as, intuitively, the server
 415 will have comparatively less “catching up” to do with the tasks’ demand) – and vice versa.

416 All these different interdependencies between the parameters of different servers and their ordering,
 417 complicate the task of devising good heuristics for ordering the servers inside the timeslot and



■ **Figure 4** The two server schedules depicted are identical (since each server occupies exactly the same time windows in both), the only difference being which time instant is considered as the start of the periodic timeslot of length S . Therefore, one system being (un)schedulable implies that the other one is too.

418 assigning budgets to them for the two modes. However, as we will show in the next section, it is still
 419 possible to leverage them in a useful way, and attain good performance.

420 5 Server budget assignment heuristics

421 We consider two different scheduling arrangements (static-server budgets and dynamic server budgets)
 422 and different heuristics for assigning the server budgets in each case.

423 5.1 Static server budgets (SSB)

424 Under the static server budget (SSB) arrangement, the execution budget of a server remains the same
 425 in both modes (i.e., $X_i^L = X_i^H = X_i \forall i$). Additionally, the first server is positioned at the beginning of
 426 timeslot ($O_1^L = O_1^H = 0$) and every subsequent server starts when its predecessor ends. These properties
 427 imply that $O_i^L = O_i^H \forall i$. In other words, neither the starting offset nor the execution budget of any
 428 server ever changes. Consequently, there is no need to apply the analysis formulated in Section 4.
 429 Rather, we can apply the original AMC-max schedulability test [5], for sizing each server, with the
 430 addition of a top-priority fake H-task τ_f with attributes $C_f^L = C_f^H = S - X_i$ and $T_f = S$. The minimum
 431 feasible server budget X_i for a given server can be identified via binary search [3, 34] over the interval
 432 $(0, S]$. Note that for this arrangement, each server can be sized independently of other servers and
 433 their attributes. Moreover, the order in which the servers are arranged on a processor is irrelevant.

434 The previously mentioned independence between servers also holds when our analysis is used
 435 instead of the original AMC-max analysis to test the feasibility while sizing servers with static budgets.
 436 Indeed, the offsets of the server under analysis (and all other servers) can be disregarded because
 437 any static-budget server execution pattern can be transformed via shift-rotation along the temporal
 438 axis into an equivalent schedule where the server under consideration has a given offset (which can
 439 conveniently be $O_i^L = O_i^H = 0$), as Figure 4 illustrates. Crucially, the unavailability intervals for the
 440 server all have a duration of $S - X_i$ and occur strictly periodically with a period of S .

441 Nevertheless, even if our new analysis can accommodate static server budgets as a special case,
 442 it does not necessarily perform better than the original AMC-max for this arrangement because of
 443 the slight pessimism introduced by upper-bounding the interferences from the unavailability of the
 444 server in L-mode and in H-mode separately from each other (see Equation 9). In any case, for greater
 445 insight, in Section 6, we plot results for SSB using either the original AMC-max or the new analysis.

446 5.2 Dynamic server budgets (DSB)

447 Under this arrangement, for which our analysis was developed, *a server can have a different budget*
 448 *after a mode switch*. As in the SSB arrangement, we assume that the servers execute back-to-back
 449 and the first server is aligned with the start of the timeslot. Both the L-mode and H-mode budgets of
 450 a server are computed offline. To compute the minimum feasible budget X_i^H of a server in H-mode,

451 we need to know its starting offsets in each mode (O_i^L and O_i^H) and its L-mode budget X_i^L . In turn,
 452 the offset of a given server in a given mode can only be computed if we already know the budgets of
 453 all predecessor servers. This implies that the order in which the servers are arranged in the schedule
 454 is already decided. In Section 5.2.1, we discuss heuristics for selecting that ordering.

455 For the DSB arrangement, we explore two different heuristics. First, we consider a simple
 456 heuristic that assigns L-mode budgets (X_i^L) to all servers, proportionally to the minimum feasible
 457 L-mode budget X_i^{min} for each server (identified with a binary search algorithm [34]). All server
 458 offsets and H-mode budgets are eventually computed from that set of L-mode budgets, directly or
 459 indirectly. As a second option, we explore a metaheuristic (Simulated Annealing [28]), which accepts
 460 the output of the previous heuristic as a starting solution (if not already feasible), and tries to mutate
 461 the original set of X_i^L budgets until it becomes feasible.

462 Simulated annealing attempts to replace the current solution of a problem with another (randomly
 463 obtained) solution in each iteration. A candidate solution that improves on the current one is always
 464 accepted. However, occasionally, the algorithm will also accept a “worse” candidate solution with
 465 a probability that depends on the value of a probability function. This function takes as parameters
 466 a variable Θ (dubbed “the temperature”) and the difference of the “utilities” of the current solution
 467 and the candidate solution. Higher temperatures and lower reduction in utility raise the acceptance
 468 probability for a “worse” solution. Occasionally accepting “worse” solutions helps avoid the pitfall
 469 of getting stuck at a local optimum of the optimization problem. The temperature Θ is gradually
 470 decreased with the number of iterations. In our particular problem, a solution is represented by the
 471 set of X_i^L values (which uniquely determines all eventual O_i^L and O_i^H offsets and X_i^H budgets). As
 472 utility of a given solution, we define the sum of the X_i^H budgets calculated separately for each server,
 473 assuming the corresponding X_i^L value and $O_i^L = O_i^H = 0$.

474 In more detail, the pseudocode for both heuristics is presented in Algorithm 1.

475 **Initial Phase (Simple heuristic):** Initially, we determine the minimum feasible L-mode budget
 476 X_i^{min} for each server. To do that, for each server separately, we assume that its H-mode budget is equal
 477 to S (the entire timeslot) and its starting offsets are equal to zero (i.e., $X_i^H = S$ and $O_i^L = O_i^H = 0$).
 478 With these assumptions, we compute, using our new analysis as feasibility test, the corresponding
 479 minimum feasible L-mode size X_i^{min} for each server with a binary search algorithm [34]. If the
 480 sum of X_i^{min} for all servers is greater than S or if any of the servers are infeasible with the maximal
 481 H-mode server size of S , we declare failure as the system is provably unschedulable, with any
 482 assignment of server budgets. Otherwise, once X_i^{min} has been computed for all servers, we set
 483 $X_i^L = X_i^{min} * \frac{S}{\sum_{\forall i} X_i^{min}}$ for each server. The factor $\frac{S}{\sum_{\forall i} X_i^{min}}$ proportionally scales up the X_i^{min}
 484 value of each server to fill up entirely the L-budget timeslot S . So by construction, $\sum_{\forall i} X_i^L \leq S$.

485 The initial X_i^L server budgets in L-mode are in turn used to compute the actual H-mode server
 486 budgets. With a server order given a priori, the H-mode budget (X_i^H) of any i^{th} server can be
 487 computed with a binary search algorithm assuming offsets of $O_i^L = \sum_{j=1}^{i-1} X_j^L$ and $O_i^H = \sum_{j=1}^{i-1} X_j^H$.
 488 If for the computed X_i^H values it holds that $\sum_{\forall i} X_i^H \leq S$, we declare a success. Otherwise, we try
 489 the metaheuristic (Simulated Annealing), implemented in the main loop:

490 **Main Loop (Simulated Annealing):** In any iteration k , two servers (\tilde{P}_a and \tilde{P}_b) are selected
 491 randomly. This heuristic increments X_a^L of \tilde{P}_a and decrements X_b^L of \tilde{P}_b by a same value of β , where
 492 β represents the server variation length parameter for this iteration. Adding and subtracting the same
 493 value of β from two selected servers keeps the sum of L-mode server budgets in the k^{th} iteration
 494 equal to that of the $(k - 1)^{th}$ iteration, i.e., $\sum_{\forall i} X_i^L(k) = \sum_{\forall i} X_i^L(k - 1)$.

495 By construction, the heuristic keeps $\sum_{\forall i} X_i^L \leq S$. Hence, the parameter β is computed through
 496 Algorithm 2 in such a way that this condition is never violated. In Algorithm 2, initially, $\beta_{max} =$
 497 $-(\Delta * S) + (2 * \gamma)$ gives the maximum length to vary in this iteration, where $\Delta \in (0, 1]$ is an input
 498 parameter, and γ is a randomly generated variable in $(0, \Delta * S]$. Afterwards, β_a and β_b are selected

Algorithm 1 Simple heuristic and Simulated Annealing algorithm

Input: Sorted \tilde{P} , S , Δ , Θ and Cooling Rate

Output: $X_i^L, X_i^H, \forall i$

1: **Initial Phase:**

2: Generate X_i^{min} for each server \tilde{P}_i using a binary search algorithm [34] assuming $X_i^H = S$.

3: **if** ($\sum_{\forall i} X_i^{min} > S$ || any \tilde{P}_i infeasible with $X_i^H = S$) **then return** Failure

4: **else**

▷ scale up X_i^L values

5: Set $X_i^L = X_i^{min} \times \frac{S}{\sum_{\forall i} X_i^{min}}, \forall i$

6: Compute X_i^H for each server, given X_i^L values $\forall i$, with offsets

7: **if** ($\sum_{\forall i} X_i^H \leq S$) **then return** $X_i^L, X_i^H, \forall i$

8: **Main Loop:**

9: **while** ($\Theta > 1$) **do**

10: Select two random servers \tilde{P}_a and \tilde{P}_b

11: Compute β for servers \tilde{P}_a and \tilde{P}_b with Algorithm 2

12: Set $X_a^L(k) = X_a^L(k-1) + \beta$ and $X_b^L(k) = X_b^L(k-1) - \beta$

13: Set other servers $X_i^L(k) = X_i^L(k-1), \forall i \notin a, b$

14: Compute $X_i^H(k), \forall i$ in k^{th} iteration with offsets

15: **if** ($\sum_{\forall i} X_i^H \leq S$) **then return** $X_i^L(k), X_i^H(k), \forall i$

16: **else**

17: Compute $X_i^H(0, k), \forall i$ with $O_i^L = O_i^H = 0$ offset in iteration k

18: **if** ($\sum_{\forall i} X_i^H(0, k) < \sum_{\forall i} X_i^H(0, k-1)$) **then**

19: Keep $X_i^L(k), X_i^H(0, k), \forall i$ for $k+1^{th}$ iteration

20: **else if** ($z \in (0, 1] \leq e^{-\frac{\sum_{\forall i} X_i^H(0, k-1) - \sum_{\forall i} X_i^H(0, k)}{\Theta}}$) **then**

21: Keep $X_i^L(k), X_i^H(0, k), \forall i$ for $k+1^{th}$ iteration

22: **else**

23: Discard this and keep $(k-1)^{th}$ iteration solution

24: $\Theta = \Theta * (1 - \text{Cooling Rate})$

25: $k = k + 1$

26: On while loop termination without success, **return** Failure

499 such that $X_a^L + \beta_a$ remains in $[X_a^{min}, S]$ and $X_b^L - \beta_b$ remains in $[X_b^{min}, S]$. Between β_a and β_b ,
 500 the one that gives the least change in server size is selected as β .

501 After selecting β , $X_a^L(k)$ and $X_b^L(k)$ are updated to $X_a^L(k-1) + \beta$ and $X_b^L(k-1) - \beta$, respectively.

502 All other servers get the previous-iteration values, i.e., $X_i^L(k) = X_i^L(k-1), \forall i \neq a, i \neq b$.

503 Once L-server budgets are available for the k^{th} iteration, the corresponding H-mode server sizes

504 are computed (employing binary search and our analysis as schedulability test), using the offsets

505 $O_i^L(k) = \sum_{j=1}^{i-1} X_j^L(k)$ and $O_i^H(k) = \sum_{j=1}^{i-1} X_j^H(k)$, for any server \tilde{P}_i .

506 If the process of computing $X_i^H(k)$ with the above offsets for the k^{th} iteration is successful and

507 $\sum_{\forall i} X_i^H(k) \leq S$, we declare a success and exit the loop. Otherwise, the ‘‘utility’’ of the current

508 solution is computed. For that purpose, we calculate for each server what its least feasible H-mode

509 budget would be ($X_i^H(0, k)$), with the current $X_i^L(k)$ budget and assuming $O_i^L = O_i^H = 0$. The

510 utility of the solution is the sum of those X_i^H values. If $\sum_{\forall i} X_i^H(0, k) < \sum_{\forall i} X_i^H(0, k-1)$ (i.e., if it

511 is a ‘‘better’’ solution than the previous iteration), or if a randomly generated number z in $(0, 1]$ is less

512 than or equal to acceptance probability of $e^{-\frac{\sum_{\forall i} X_i^H(0, k-1) - \sum_{\forall i} X_i^H(0, k)}{\Theta}}$ (i.e., occasionally accepting the

513 worse solution), then the $X_i^L(k)$ and $X_i^H(0, k), \forall i$ are accepted for the $(k+1)^{th}$ iteration. Otherwise,

514 these values are discarded, and $X_i^L(k-1)$ and $X_i^H(0, k-1), \forall i$ are used for the $(k+1)^{th}$ iteration.

Algorithm 2 Server variation length parameter (β) computation algorithm

Input: $\tilde{P}_a, \tilde{P}_b, S$ and Δ **Output:** Server size variation factor β

```
1:  $\beta_{max} = -(\Delta * S) + (2 * \gamma)$   $\triangleright$  Parameter  $\Delta \in (0, 1]$ ;  $\gamma$  uniformly distributed over  $(0, \Delta * S]$ .
2: if ( $X_a^L + \beta_{max} > S$ ) then
3:    $\beta_a = S - X_a^L$ 
4: else if ( $X_a^L + \beta_{max} < X_a^{min}$ ) then
5:    $\beta_a = X_a^{min} - X_a^L$ 
6: else
7:    $\beta_a = \beta_{max}$ 
8: if ( $X_b^L - \beta_{max} > S$ ) then
9:    $\beta_b = X_b^L - S$ 
10: else if ( $X_b^L - \beta_{max} < X_b^{min}$ ) then
11:    $\beta_b = X_b^L - X_b^{min}$ 
12: else
13:    $\beta_b = \beta_{max}$ 
14: if ( $\beta_{max} \geq 0$ ) then
15:    $\beta = \min\{\beta_a, \beta_b\}$ 
16: else
17:    $\beta = \max\{\beta_a, \beta_b\}$ 
return  $\beta$ 
```

515 Finally, if the system cools down without finding any feasible solution, a failure is declared.

516 5.2.1 Server ordering heuristics

517 To achieve efficient dynamic server budget assignment, the order of servers is an important initial
518 step. We propose to sort the servers in non-decreasing order of $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$, where $U^H(\tilde{P}_i)$
519 and $U^L(\tilde{P}_i)$ represent the H and L-mode utilisation of server \tilde{P}_i , respectively. The reason we chose
520 this ordering was because it would result in a server ordering that would approximate an ordering
521 by non-decreasing $X_i^H - X_i^L$. Recall that in Section 4.2.1, we identified that ordering servers by
522 non-decreasing $X_i^H - X_i^L$ would likely promote good performance. However one can only confirm
523 whether a given server ordering meets this property *a posteriori*, due to the interdependencies of the
524 servers' parameters with each other and the ordering. Therefore, we simply chose $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$
525 (which is verifiable *a priori*) as a good proxy. In our set of experiments, we also experimented with
526 server ordering by non-increasing $\frac{U^H(\tilde{P}_i)}{U^L(\tilde{P}_i)}$, but it performed slightly worse.

527 6 Evaluation

528 6.1 Experimental setup

529 We have developed a Java tool to implement the proposed analysis and explore the scheduling per-
530 formance of different scheduling arrangements and budget assignment heuristics and the effectiveness
531 of the schedulability analysis. Our tool has two modules. The first module generates the synthetic
532 workload (task sets and servers) for the given platform parameters. A second module performs the
533 feasibility analysis with the proposed techniques.

534 **Task set generation:** Task periods are generated with a log-uniform distribution in the 10-100

■ **Table 1** Overview of Parameters

Parameters	Values	Default
Number of Servers (q)	{2, 3, 4, 5, 6}	3
Task-set size (n)	{9, 12, 15, 18, 21, 24}	3/server
H-tasks share	{20 : 10 : 80}%	30%
HWCET scaling factor (κ)	{2 : 1 : 6}	3
Temperature (Θ)	{.025, .05, .1, .5, 1, 5, 10}*1000	10000
Cooling rate	{.001, .005, .01, .05, .1, .3, .5, .7, .9, 1}	0.005
Server size variation (β)	{.001, .005, .01, .02, .05, .1, .2, .3, .4, .5, .6, .7, .8, .9}	0.5
Server ordering	$\{U^H(\tilde{P}_i) - U^L(\tilde{P}_i), \frac{U^H(\tilde{P}_i)}{U^L(\tilde{P}_i)}\}$	1 st
Inter-arrival time (T_i)	10ms to 100ms	N/A
Nominal utilisation	{0.1 : 0.1 : 1}	N/A

535 msec range. We generate implicit-deadline tasks ($D_i = T_i$), even though the analysis holds for
536 the more general constrained deadline model ($D_i \leq T_i$). The given target L-mode utilisation is
537 distributed among tasks by the UUnifast algorithm [8, 14] in an unbiased way. Then, a task's L-WCET
538 C_i^L is computed to $T_i * U_i^L$, where U_i^L is the task's L-mode utilisation. The fraction of H-tasks is a
539 user-defined parameter. The H-WCET of a task is a linearly scaled up value of its L-WCET, according
540 to an input parameter κ (i.e., $C_i^H = \kappa C_i^L$). Tasks are assigned priorities based on their arrival rates
541 (Deadline Monotonic). The generated task set is indexed in order of increasing deadlines and tasks
542 are assigned to servers using round-robin. To better utilise the system's resources, one can explore
543 the problem of efficiently assigning tasks-to-servers. However, we assume that the designer may not
544 have control over this, since the grouping of tasks is functional, on the basis of application. Finally,
545 for each system, the timeslot length (S), corresponding to the common period of all servers, is set to
546 the shortest task interarrival time in the task set.

547 The target L-mode system utilisation is varied within a range of (0, 1]. Different random class
548 objects are used to generate period and utilisation values. Each random object is seeded with a
549 different odd number and reused in successive replications [25]. For each set of input parameters, we
550 generate 1000 random task sets. The default values of the aforementioned parameters are: $q = 3$,
551 $n = q * 3$, 30% H-tasks and $\kappa = 3$. The range of values considered for all the parameters is presented
552 in Table 1. By default the servers are sorted in order of non-decreasing $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$. The
553 triple given in this table corresponds to {minimum : increment granularity : maximum} values of
554 a parameter. In our experiments, we vary one parameter from Table 1 at a time, while the others
555 conform to their default values.

556 Instead of providing plots comparing the approaches in terms of scheduling success ratio (i.e., the
557 fraction of task-sets deemed schedulable under the respective schedulability test), we condense this
558 information by providing plots of *weighted schedulability*. This performance metric [7, 10] condenses
559 what would have been three-dimensional plots into two dimensions. It is a weighted average that gives
560 more weight to task-sets with higher utilisation, which are supposedly harder to schedule. Specifically,
561 using notation from [10], let $S_y(\tau, p)$ represent the result (0 or 1) of the schedulability test y for a
562 given task-set τ with an input parameter p . Then $W_y(p)$, the weighted schedulability for that test
563 y as a function p , is $W_y(p) = \sum_{\forall \tau} (U(\tau) * S_y(\tau, p)) / \sum_{\forall \tau} U(\tau)$, where $U(\tau)$ is the utilisation of
564 task set τ . In this paper, the weighing is according to the L-mode schedulability. Nevertheless, the
565 schedulability success ratio plots are presented in Appendix A.


566 6.2 Scheduling arrangements and server ordering heuristics

567 The following scheduling arrangements and heuristics are compared in this section.


- 568 ■ **Static Server Budgets (SSB)**: This scheduling arrangement (presented in Section 5) maintains
569 the same size of a server in both modes. We use the original AMC-max analysis [5] for the
570 schedulability test, for fair comparison. A server's budget is set to the minimal value that ensures
571 schedulability, according to sensitivity analysis (binary search).
- 572 ■ **Static Server Budgets with our analysis (SSBO)**: This is the same as SSB, except that in place
573 of the original AMC-max test, our new analysis is used with server offsets set to zero, i.e.,
574 $O_i^L = O_i^H = 0 \forall i$. Comparing SSBO with SSB assess the pessimism, compared to AMC-max,
575 when the new analysis deals with the special case of $X_i^L = X_i^H \forall i$. Any such pessimism can be
576 attributed to the independent bounding of the fake task's interference in the two modes, in the
577 design of our analysis, to be able to analyse the general case of dynamic budgets.
- 578 ■ **Dynamic Server Budgets with Simple heuristic (SH)**: This approach for dynamic server
579 budgets declares a success, if the system is feasible using the simple heuristic for assigning
580 budgets that we described in Section 5, i.e., the "Initial Phase" from Algorithm 1. Our new
581 analysis is used for schedulability testing.
- 582 ■ **Dynamic Budgets with Simulated Annealing (SA)**: The simulated annealing heuristic is also
583 presented in Section 5, in the latter part of Algorithm 1. Once again, our new analysis is used for
584 schedulability testing. The corresponding curve plots success if a task set is either schedulable
585 using just SH, or, if it is not schedulable by SH, but the metaheuristic in the second stage is able to
586 find a feasible assignment of server budgets for the two modes. The variation in three parameters
587 (Temperature, Cooling Rate and Δ) used to configure the simulated annealing metaheuristic is
588 presented in Table 1.
- 589 ■ **Simulated Annealing with all possible server orderings (SAAO)**: Instead of picking a pre-
590 defined server ordering, all possible orderings in which the servers can be arranged are tested
591 for a feasible solution. For each of them, the SA heuristic (Dynamic Budgets with Simulated
592 Annealing) explained earlier is applied, with the specified maximum number of iterations, until
593 success. The SAAO heuristic allows us: (a) to analyze the effect of the order in which the servers
594 are arranged, and (b) to quantify the quality of the default ordering of $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$ and,
595 indirectly, the validity of the intuition behind its selection.
- 596 ■ **Simple heuristic assuming all possible server orderings (SHAO)**: It is similar to SH except
597 that all possible ordering of servers on the given processor are checked instead of the default
598 order.

599 6.3 Results


600 Figures 5 and 6 present the weighted schedulability for different number of servers. The number of
601 tasks per server is constant (3/Server) in Figure 5, while the total number of tasks to distribute among
602 servers via a round robin policy is constant (task set size = 18) in Figure 6. Increasing the server
603 count results in lower schedulability, due to the reduced average per-server budget in a given time
604 slot. All curves follow this trend. The difference between SAAO and SA is small which indicates that
605 the selected server ordering by non-decreasing ($U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$) is a reasonable choice. At low
606 server counts, SH performs similar to SA as the scaling of L-mode budgets to fully utilise the timeslot
607 of length S helps finding more feasible H-mode budgets, already in the initial phase. An increase in
608 the number of servers makes it harder to fit the servers in the timeslot and hence, the main loop in SA
609 (Algorithm 1) becomes useful. Similar behaviour can be seen when comparing SHAO and SAAO. In
610 the majority of the cases, SHAO performs better than SA, which indicates that the order of servers




■ **Figure 5** (Fixed number of tasks per server)



■ **Figure 6** (Fixed number of total tasks)



■ **Figure 7**



■ **Figure 8**

611 has high impact on the feasibility of the solution. SSB and SSBO behave almost the same. The slight
 612 improvement of SSBO over SSB, where present, is attributed to the optimisation in (13), where all
 613 L-mode interference is upper-bounded by s . This seems to mask the pessimism from the independent
 614 bounding of the fake task's interference in the two modes. SSB and SSBO perform similarly when
 615 the servers are fewer. However, their slight performance difference increases with more servers in the
 616 system. On the other hand, the difference between SA and SSB increases with more servers. This
 617 indicates that the proposed analysis performs better with its intended dynamic setting. In general,
 618 SAAO and SA display superior performance to SSB and other variants.


619 The number of tasks per server is larger in Figure 6 when compared to Figure 5. For a given
 620 system utilisation, many light tasks per server are, more often than not, easier to schedule compared
 621 to fewer, heavy tasks per server. Hence, the weighted schedulability is slightly higher in Figure 6
 622 against Figure 5. This observation is consistent with the result shown in Figure 7, as the weighted
 623 schedulability improves with a larger task set size. The number of servers is constant in Figure 7
 624 and the increase in task set size only increases the number of tasks per server, and consequently, the
 625 improvement in weighted schedulability.

626 A higher number of H-tasks (Figure 8) in a task set and a higher H-mode utilisation scaling
 627 factor κ (Figure 9) both increase the H-mode utilisation, making the task sets harder to schedule.
 628 Hence, the weighted schedulability decreases with an increase in these parameters for all heuristics.
 629 A higher temperature and a lower cooling rate increases the number of solution mutations during the
 630 iterations of the main loop in SA, and appear to improve the weighted schedulability, as shown in
 631 Figures 10 and 11, respectively. A larger value of Δ provides a bigger area of the design space to
 632 explore in each iteration of the simulated annealing metaheuristic, and, from the results, this improves
 633 the weighted schedulability, as shown in Figure 12. The order of servers is important for both SH
 634 and SA. The servers sorted in non-increasing order of $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$ show up to 0.5% and 0.8%
 635 better schedulability success ratio than non-increasing order of $\frac{U^H(\tilde{P}_i)}{U^L(\tilde{P}_i)}$, for SA and SH, respectively.


636 When compared to the baseline SSB, the absolute difference in terms of *schedulability success*

■ **Table 2** Maximum absolute difference in schedulability ratio of all heuristics from baseline SSB.


SA	SH	SAAO	SSBO	SHAO
39.6%	27.9%	52.8%	5.6%	40.4%




■ **Figure 9**




■ **Figure 10**



■ **Figure 11**



■ **Figure 12**



■ **Figure 13**

637 *ratio* of all heuristics is presented in Table 2. The change in server's budget can achieve up to 52.8%
 638 more in schedulability success ratio. For more insights, please see the schedulability success ratio
 639 plots in Appendix A.

640 **7** Conclusions

641 In this paper, we proposed a new schedulability analysis for mixed-criticality uniprocessor systems
 642 executing periodic servers (varying budgets in response to mode switch) in a cyclic executive manner
 643 and use the AMC-max scheduling policy to schedule the tasks within each server. Our proposed
 644 approach provides strict temporal isolation among applications with additional ability to efficiently
 645 utilise the available execution capacity across mode switches. We also proposed number of heuristics

646 that assigns budgets to servers in both modes and define the order of the servers in the cyclic executive
647 schedule. The experimental evaluation with synthetic task-sets showed by varying budgets in response
648 to mode switch improves schedulability ratio by up to 52.8%, compared to the baseline static server
649 budget algorithm. Even with a simple heuristic, we can achieve up to 27% of improvements in
650 schedulability ratio. The order of the servers in the cyclic executive schedule has high impact on the
651 schedulability ratio and the proposed heuristic to select the ordering of servers performs well in our
652 experiments. In the future, we intend to extend this approach to multicore platforms and include the
653 effect of other shared resources in the schedulability analysis.

654 ——— References ———

- 655 1 Certification authorities software team (cast), position paper (cast-32a) multicore processors. Certification
656 authorities in North and South America, Europe, and Asia, November 2016.
- 657 2 AERONAUTICAL RADIO, INC. Avionics Application Software Standard Interface, Part 1, Required
658 Services, ARINC SPECIFICATION 653P1-3 edition, November 2010.
- 659 3 Muhammad Ali Awan, Konstantinos Bletsas, Pedro F. Souto, and Eduardo Tovar. Semi-partitioned mixed-
660 criticality scheduling. In Proceedings of the 30th International Conference Architecture of Computing
661 Systems, pages 205–218, 2017.
- 662 4 Muhammad Ali Awan and Stefan M. Petters. Intra-task device scheduling for real-time embedded systems.
663 In Journal of Systems Architecture, 2013.
- 664 5 S. K. Baruah, A. Burns, and R. I. Davis. Response-time analysis for mixed criticality systems. In
665 Proceedings of the 32nd IEEE Real-Time Systems Symposium, pages 34–43, 2011.
- 666 6 Sanjoy Baruah and Alan Burns. Implementing mixed criticality systems in Ada. In 16th Ada-Europe
667 Conference, pages 174–188, 2011.
- 668 7 Andrea Bastoni, Björn Brandenburg, and James Anderson. Cache-related preemption and migration
669 delays: Empirical approximation and impact on schedulability. Proceedings of OSPERT, pages 33–44,
670 2010.
- 671 8 E. Bini and G.C. Buttazzo. Measuring the performance of schedulability tests. Journal of Real-Time
672 Systems, 30(1-2):129–154, 2009.
- 673 9 Konstantinos Bletsas, Muhammad Ali Awan, Pedro F. Souto, Benny Akesson, Alan Burns, and Eduardo
674 Tovar. Decoupling criticality and importance in mixed-criticality scheduling. In Proceedings of the 6th
675 Workshop on Mixed-Criticality Systems, 2018.
- 676 10 A. Burns and R.I. Davis. Adaptive mixed criticality scheduling with deferred preemption. In Proceedings
677 of the 35rd IEEE Real-Time Systems Symposium, pages 21–30, Dec 2014.
- 678 11 Alan Burns and Robert Davis. Mixed criticality systems – a review (11th edition). Technical report,
679 Department of Computer Science, University of York, UK, Aug. 2018.
- 680 12 Alan Burns and Robert I. Davis. A survey of research into mixed criticality systems. ACM Computing
681 Surveys, 50(6):82:1–82:37, November 2017.
- 682 13 Micaiah Chisholm, Namhoon Kim, Stephen Tang, Nathan Otterness, James H. Anderson, F. Donelson
683 Smith, and Donald E. Porter. Supporting mode changes while providing hardware isolation in mixed-
684 criticality multicore systems. In Proceedings of the 25th Conference Real-Time and Networked Systems,
685 RTNS '17, pages 58–67, 2017.
- 686 14 Robert I. Davis and Alan Burns. Priority assignment for global fixed priority pre-emptive scheduling
687 in multiprocessor real-time systems. In Proceedings of the 30th IEEE Real-Time Systems Symposium,
688 pages 398–409, 2009.
- 689 15 Pontus Ekberg and Wang Yi. Bounding and shaping the demand of generalized mixed-criticality sporadic
690 task systems. Journal of Real-Time Systems, 50(1):48–86, 2014.
- 691 16 C. Evripidou and A. Burns. Scheduling for mixed-criticality hypervisor systems in the automotive domain.
692 In Proceedings of the 4th Workshop on Mixed-Criticality Systems, 2016.
- 693 17 Lipari Giuseppe and C. Buttazzo Giorgio. Resource reservation for mixed criticality systems. In Workshop
694 on Real-Time Systems: The past, The present, and the future, March 2013.

- 695 18 X. Gu and A. Easwaran. Dynamic budget management with service guarantees for mixed-criticality
696 systems. In Proceedings of the 37rd IEEE Real-Time Systems Symposium, pages 47–56, 2016.
- 697 19 Xiaozhe Gu, Arvind Easwaran, Kieu-My Phan, and Insik Shi. Resource efficient isolation mechanisms in
698 mixed-criticality scheduling. In Proceedings of the 27th Euromicro Conference on Real-Time Systems,
699 pages 13–24, 2015.
- 700 20 Fei Guan, Long Peng, Luc Perneel, Hasan Fayyad-Kazan, and Martin Timmerman. Adaptive reservation
701 into mixed-criticality systems. Scientific Programming, 2017, 2017.
- 702 21 J. L. Herman, C. J. Kenna, M. S. Mollison, J. H. Anderson, and D. M. Johnson. Rtos support for multicore
703 mixed-criticality systems. In Proceedings of the 18th IEEE Real-Time and Embedded Technology and
704 Applications Symposium, pages 197–208, April 2012.
- 705 22 Biao Hu, Kai Huang, Gang Chen, Long Cheng, and Alois Knoll. Adaptive workload management in
706 mixed-criticality systems. ACM Transactions on Embedded Computing Systems, 16, 10 2016. doi :
707 10.1145/2950058.
- 708 23 Biao Hu, Lothar Thiele, Pengcheng Huang, Kai Huang, Christoph Griesbeck, and Alois Knoll. FFOB:
709 Efficient online mode-switch procrastination in mixed-criticality systems. Journal of Real-Time Systems,
710 2018. doi:10.1007/s11241-018-9323-x.
- 711 24 Huang-Ming Huang, Christopher Gill, and Chenyang Lu. Implementation and evaluation of mixed-
712 criticality scheduling approaches for periodic tasks. In Proceedings of the 18th IEEE Real-Time and
713 Embedded Technology and Applications Symposium, pages 23–32, 2012.
- 714 25 Raj Jain. The art of computer systems performance analysis - techniques for experimental design,
715 measurement, simulation, and modeling. Wiley professional computing. Wiley, 1991.
- 716 26 M. Joseph and P. Pandya. Finding Response Times in a Real-Time System. The Computer Journal,
717 29(5):390–395, 1986.
- 718 27 Namhoon Kim, Bryan C. Ward, Micaiah Chisholm, Cheng-Yang Fu, James H. Anderson, and F. Donelson
719 Smith. Attacking the one-out-of-m multicore problem by combining hardware management with mixed-
720 criticality provisioning. Journal of Real-Time Systems, 53(5):709–759, 2017.
- 721 28 S. Kirkpatrick, C. D. Gelatt, , and M. P. Vecchi. Optimization by simulated annealing. Science, 220:671—
722 -680, 1983.
- 723 29 Yann-Hang Lee, K.P. Reddy, and C.M. Krishna. Scheduling techniques for reducing leakage power in
724 hard real-time systems. In Proceedings of the 15th Euromicro Conference on Real-Time Systems, pages
725 105–112, Jul. 2003.
- 726 30 E. Missimer, K. Missimer, and R. West. Mixed-criticality scheduling with I/O. In Proceedings of the 28th
727 Euromicro Conference on Real-Time Systems, pages 120—130, 2016.
- 728 31 Alessandro Vittorio Papadopoulos, Enrico Bini, Sanjoy Baruah, and Alan Burns. AdaptMC: A control-
729 theoretic approach for achieving resilience in mixed-criticality systems. In Proceedings of the 30th
730 Euromicro Conference on Real-Time Systems, pages 14:1–14:22, 2018.
- 731 32 Alessandro Vittorio Papadopoulos, Martina Maggio, Alberto Leva, and Enrico Bini. Hard real-time
732 guarantees in feedback-based resource reservations. Journal of Real-Time Systems, 51(3):221–246, 2015.
- 733 33 Jiankang Ren and Linh Thi Xuan Phan. Mixed-criticality scheduling on multiprocessors using task
734 grouping. In Proceedings of the 27th Euromicro Conference on Real-Time Systems, pages 25–34, 2015.
- 735 34 Paulo Baltarejo Sousa, Konstantinos Bletsas, Eduardo Tovar, Pedro Souto, and Benny Åkesson. Unified
736 overhead-aware schedulability analysis for slot-based task-splitting. Journal of Real-Time Systems,
737 50(5-6):680–735, 2014.
- 738 35 S. Vestal. Preemptive scheduling of multi-criticality systems with varying degrees of execution time
739 assurance. In Proceedings of the 28th IEEE Real-Time Systems Symposium, 2007.

740 Appendix A

741 The following figures present the schedulability ratio of the proposed heuristics with different
742 parameters. The variation in a parameter is given in the label of a figure, while the other parameters
743 are considered to be default.




Figure 14 Cooling rate = 0.001




Figure 15 Cooling rate = 0.005




Figure 16 Cooling rate = 0.01




Figure 17 Cooling rate = 0.05




Figure 18 Cooling rate = 0.1




Figure 19 Cooling rate = 0.3




Figure 20 Cooling rate = 0.5




Figure 21 Cooling rate = 0.7




Figure 22 Cooling rate = 0.9







Figure 23 Cooling rate = 1




■ Figure 24 HWCET scaling = 2




■ Figure 25 HWCET scaling = 3




■ Figure 26 HWCET scaling = 4




■ Figure 27 HWCET scaling = 5




■ Figure 28 HWCET scaling = 6




■ **Figure 29** H-tasks share = 20%




■ **Figure 30** H-tasks share = 30%




■ **Figure 31** H-tasks share = 40%




■ **Figure 32** H-tasks share = 50%



■ **Figure 33** H-tasks share = 60%



■ **Figure 34** H-tasks share = 70%



■ **Figure 35** H-tasks share = 80%




Figure 36 Fixed tasks per server, Servers = 2




Figure 37 Fixed tasks per server, Servers = 3




Figure 38 Fixed tasks per server, Servers = 4




Figure 39 Fixed tasks per server, Servers = 5




Figure 40 Fixed tasks per server, Servers = 6




Figure 41 Fixed task-set size, servers = 2




Figure 42 Fixed task-set size, servers = 3




Figure 43 Fixed task-set size, servers = 4




Figure 44 Fixed task-set size, servers = 5




Figure 45 Fixed task-set size, servers = 6




Figure 46 Task-set size = 9




Figure 47 Task-set size = 12




Figure 48 Task-set size = 15




Figure 49 Task-set size = 18




Figure 50 Task-set size = 21




Figure 51 Task-set size = 24




Figure 52 Server order = $U^H(\tilde{P}_i) - U^L(\tilde{P}_i)$




Figure 53 Server order = $U^H(\tilde{P}_i)/U^L(\tilde{P}_i)$




Figure 54 Server size variation $\beta = 0.001$




Figure 55 Server size variation $\beta = 0.005$




Figure 56 Server size variation $\beta = 0.01$




Figure 57 Server size variation $\beta = 0.02$




Figure 58 Server size variation $\beta = 0.05$




Figure 59 Server size variation $\beta = 0.1$




Figure 60 Server size variation $\beta = 0.2$




Figure 61 Server size variation $\beta = 0.3$




Figure 62 Server size variation $\beta = 0.4$




Figure 63 Server size variation $\beta = 0.5$




Figure 64 Server size variation $\beta = 0.6$




Figure 65 Server size variation $\beta = 0.7$




Figure 66 Server size variation $\beta = 0.8$




Figure 67 Server size variation $\beta = 0.9$




Figure 68 Temperature = 25




Figure 69 Temperature = 50




Figure 70 Temperature = 100




Figure 71 Temperature = 500




Figure 72 Temperature = 1000





Figure 73 Temperature = 5000



■ Figure 74 Temperature = 10000