



CISTER

Research Centre in
Real-Time & Embedded
Computing Systems

Conference Paper

Worst-case Stall Analysis for Multicore Architectures with Two Memory Controllers

Muhammad Ali Awan*

Pedro F. Souto

Konstantinos Bletsas*

Benny Åkesson*

Eduardo Tovar*

*CISTER Research Centre

CISTER-TR-180401

2018/07/03

Worst-case Stall Analysis for Multicore Architectures with Two Memory Controllers

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

*CISTER Research Centre

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8321159

E-mail: muaan@isep.ipp.pt, ksbs@isep.ipp.pt, kbake@isep.ipp.pt, emt@isep.ipp.pt

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

Abstract


In multicore architectures, there is potential for contention between cores when accessing shared resources, such as system memory. Such contention scenarios are challenging to accurately analyse, from a worst-case timing perspective. One way of making memory contention in multicore more amenable to timing analysis is the use of memory regulation mechanisms. It restricts the number of accesses performed by any given core over time by using periodically replenished per-core budgets. Typically, this assumes that all cores access memory via a single shared memory controller. However, ever-increasing bandwidth requirements have brought about architectures with multiple memory controllers. These control accesses to different memory regions and are potentially shared among all cores. While this presents an opportunity to satisfy bandwidth requirements, existing analysis designed for a single memory controller are no longer safe. This work formulates a worst-case memory stall analysis for a memory-regulated multicore with two memory controllers. This stall analysis can be integrated into the schedulability analysis of systems under fixed-priority partitioned scheduling. Five heuristics for assigning tasks and memory budgets to cores in a stall-cognisant manner are also proposed. We experimentally quantify the cost in terms of extra stall for letting all cores benefit from the memory space offered by both controllers as well as also evaluate the five heuristics for different system characteristics.

1 Worst-case Stall Analysis for Multicore 2 Architectures with Two Memory Controllers

3 **Muhammad Ali Awan**

4 CISTER Research Centre and ISEP, Porto, Portugal


5 muaan@isep.ipp.pt

6  <https://orcid.org/0000-0001-5817-2284>

7 **Pedro F. Souto**

8 University of Porto, Faculty of Engineering and CISTER Research Centre, Porto, Portugal


9 pfs@fe.up.pt

10  <https://orcid.org/0000-0002-0822-3423>

11 **Konstantinos Bletsas**

12 CISTER Research Centre and ISEP, Porto, Portugal


13 ksbs@isep.ipp.pt

14  <https://orcid.org/0000-0002-3640-0239>

15 **Benny Akesson**

16 Embedded Systems Innovation, Eindhoven, the Netherlands


17 benny.akesson@tno.nl

18  <https://orcid.org/0000-0003-2949-2080>

19 **Eduardo Tovar**

20 CISTER Research Centre and ISEP, Porto, Portugal

21 emt@isep.ipp.pt

22  <https://orcid.org/0000-0001-8979-3876>

23 — Abstract —

24 In multicore architectures, there is potential for contention between cores when accessing shared
25 resources, such as system memory. Such contention scenarios are challenging to accurately ana-
26 lyse, from a worst-case timing perspective. One way of making memory contention in multicores
27 more amenable to timing analysis is the use of memory regulation mechanisms. It restricts the
28 number of accesses performed by any given core over time by using periodically replenished per-
29 core budgets. Typically, this assumes that all cores access memory via a single shared memory
30 controller. However, ever-increasing bandwidth requirements have brought about architectures
31 with multiple memory controllers. These control accesses to different memory regions and are
32 potentially shared among all cores. While this presents an opportunity to satisfy bandwidth
33 requirements, existing analysis designed for a single memory controller are no longer safe.

34 This work formulates a worst-case memory stall analysis for a memory-regulated multicore
35 with two memory controllers. This stall analysis can be integrated into the schedulability analysis
36 of systems under fixed-priority partitioned scheduling. Five heuristics for assigning tasks and
37 memory budgets to cores in a stall-cognisant manner are also proposed. We experimentally
38 quantify the cost in terms of extra stall for letting all cores benefit from the memory space offered
39 by both controllers, and also evaluate the five heuristics for different system characteristics.

40 **2012 ACM Subject Classification** Computer systems organization → Real-time systems, Com-
41 puter systems organization → Real-time operating systems, Computer systems organization →
42 Real-time system architecture

43 **Keywords and phrases** multiple memory controllers, memory regulation, multicore

44 **Digital Object Identifier** 10.4230/LIPIcs.ECRTS.2018.2



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

30th Euromicro Conference on Real-Time Systems (ECRTS 2018).

Editor: Sebastian Altmeyer; Article No. 2; pp. 2:1–2:22

Leibniz International Proceedings in Informatics



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

45 **Funding** This work was partially supported by National Funds through FCT (Portuguese Found-
46 ation for Science and Technology), within the CISTER Research Unit (CEC/04234).

47 **1** Introduction

48 The strong trend towards increasing integration in hardware for embedded real-time systems
49 has led to multicores becoming mainstream platforms of choice for such systems. Multicores
50 have significant advantages in terms of computing power, energy usage and weight over
51 single-cores. Yet, one issue with multicores is that worst-case timing analysis becomes more
52 complicated. In particular, the fact that multiple cores contend for the same shared system
53 resources (buses, caches, memory) must be accounted for [8].

54 Focusing specifically on the problem of main memory contention, we note various research
55 efforts [21, 22, 15, 10, 5, 11, 13, 20, 14, 3] that employ *memory regulation* to make the memory
56 access patterns of the different cores more amenable to worst-case timing analysis. Under
57 memory regulation schemes, each core gets an associated periodically-replenished memory
58 access budget. When a core attempts to issue more memory accesses than its budget, it gets
59 temporarily stalled, until the next replenishment.

60 However, engineering practice forges ahead and analysis has to catch up. In recent years,
61 in response to memory bandwidth often becoming a performance bottleneck, multicore chips
62 that integrate, not one, but two memory controllers, have become commercially available.
63 In such platforms, both controllers are accessible by all cores, with little to no difference
64 in latency. Examples include various multicore processors from the NXP QorIQ series [16],
65 ranging from the P5020 with 2 cores to the P4080 with 8 cores. For existing approaches
66 to apply to systems with multiple controllers, one could statically map cores to memory
67 controllers and apply the analyses to each partition independently. This simple approach
68 efficiently reduces contention between cores. Still, it may be hard to find a partition such
69 that no tasks depend on data from the memory space of the other memory controller. Core-
70 to-controller partitioning also reduces flexibility in bandwidth allocation, as a partition's
71 bandwidth requirements must be met by just the associated memory controller. In cases
72 when no such partitions can be found, there are currently no good solutions, because existing
73 approaches can be *unsafe* when applied to platforms with two controllers. The reason is that
74 the worst-case memory access pattern for each controller in isolation will not necessarily lead
75 to the worst-case stall, as we demonstrate in Section 5. This reality motivated the present
76 work, whose main contributions are the following:

77 First, we show via counter-examples that existing techniques for upper-bounding the
78 memory stall, conceived for memory-regulated architectures with a single memory controller,
79 are not necessarily safe in the presence of multiple controllers. Our second and more important
80 contribution is new worst-case memory stall analysis for architectures with two memory
81 controllers, shared by all cores. This analysis, which presumes fixed task-to-core partitioning
82 and fixed-priority scheduling, can then be integrated to the schedulability analysis for the
83 system. Finally, we explore five different stall-cognisant heuristics for combined memory-
84 bandwidth-to-core assignment and task-to-core assignment and evaluate their performance
85 in terms of schedulability via experiments with synthetic task sets capturing different system
86 characteristics. These experiments also highlight the performance implications of having
87 fully shared memory controllers vs. partitioning the controllers to different cores, in cases
88 when the latter arrangement would be viable from the application perspective (i.e., no data
89 sharing across memory domains).

90 Next, in Section 2, we discuss related work. Section 3 defines our system model and

91 Section 4 discusses some relevant existing results from the single-controller case. Section 5
92 contains our analysis. Section 6 describes five proposed stall-cognisant task-to-core assignment
93 heuristics. Section 7 provides an experimental evaluation of our analysis and heuristics in
94 terms of theoretical schedulability using synthetic task sets. Section 8 concludes the paper.

95 **2 Related work**

96 Several software-based approaches for mitigating memory interference in multi-core plat-
97 forms [21, 22, 15, 10, 5, 11, 3] have been proposed in recent years. These approaches consider
98 a periodic server implemented in software that manages the memory budgets of the cores.
99 This is combined with run-time monitoring through performance counters that keep track of
100 the number of memory accesses and with an enforcement mechanism that suspends tasks
101 whenever they exhaust their budget. Our work is similar to these, as it exploits such a
102 memory throttling mechanism to enforce budgets on memory requests.

103 The memory regulation techniques used to mitigate the interference on shared memory
104 controllers introduce new stalls and the existing analyses are unsafe unless adapted to
105 account for them. Some efforts in this direction exist for partitioned fixed-priority schedul-
106 ing [21, 13] and hierarchical scheduling in [5]. Mancuso et al. [13], under their Single-Core
107 Equivalence framework [18], addressed the problem of fixed-priority partitioned schedulability
108 on a multicore. They employ the periodic software-based memory regulation mechanism
109 MemGuard [22] to ensure that each core gets an equal share of memory bandwidth in each
110 regulation interval (or period) and stalls until the end of the regulation period once the
111 budget has been depleted. Such stalls, resulting from the memory regulation together with
112 contention stalls are integrated into the schedulability analysis in [13].

113 Even if equal sharing of memory bandwidth is simple and facilitates porting applications
114 from a single-core to multi-core platforms (by making the analysis akin to that for a single-
115 core), it is inefficient when the memory requirements of the applications on different cores are
116 diverse. Yao et al. [20], and Pellizzoni and Yun [17] generalise the arrangement along with
117 the analysis to uneven memory budgets per core. The former approach considers round-robin
118 memory arbitration, whereas the latter proposes a new analysis for First-Ready First Come
119 First Served memory scheduling. Recently, Mancuso et al. [14] improved their memory stall
120 analysis by considering the exact memory bandwidth distribution on other cores. However,
121 all these approaches are designed to work with a *single memory controller* and are unsafe
122 with more than one memory controller. The reason is that the worst-case memory access
123 pattern for each controller in isolation no longer necessarily leads to the worst-case stall, as
124 we show in Section 5. In contrast, our work provides a worst-case memory stall analysis for a
125 memory-regulated multicore platform with two memory controllers and incorporates this stall
126 analysis in the schedulability analysis for fixed-priority partitioned preemptive scheduling.
127 We also present five memory bandwidth allocation and task-to-core assignment heuristics.

128 To summarise, existing works on memory regulation rely on an assumption of a single
129 memory controller. Here, we expand the state-of-the-art by proposing memory stall analysis,
130 when each core can access two controllers, facilitating data sharing among applications
131 and allowing more flexible use of bandwidth. We allow uneven distribution of the memory
132 bandwidth of each controller to available cores. Each core is scheduled under fixed-priority
133 preemptive scheduling, assuming a round-robin memory arbitration policy on both controllers.

134 3 System Model

135 We consider a platform with m identical cores (P_1 to P_m) and 2 memory controllers on the
 136 same chip, both uniformly accessible by all cores. The sets of memory regions accessible by
 137 the two controllers are non-overlapping. Examples of platforms with 2-8 identical cores and
 138 two memory controllers include NXP QorIQ P-series P4040, P4080, P5020 and P5040 [16].

139 Assume a set of n sporadic tasks, τ_1 to τ_n . Each task has a minimum interarrival time
 140 T_i , a deadline $D_i \leq T_i$, and a worst-case execution time (WCET) of C_i . Like Yao et al. [20],
 141 we assume that CPU computation and memory access do not overlap in time. Each task
 142 can access memory via both controllers. Therefore, $C_i = C_i^e + C_i^{m1} + C_i^{m2}$, where C_i^e is the
 143 worst-case CPU computation time and C_i^{m1} and C_i^{m2} are the worst-case total memory access
 144 times of a task via each respective controller in isolation.

145 The tasks are partitioned to the cores (no migration) and fixed-priority scheduling is used.
 146 For the memory controllers and their interconnects, we assume a round-robin policy [22, 20].
 147 The last-level cache (furthest from the cores) is either private or partitioned to each core. Like
 148 Yao et al. [20], we assume that access to main memory is regulated, e.g., by Memguard [22]
 149 or in hardware. We also require performance monitoring counters to count the number
 150 of memory accesses issued to each controller from each core. As in [20], we assume each
 151 memory access takes a constant time L . This allows us to specify P and C_i^e , C_i^{m1} and C_i^{m2}
 152 as multiples of L . Our model is agnostic w.r.t. the points in time when memory accesses may
 153 occur within the activation of a task and hence imposes no particular programming model.

154 Memory accesses are regulated as follows. Each core i has a *memory access budget* $Q1_i$
 155 for memory controller 1, which is the maximum allowed memory access time (measured in
 156 multiples of L) via that controller, within a *regulation period* of length P . Likewise, it has
 157 a budget $Q2_i$ for controller 2. These budgets are set at design time and may be different.
 158 A core i that consumes its memory access budget for a given memory controller within a
 159 regulation period is *stalled* until the start of the next regulation period¹. Regulation periods
 160 on all cores are synchronised. The *memory bandwidth share* of core i on controller 1 is
 161 $b1_i = \frac{Q1_i}{P}$. Similarly for $b2_i$ and controller 2. By design, $\sum_i b1_i \leq 1$ and $\sum_i b2_i \leq 1$, i.e., the
 162 bandwidth of any controller is not overcommitted.

163 4 Relevant existing results from the single-controller case

164 We now summarise some existing results from [20], for a similar, albeit single-controller,
 165 system, in order to later show why those do not apply, and new analysis is needed.

166 The technique in [20] calculates a worst-case stall term for each task, which is added to
 167 the right hand side of the standard worst-case response time (WCRT) recurrence relation
 168 for fixed priorities. For ease of presentation, the authors assume that there is a single task
 169 running on the core under consideration. Later on, for the case when many tasks are assigned
 170 to a core, they explain how to equivalently model the considered task τ_i and all higher-priority
 171 tasks as a single synthetic task, in order to apply their stall analysis and derive the worst-case
 172 stall term for τ_i . Below, we similarly assume a single task per core.

173 A memory request may stall either (i) because of requests from other cores, contending
 174 for the memory controller simultaneously (a case of **contention stall**) or (ii) because the

¹ On practical grounds, we assume that a core is stalled immediately after the Q^{th} memory access in a regulation period via the respective controller is served. Yao et al [20], more generously, assume that it is stalled immediately before attempting a $(Q + 1)^{th}$ access within the same regulation period.

175 issuing core has exhausted its budget for the current regulation period (a **regulation stall**).

176 Yao et al. identify worst-case patterns for memory accesses and computation within a
 177 single regulation period, characterised by maximum stall with the fewest memory accesses.
 178 Next, they use these patterns as main “building blocks” for the worst-case pattern for the
 179 entire task activation, over multiple regulation periods. In more detail:

180 **Case $b_i \leq 1/m$ (regulation dominant):** If $b_i \leq 1/m$, i.e., if the task’s bandwidth share
 181 is “fair” at most, then a task incurs worst-case stall when all its memory accesses are clustered
 182 at the start of its activation, before any computation. Another pessimistic assumption is
 183 that the task is released just after a regulation stall, so it waits for $(P - Q_i)$ until the next
 184 regulation period. The task will incur a stall of $(P - Q_i)$ within each of the next $\lfloor \frac{C_i^m}{Q} \rfloor$
 185 regulation periods; whether this is entirely due to a regulation stall or partially also due
 186 to contention from other cores is irrelevant. Afterwards, any remaining memory accesses
 187 (which are too few to trigger a regulation stall), can each incur a worst-case contention stall
 188 of $(m-1)$, i.e., one contending access from each other core due to round robin arbitration.

189 **Case $b_i > 1/m$ (contention dominant):** In this case, the smallest number of memory
 190 accesses per period a core must issue to get the maximum stall is $RBS \stackrel{\text{def}}{=} \frac{P_i - Q_i}{m-1}$, and occurs
 191 when the remaining budget is shared evenly among the other cores. From the assumption
 192 of the case, $b_i > 1/m$, it follows that $RBS < Q_i$. Therefore, the worst-case pattern for one
 193 regulation period involves $c_i^m = RBS$ accesses, each suffering a maximum contention stall of
 194 $(m - 1)$, for a total stall of $P - Q_i$. This leaves $Q_i - RBS$ time units not filled by memory
 195 accesses or respective stalls. These are filled with computation; if memory accesses were
 196 added instead, they would incur no stall. To bound the stall for the entire task activation,
 197 this pattern is applied to as many regulation periods as possible. Two subcases exist: either
 198 memory accesses or computation will run out first.

199 Due to space constraints, we refer to [20] for details. Meanwhile, some insights driving
 200 Yao’s analysis, for single-controller systems, are codified via the following lemmas from [20]:

201 **► Lemma 1.** *Considering the stall of a core due to memory regulation alone, the worst-case*
 202 *memory access pattern of one task is when all accesses within the task are clustered, and the*
 203 *stall is upper bounded by $P - Q_i$ for each regulation period P .*

204 **► Lemma 2.** *If the memory is not overloaded and the regulation periods are the same and*
 205 *synchronized, the stall due to inter-core memory contention alone on each core i with assigned*
 206 *budget Q_i is upper-bounded by $P - Q_i$ for every regulation period P .*

207 **► Lemma 3.** *Considering the contention stall alone, the maximum stall for core i with*
 208 *budget Q_i is obtained when the remaining budget $P - Q_i$ is evenly distributed among all other*
 209 *cores and they generate the maximum amount of accesses.*

210 **5 Analysis**

211 In this section, we formulate the main contribution of this paper: a stall analysis for multicores
 212 with two memory controllers, which leverages on Yao et al [20] stall and schedulability analysis
 213 for multicores with a single memory controller. First, we look at Lemmas 1 to 3 and Yao’s
 214 analysis in general, and examine what holds over from [20] and what does not. For readability,
 215 we omit the core (task) index, since it is implied. Table 1 summarizes the symbols used.

216 **5.1 What holds over from Yao’s analysis and what does not**

217 When we have multiple controllers, with an assigned memory budget Q_j for each, Lemma 1
 218 can be generalized as follows:

■ **Table 1** Symbols used in the analysis

$Q1, Q2$	memory budget on controllers 1 and 2, respectively
C^{m1}, C^{m2}	maximum number of memory accesses via controllers 1 and 2, respectively
C^e	worst-case computation time
P	regulation period
m	number of cores
$b1, b2$	core memory bandwidth shared on controllers 1 and 2, respectively
$RBS1, RBS2$	remaining budget share on controllers 1 and 2, respectively
c^{m1*}, c^{m2*}	worst-case number of accesses per period in contention-dominant case
$K1^*$	number of regulation periods of phase 1 in contention-dominant case
$\hat{C}^e, \hat{C}^{m1}, \hat{C}^{m2}$	task computation parameters after phase 1 (in contention-dominant case)
$\Delta\rho^*$	worst-case reduction in regulation stalls w.r.t. maximum regulation stalls in the third case (regulation is dominant only for one controller)
ΔC^e	additional "computation" added to contention-only phase by reducing the number of regulation stalls by 1
ΔC_c^{m2*}	additional number of contention stalls required when moving ΔC^e to ensure that the total stall is larger with one less regulation stall on controller 1
$\Delta C_c^{m2}(max)$	maximum number of additional contention stalls obtained by moving ΔC^e to the contention-only phase
$\Delta C_c^{m2}(min)$	minimum number of additional contention stalls obtained by moving ΔC^e to the contention-only phase
$r^m = \frac{C^{m2}}{C^{m1}}$	ratio of memory accesses to each controller
$C_{\bar{e}}^{m2}$	number of memory accesses via controller 2 without contention
$single()$	worst-case single controller stall according to Yao's analysis, ignoring the regulation stall at the beginning of the execution

219 ► **Lemma 4.** *Considering the stall of a core due to memory regulation alone **on controller***
 220 ***j** , with budget Qj , the worst-case memory access pattern of one task is when all accesses*
 221 ***via controller j** within the task are clustered, and the stall is upper bounded by $P - Qj$ for*
 222 *each regulation period P .*


223 A corollary of this lemma is that the regulation stall on controller j is maximum when there
 224 are no memory accesses to a second controller in that period. Note also that a core can only
 225 regulation-stall on at most one memory controller in a given regulation period.

226 With multiple controllers Lemmas 2 and 3 apply to *each controller separately*. Furthermore,
 227 because a core may access memory via multiple controllers in a single regulation period, a
 228 consequence of Lemma 2 is the following:

229 ► **Lemma 5.** *If the memory is not overloaded and the regulation periods are the same and*
 230 *synchronized, the stall due to inter-core memory contention alone on each core i with assigned*
 231 *budget Qj_i on controller j is upper-bounded by $\min\left(\sum_j (P - Qj_i), \frac{P}{m} \cdot (m - 1)\right)$ for every*
 232 *regulation period P .*

233 When there are multiple memory controllers, the maximum contention stall may occur when
 234 there are accesses via more than one controller. The first argument to the min operator in
 235 the above expression sums up the contention stall from each controller according to Lemma 2.
 236 The second argument expresses the fact that no more than P/m accesses (irrespective via
 237 which controller) can all suffer the worst-case per-access contention stall of $(m - 1)$ because
 238 of round robin arbitration. Both terms independently bound the contention stall.

239 When there are multiple shared controllers and we try to upper-bound the stall over
 240 multiple regulation periods, Yao's analysis may not be safe, i.e., it may underestimate the
 241 worst-case stall, as illustrated by the example of Figure 1. Execution i) has the worst-case
 242 stall, according to Yao's stall analysis, when in a regulation period all memory accesses are



■ **Figure 1** As shown in this example, the worst-case total stall is when there are memory accesses via more than one controller in the same regulation period.

243 via the same controller. In each period, the first two memory accesses suffer the maximum
 244 stall. However the remaining 4 memory accesses suffer no stall, because the maximum stall
 245 in every regulation period is 6, $P - Qi$, and it occurs in the first two memory accesses of the
 246 respective regulation period. Execution ii) shows the worst-case stall when there are accesses
 247 via both controllers in the same period. In each period, we have 2 memory accesses *via each*
 248 *controller* and each of these accesses suffers the maximum contention stall, $m - 1$. This is
 249 because the contention stall on accesses via one controller does not affect the contention stall
 250 on accesses via the other controller. Thus, in execution ii) all memory accesses suffer the
 251 maximum contention stall, whereas in execution i) only a third does.

252 5.2 Two-controller Task Stall Analysis

253 Having shown the need for a new analysis, we consider several cases depending on the values
 254 of $b1$ and $b2$. Some entail sub-cases. More specifically, we consider 3 cases:

- 255 1. $b1 \leq \frac{1}{m} \wedge b2 \leq \frac{1}{m}$
- 256 2. $b1 > \frac{1}{m} \wedge b2 > \frac{1}{m}$
- 257 3. remaining cases, i.e. $(b1 \leq \frac{1}{m} \wedge b2 > \frac{1}{m}) \vee (b1 > \frac{1}{m} \wedge b2 \leq \frac{1}{m})$

258 5.2.1 Case 1: $b1 \leq \frac{1}{m} \wedge b2 \leq \frac{1}{m}$

259 In this case, for each controller, the worst case occurs when there is a regulation stall, as
 260 shown in [20]. By Lemma 4, the following execution suffers the worst-case stall. In a first
 261 phase, there is the longest sequence of consecutive periods with regulation stalls on controller
 262 1, followed by a second phase consisting of the longest sequence of consecutive periods with
 263 regulation stalls on controller 2. Finally, there is a third phase with the remaining memory
 264 accesses via each controller, $C^{mi} \bmod Qi$, that suffer the maximum contention stall per
 265 memory access, $m - 1$, and any computation. Because in each of the two first phases all
 266 memory accesses are via a single controller, we can use Yao's stall analysis to compute an
 267 upper bound on the stall in each of these phases. The upper bound of the total stall can
 268 then be computed by adding the upper bounds for each phase. I.e.:

$$\begin{aligned}
 \text{Stall} = & \text{single}(C^m = \left\lfloor \frac{C^{m1}}{Q1} \right\rfloor \cdot Q1, C^e = 0, Q = Q1, P = P, m = m) \\
 & + \text{single}(C^m = \left\lfloor \frac{C^{m2}}{Q2} \right\rfloor \cdot Q2, C^e = 0, Q = Q2, P = P, m = m) \\
 & + (C^{m1} \bmod Q1 + C^{m2} \bmod Q2) \cdot (m - 1)
 \end{aligned} \tag{1}$$

270 where $\text{single}()$ is the stall based on Yao's (single controller) stall analysis for the respective
 271 set of parameter values [20].

272 **5.2.2 Case 2:** $b1 > \frac{1}{m} \wedge b2 > \frac{1}{m}$

273 In this case, according to Yao's analysis, for each controller, the worst case occurs when there
 274 is maximum contention stall in a regulation period with the minimum number of memory
 275 accesses. However, as shown in Figure 1, in this case the worst-case stall may occur when
 276 a task accesses memory via different controllers in the same regulation period. Therefore,
 277 the worst-case memory access pattern of a task in this case has 3 phases, as illustrated in
 278 Figure 2 i):

279 **Phase 1** In this phase, every regulation period incurs the maximum contention stall. This
 280 phase terminates when the task runs out of memory accesses via some controller, and
 281 therefore cannot sustain the maximum contention stall any more. In Figure 2 i), this
 282 phase spans the two first periods, and, in each period, there are $RBS1$ and $RBS2$ memory
 283 accesses via the respective controller.

284 **Phase 3** In this phase, all accesses are via a single controller. This phase may not exist, if the
 285 task runs out of memory accesses via both controllers in the same regulation period. In
 286 Figure 2 i), this is the 4th and last period and has memory accesses only via controller 1.

287 **Phase 2** This "middle" phase may also not exist, but if it exists, it has only one regulation
 288 period. In this phase, we have memory accesses via both controllers, but either there are
 289 not enough memory accesses via at least one of the controllers to ensure the maximum
 290 contention stall in that period, or there is not enough execution to fill the complete period.
 291 In Figure 2 i), this is the 3rd period, and has only one memory access via controller 2.


292 According to Lemma 5, there are two main cases for the maximum contention stall in a
 293 regulation period. We analyse each of these cases separately.

294 **5.2.2.1 Sub-case 1:** $(P - Q1) + (P - Q2) < \frac{P}{m} \cdot (m - 1)$

295 In this case, the maximum contention stall in a regulation period occurs when a task
 296 performs $RBS1$ memory accesses via controller 1 and $RBS2$ memory accesses via controller
 297 2. Therefore, the maximum stall per period is $(RBS1 + RBS2) \cdot (m - 1) = (P - Q1) + (P - Q2)$.
 298 Because the task is non preemptive and $(P - Q1) + (P - Q2) < \frac{P}{m} \cdot (m - 1)$, by the definition
 299 of the sub-case, there is a "hole" of size $P - (RBS1 + RBS2) \cdot m$ that must be filled with
 300 execution, i.e. either computation or memory accesses. An execution in which computation
 301 fills as many of these holes as possible suffers the maximum stall, because any additional
 302 memory accesses in these periods suffer no contention stall. This will minimize the number of
 303 memory accesses without contention in Phase 1, increasing the number of memory accesses
 304 in latter phases, and possibly their stall. Similar reasoning can be applied to Phase 2, as well.

305 Figure 2 illustrates an execution pattern that leads to the worst-case stall, based on the
 306 above observations. In execution i) there is enough computation to fill in the holes in Phases
 307 1 (the first two periods) and 2. However, there is not enough computation to ensure that all
 308 memory accesses suffer contention: in the 4th and last period, which belongs to Phase 3,
 309 there are 4 memory accesses via controller 1 that do not suffer any contention. In execution
 310 ii) there is not enough computation to fill the holes in Phase 2, and therefore, we have 6
 311 memory accesses via controller 1 in Phase 2, the 3rd period, that do not suffer any contention,
 312 and there is no 3rd Phase. In execution iii) there is no Phase 2, because all memory accesses
 313 via controller 2 are used to fill the holes in Phase 1. Phase 3 consists only of a single memory
 314 access via controller 1. Finally, in execution iv) there is not enough computation, and Phase
 315 1, like Phase 2, has only one period, and there is no Phase 3.

316 It can be shown, by case analysis, that in any of these executions swapping any com-
 317 putation or memory access in one regulation period with computation or memory accesses



■ **Figure 2** Example execution patterns with worst case stall, for the contention-dominant case when $(P - Q_1) + (P - Q_2) < \frac{P}{m} \cdot (m - 1)$

318 in later regulation periods does not lead to an increase in the total stall, and therefore the
 319 execution pattern shown suffers the maximum stall. The following stall analysis is based on
 320 the execution pattern shown in Figure 2.

321 In order to reuse the analysis in other cases below, let c^{m1*} and c^{m2*} be the minimum
 322 values of c^{m1} and c^{m2} , respectively, that maximize the contention stall in a regulation period,
 323 assuming that any holes are filled with computation. Note that by Lemma 5, it must be
 324 $c^{m1*} \leq RBS1$ and $c^{m2*} \leq RBS2$. In this sub-case, they are $RBS1$ and $RBS2$, respectively.

325 In our analysis, we consider Phase 1 separately from the remaining phases, if any.

326 **Phase 1 stall:** In Phase 1, the contention stall in every regulation period is maximum
 327 and equal to $(c^{m1*} + c^{m2*}) \cdot (m - 1)$. The total stall in this phase is:

$$328 \quad Stall1 = K1^* \cdot (c^{m1*} + c^{m2*}) \cdot (m - 1) \quad (2)$$

$$329 \quad \text{where: } K1^* = \min \left(\left\lfloor \frac{C^{m1}}{c^{m1*}} \right\rfloor, \left\lfloor \frac{C^{m2}}{c^{m2*}} \right\rfloor, \left\lfloor \frac{C^e + C^{m1} + C^{m2}}{P - (c^{m1*} + c^{m2*}) \cdot (m - 1)} \right\rfloor \right) \quad (3)$$

330

331 is the number of regulation periods in Phase 1. Indeed, to sustain maximum contention stall
 332 in every regulation period of Phase 1, the task must have both:

- 333 1. Enough memory accesses via controller 1, i.e. $K1^* \leq \left\lfloor \frac{C^{m1}}{c^{m1*}} \right\rfloor$.
- 334 2. Enough memory accesses via controller 2, i.e. $K1^* \leq \left\lfloor \frac{C^{m2}}{c^{m2*}} \right\rfloor$.
- 335 3. Enough execution, since when a core is not stalled it must be either computing or accessing
 336 memory, i.e. in every Phase 1 period a task must execute for $P - (c^{m1*} + c^{m2*}) \cdot (m - 1)$.

337 Therefore, $K1^* \leq \left\lfloor \frac{C^e + C^{m1} + C^{m2}}{P - (c^{m1*} + c^{m2*}) \cdot (m - 1)} \right\rfloor$.

338 We use the minimum of these 3 values, because this is the largest possible number of periods
 339 in Phase 1 and, as argued above, this leads to the worst-case stall.

340 **Remaining stall:** Without loss of generality, let $\left\lfloor \frac{C^{m1}}{c^{m1*}} \right\rfloor \geq \left\lfloor \frac{C^{m2}}{c^{m2*}} \right\rfloor$, i.e. controller 2 runs
 341 out of memory accesses entirely in Phase 2 the latest. (The other case is symmetric.)

342 To analyse the stall in Phases 2 and 3, if any, we consider the stall of each controller
 343 separately. Since memory accesses via controller 2 occur only in Phase 2 (which has at most
 344 one regulation period) and not in Phase 3, the contention stall on controller 2 can be upper
 345 bounded by $\min(\hat{C}^{m2}, RBS2) \cdot (m - 1)$, where \hat{C}^{m2} is the number of memory accesses via
 346 controller 2 in Phase 2, if any. Observe that these memory accesses and respective stall
 347 can be taken into account as computation in the analysis of the stall of memory accesses
 348 via controller 1, in Phase 2. Furthermore, in Phase 3, if any, all memory accesses are via
 349 controller 1, only. Therefore, we apply Yao's stall analysis to compute the stall of memory
 350 accesses via controller 1 in Phases 2 and 3, if they exist.

351 So, to complete analysis of this case, we compute \hat{C}^{m2} , as well as parameters for Yao's
 352 single controller stall analysis. Since in the latter we consider the remaining memory accesses
 353 via controller 2, \hat{C}^{m2} , and respective stall, if any, as computation, C^e is obtained by adding to
 354 that value the remaining computation, \hat{C}^e , i.e. the task computation that was not performed
 355 in Phase 1. Finally, the value of C^m to use in the single controller analysis is the number of
 356 memory accesses via controller 1 that were not performed in Phase 1, \hat{C}^{m1} , if any. Thus,

$$\begin{aligned}
 Stall &= Stall1 + \min(\hat{C}^{m2}, RBS2) \cdot (m - 1) \\
 &\quad + \text{single}(C^e = \hat{C}^{m2} + \min(\hat{C}^{m2}, RBS2) \cdot (m - 1) + \hat{C}^e, \\
 357 \quad C^m = \hat{C}^{m1}, Q = Q1, P = P, m = m)
 \end{aligned} \tag{4}$$

358 where $Stall1$ is given by (2). Next, we derive the expressions for \hat{C}^e , \hat{C}^{m1} and \hat{C}^{m2} .

359 In every Phase 1 period a task must execute, i.e. either compute or access memory, when
 360 it is not stalled. Thus, in addition to the $c^{m1*} + c^{m2*}$ memory accesses that lead to the
 361 maximum stall in a regulation period, a task may have to execute for the remaining time:
 362 $P - (c^{m1*} + c^{m2*}) \cdot m$. As we have argued, the total stall will be maximum in executions
 363 where computation fills as many of these "holes" as possible. Thus:

$$\hat{C}^e = \max(0, C^e - K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m)) \tag{5}$$

366 If there is enough computation to fill all these holes, $C^e \geq K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m)$,
 367 then $\hat{C}^{m1} = C^{m1} - K1^* \cdot c^{m1*}$ and $\hat{C}^{m2} = C^{m2} - K1^* \cdot c^{m2*}$.

368 If there is not enough computation to fill all these holes, then the remaining holes,
 369 $K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m) - C^e$, will be filled with memory accesses. Thus, the total
 370 number of memory accesses that will occur in the remaining phases, if any, is:

$$\begin{aligned}
 \hat{C}^m &= C^{m1} + C^{m2} - K1^* \cdot (c^{m1*} + c^{m2*}) - (K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m) - C^e) \\
 371 \quad &= C^{m1} + C^{m2} - (K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot (m - 1)) - C^e)
 \end{aligned} \tag{6}$$

372 To determine \hat{C}^{m1} and \hat{C}^{m2} , we distinguish two cases, depending on the value of $K1^*$.

373 If $K1^* = \left\lfloor \frac{C^{m2}}{c^{m2*}} \right\rfloor \left(\leq \left\lfloor \frac{C^{m1}}{c^{m1*}} \right\rfloor \right)$, then an execution that has at least $\min(C^{m1} - K1^* \cdot$
 374 $c^{m1*}, RBS1, \hat{C}^m)$ controller 1 memory accesses in the first period of the remaining phases,
 375 will suffer maximum stall, because all these memory accesses suffer maximum contention
 376 stall. The first bound is the number of memory accesses not used to ensure maximum stall in
 377 Phase 1, the second bound is the maximum number of accesses via controller 1 that can suffer
 378 maximum stall in a regulation period, and the third bound is the number of memory accesses
 379 in the remaining phases. This ensures that controller 2 runs out of memory accesses before
 380 controller 1, as shown in Figure 2 iii). Thus the number of memory accesses via controller
 381 2 in Phase 2 is $\hat{C}^{m2} = \min(\hat{C}^m - \min(C^{m1} - K1^* \cdot c^{m1*}, RBS1, \hat{C}^m), C^{m2} - K1^* \cdot c^{m2*})$
 382 i.e. the number of memory accesses via controller 2 in Phase 2 is the number of memory
 383 accesses not used to fill the holes in Phase 1, discounted by the minimum number of memory
 384 accesses via controller 1 that suffer maximum contention in Phase 2, and upper-bounded
 385 by the maximum number of controller 2 memory accesses that are not necessary to ensure
 386 maximum stall in Phase 1. Finally, $\hat{C}^{m1} = \hat{C}^m - \hat{C}^{m2}$.

387 If $K1^* = \left\lfloor \frac{C^e + C^{m1} + C^{m2}}{P - (c^{m1*} + c^{m2*}) \cdot (m - 1)} \right\rfloor$, there is not enough execution to complete the $K1^* + 1$ st
 388 regulation period, if any – the execution has at most one regulation period after Phase 1.

389 In this case, the total stall is maximum in executions where the number of contention
 390 stalls in the last period is maximum. However, there cannot be more than $RBS1$ ($RBS2$)

391 contention stalls on controller 1 (2, respectively) in this period. Like in the previous sub-case,
 392 an execution with at least $\min(C^{m1} - K1^* \cdot c^{m1*}, RBS1, \hat{C}^m)$ controller 1 memory accesses
 393 in Phase 2, guarantees that controller 2 runs out of memory accesses no later than controller
 394 1, and suffers maximum stall, because all these memory accesses suffer maximum contention
 395 stall. Thus, the expressions we derived for \hat{C}^{m1} and \hat{C}^{m2} in the previous sub-case, are also
 396 valid for this one. Summarizing, we get the following expressions:

$$397 \quad \hat{C}^{m2} = \begin{cases} C^{m2} - K1^* \cdot c^{m2*}, & \text{if } C^e \geq K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m) \\ \min(\hat{C}^m - \min(C^{m1} - K1^* \cdot c^{m1*}, RBS1, \hat{C}^m), C^{m2} - K1^* \cdot c^{m2*}), & \text{o.w} \end{cases} \quad (7)$$

$$398 \quad \hat{C}^{m1} = \begin{cases} C^{m1} - K1^* \cdot c^{m1*} & \text{if } C^e \geq K1^* \cdot (P - (c^{m1*} + c^{m2*}) \cdot m) \\ \hat{C}^m - \hat{C}^{m2} & \text{otherwise} \end{cases} \quad (8)$$

400 **5.2.2.2 Sub-case 2:** $(P - Q1) + (P - Q2) \geq \frac{P}{m} \cdot (m - 1)$

401 In this case (by the definition of *RBS*), $RBS1 + RBS2 \geq \frac{P}{m}$, and therefore it is possible
 402 to guarantee maximum contention stall in a period, without any computation or memory
 403 accesses without contention. To ensure the maximum stall, the memory accesses should be
 404 distributed in a “balanced” way so that both controllers run out of memory access at more
 405 or less the same time, thus ensuring that all C^m memory access suffers the maximum stall.

406 Let c^{m1*} and c^{m2*} be the number of memory accesses via controllers 1 and 2 per regulation
 407 period that maximize the contention stall in a period. The goal is then to ensure:

$$408 \quad \frac{C^{m1}}{c^{m1*}} = \frac{C^{m2}}{c^{m2*}} \Rightarrow c^{m2*} = \frac{C^{m2}}{C^{m1}} \cdot c^{m1*} \Rightarrow c^{m2*} = r^m c^{m1*}, \text{ where: } r^m \stackrel{\text{def}}{=} \frac{C^{m2}}{C^{m1}} \quad (9)$$

409 Without loss of generality, assume $r^m < 1$; the other case is symmetrical. Then it must be:

$$410 \quad c^{m1*} + c^{m2*} = \frac{P}{m} \Rightarrow (1 + r^m) \cdot c^{m1*} = \frac{P}{m} \Rightarrow c^{m1*} = \frac{P}{m \cdot (1 + r^m)} \quad (10)$$

$$411 \quad c^{m2*} = r^m \cdot c^{m1*} \Rightarrow c^{m2*} = r^m \cdot \frac{P}{m \cdot (1 + r^m)} \quad (11)$$

412 We now consider three sub-cases:


413 **Sub-case $c^{m1*} \leq RBS1 \wedge c^{m2*} \geq 1$:** In this case it is possible to ensure that all memory
 414 accesses suffer the maximum contention stall, even without any computation. Thus:

$$415 \quad Stall = (C^{m1} + C^{m2}) \cdot (m - 1) \quad (12)$$

416 Note that even though c^{m1*} or c^{m2*} may be fractional, these are average values. This means
 417 that in an execution with worst-case stall, the number of memory accesses via any controller
 418 may not be the same across all the regulation periods. However, there is an execution such
 419 that $c^{m1} + c^{m2} = \frac{P}{m}$, in all but possibly the last regulation period, and $c^{m1} \leq RBS1$ and
 420 $c^{m2} \leq RBS2$ in every regulation period.

421 **Sub-case $c^{m1*} > RBS1$:** In this case, both controllers would run out of computation at
 422 the same time only if the number of memory accesses via controller 1 exceeded *RBS1*, and
 423 therefore there would be memory accesses without any contention. An execution following
 424 the pattern illustrated in Figure 2, with $c^{m1*} = RBS1$ and $c^{m2*} = \min(\frac{P}{m} - RBS1, RBS2)$
 425 will have the worst-case stall, and therefore we can apply the analysis in Section 5.2.2.1.

426 **Sub-case $c^{m2*} < 1$:** In this case, both controllers would run out of computation at
 427 the same time only if there are some periods without memory accesses via controller 2.
 428 An execution following the pattern illustrated in Figure 2, with $c^{m2*} = 1$ and $c^{m1*} =$
 429 $\min(\frac{P}{m} - 1, RBS1)$ will have the worst-case stall, and therefore we can apply the analysis in
 430 Section 5.2.2.1.
 431



■ **Figure 3** Maximizing the number of regulation stalls may not lead to the worst-case stall.

432 5.2.3 Case 3: $(b1 \leq \frac{1}{m} \wedge b2 > \frac{1}{m}) \vee (b1 > \frac{1}{m} \wedge b2 \leq \frac{1}{m})$

433 In this case, executions with the maximum number of regulation stalls do not always lead
 434 to the worst-case stall. This is shown in Figure 3. In execution i), all memory accesses
 435 via controller 1 are clustered, causing two regulation stalls on controller 1, in the first two
 436 regulation periods. All the memory accesses via controller 2, occur in the third regulation
 437 period. Of these, only the first two suffer the maximum contention stall. The remainder suffer
 438 no contention, because the memory budget of the remaining cores, $P - Qi$, is exhausted by the
 439 stalls of the first 2 memory accesses. In execution ii), there is one memory access via controller
 440 1 in each period, and thus there is no regulation stall on controller 1, but each of these
 441 accesses suffers the maximum contention stall. Furthermore, in each of the first 3 periods,
 442 there are 2 memory accesses via controller 2, each of which suffers the maximum contention
 443 stall. Thus all memory accesses via both controllers suffer the maximum contention stall,
 444 and the total stall for execution ii) exceeds that of execution i). This is counter-intuitive,
 445 because the contention stall by accesses via controller 1 in execution ii), 12, is smaller than
 446 the regulation stall, 20, caused by the same number of accesses via controller 1 in execution
 447 i). However, this loss is more than compensated by the contention stall in execution ii) of
 448 the 4 memory accesses via controller 2 that suffer no contention stall in execution i). I.e.,
 449 although we are trading off a regulation stall, $P - Qi$, for contention stalls, presumably with
 450 maximum contention stall, $Qi \cdot (m - 1) < P - Qi$, we may also be adding stall to memory
 451 accesses via the second controller that previously suffered no stall.

452 Depending on whether $b1 \leq \frac{1}{m} \wedge b2 > \frac{1}{m}$ or $b1 > \frac{1}{m} \wedge b2 \leq \frac{1}{m}$, there are two sub-cases.
 453 Because they are symmetrical, we analyse only the former.

454 5.2.3.1 Sub-case 3.1: $b1 \leq \frac{1}{m} \wedge b2 > \frac{1}{m}$

455 Figure 3 shows that the maximum number of regulation stalls does not always lead to the
 456 worst-case stall. Furthermore, it can be shown that the total stall is maximum if there are
 457 no memory accesses via the second controller in periods with a regulation stall. Thus, the
 458 following memory access pattern with two phases leads to the worst-case stall: in the first
 459 phase, there is a number, possibly 0, of consecutive periods with regulation stalls; in the
 460 second phase, the contention-only phase, there is a number of consecutive periods, possibly
 461 only 1, with contention stalls only. Thus, the problem of finding the worst-case stall reduces
 462 to that of determining the number of regulation stalls that maximizes that stall. Actually, to
 463 simplify the mathematical expressions, we use the difference, $\Delta\rho^*$, between this number and
 464 the maximum number of regulation stalls, $\left\lfloor \frac{C^{m1}}{Q1} \right\rfloor$. The total stall can then be determined

Algorithm 1 Compute stall for each task.

Input: Parameters: C^{m1} , C^{m2} , m , C^e , $Q1$, $Q2$ and P (omitting task's index for simplicity)

Output: Stall

- 1: $b1 = \frac{Q1}{P}$, $b2 = \frac{Q2}{P}$, $RBS1 = \frac{P-Q1}{m-1}$, $RBS2 = \frac{P-Q2}{m-1}$ and $C = C^e + C^{m1} + C^{m2}$
- 2: **if** ($b1 \leq \frac{1}{m} \wedge b2 \leq \frac{1}{m}$) **then** ▷ Regulation stall is dominant for both controllers
- 3: Stall = Equation (1)
- 4: **else if** ($b1 > \frac{1}{m} \wedge b2 > \frac{1}{m}$) **then** ▷ Contention stall is dominant for both controllers
- 5: **if** ($(P - Q1) + (P - Q2) < \frac{P}{m} \cdot (m - 1)$) **then**
- 6: $c^{m1*} = RBS1$, $c^{m2*} = RBS2$
- 7: Compute Stall with Algorithm 2
- 8: **else** ▷ $(P - Q1) + (P - Q2) \geq \frac{P}{m} \cdot (m - 1)$
- 9: $r^m = \frac{C^{m2}}{C^{m1}}$, $c^{m1*} = \text{Equation 10}$, $c^{m2*} = \text{Equation 11}$
- 10: **if** ($r^m < 1$) **then**
- 11: **if** ($c^{m1*} \leq RBS1 \wedge c^{m2*} \geq 1$) **then**
- 12: Stall = Equation 12
- 13: **else if** ($c^{m1*} > RBS1$) **then**
- 14: $c^{m1*} = RBS1$, $c^{m2*} = \min(RBS2, \frac{P}{m} - RBS1)$
- 15: Compute Stall with Algorithm 2
- 16: **else** ▷ $c^{m2*} < 1$
- 17: $c^{m1*} = \min(RBS1, \frac{P}{m} - 1)$, $c^{m2*} = 1$
- 18: Compute Stall with Algorithm 2
- 19: **end if**
- 20: **else** ▷ $r^m \geq 1$: symmetric of previous case, swap indices
- 21: **end if**
- 22: **end if**
- 23: **else** ▷ Regulation stall is dominant for only one controller
- 24: **if** ($b1 \leq \frac{1}{m} \wedge b2 > \frac{1}{m}$) **then**
- 25: Compute $\Delta\rho^*$ using Algorithm 3
- 26: Stall = Equation 13
- 27: **else** ▷ $b2 \leq \frac{1}{m} \wedge b1 > \frac{1}{m}$: symmetric of previous case
- 28: **end if**
- 29: **end if**
- 30: **return** Stall + $= (P - \min(Q1, Q2))$ ▷ This adds the stall when the task arrives.

Algorithm 2 Compute stall for contention dominant case.

Input: Parameters: c^{m1*} , c^{m2*} , C^{m1} , C^{m2} , m , C^e , $Q1$, $Q2$ and P (omitting task's index)

Output: Stall

- 1: $b1 = \frac{Q1}{P}$, $b2 = \frac{Q2}{P}$, $RBS1 = \frac{P-Q1}{m-1}$, $RBS2 = \frac{P-Q2}{m-1}$ and $C = C^e + C^{m1} + C^{m2}$
- 2: $K1^* = \text{Equation 3}$,
- 3: Stall1 = Equation 2
- 4: $\hat{C}^e = \text{Equation 5}$, $\hat{C}^{m1} = \text{Equation 8}$, $\hat{C}^{m2} = \text{Equation 7}$
- 5: Stall23 = $\text{single}(C^e = \hat{C}^{m2} \cdot m + \hat{C}^e, C^m = \hat{C}^{m1}, Q = Q1, P = P, m = m)$
- 6: **return** Stall = Stall1 + $\min(\hat{C}^{m2}, RBS2) \cdot (m - 1)$ + Stall23 ▷ Equation 41

465 using Yao's stall analysis:

466

467
$$\text{Stall} = \text{single}(Q = Q1, C^m = C^{m1} - C^{m1} \bmod Q1 - \Delta\rho^* \cdot Q1, C^e = 0)$$

468

$$+ ((C^{m1} \bmod Q1) + \Delta\rho^* \cdot Q1) \cdot (m - 1)$$

469

470


$$+ \text{single}(Q = Q2, C^m = C^{m2}, C^e = C^e + ((C^{m1} \bmod Q1) + \Delta\rho^* \cdot Q1) \cdot m) \quad (13)$$

471

where, for computing the stall on the memory accesses via controller 2 in the second phase, we

472

account the memory accesses via controller 1 in the second phase and respective contention



■ **Figure 4** Upper (i) and lower (ii) bounds on ΔC_c^{m2} .

473 stalls as computation, assuming that each of them suffers the maximum contention stall
 474 under round-robin, $m - 1$. Algorithms 1 and 2 detail the case analysis that we have described
 475 so far in this section. In the following, we determine the value of $\Delta \rho^*$.

476 We consider two main sub-cases depending on whether there is enough computation,
 477 including residual memory accesses via controller 1, to ensure that every memory access via
 478 controller 2 suffers maximum contention.

479 5.2.3.2 Sub-case 1: Enough computation

480 If $C^e \geq \left\lfloor \frac{C^{m2}}{RBS2} \right\rfloor \cdot (P - m \cdot RBS2) - (C^{m1} \bmod Q1) \cdot m$, then every memory access in the
 481 contention-only phase suffers maximum contention, and therefore the total stall is maximum
 482 when the number of regulation stalls is maximum, i.e. $\Delta \rho^* = 0$.

483 5.2.3.3 Sub-case 2: Not enough computation

484 In this case, as illustrated in Figure 3, if there are memory accesses in the contention-only
 485 phase that suffer no contention, the worst-case stall may occur when the number of regulation
 486 stalls is not maximum.

487 When the number of regulation stalls is decremented by one, the regulation stall reduction
 488 by $P - Q1$ is partially compensated by an increase of the contention stall via controller 1 by
 489 $Q1 \cdot (m - 1)$. If the increase in contention stall via controller 2, $\Delta stall_c^2$ is such that:

$$490 \quad \Delta stall_c^2 > \Delta stall_c^{2*} \stackrel{\text{def}}{=} P - Q1 - Q1 \cdot (m - 1) = P - Q1 \cdot m \quad (14)$$

491 then reducing the number of regulation stall leads to a larger total stall. In other words, the
 492 total stall will be worse if the increase in the number of memory accesses with maximum
 493 stall, ΔC_c^{m2} , satisfies the following inequality:

$$494 \quad \Delta C_c^{m2} > \Delta C_c^{m2*} \stackrel{\text{def}}{=} \frac{\Delta stall_c^{2*}}{m - 1} = \frac{P - Q1 \cdot m}{m - 1} \quad (15)$$

495 Like in the analysis in Section 5.2.2, to compute the stall on memory accesses via
 496 controller 2, we can view the memory accesses via controller 1 and respective contention
 497 stall as computation. Thus, we need to determine ΔC_c^{m2} when the computation in the
 498 contention-only phase increases by $\Delta C^e = Q1 \cdot m$. The challenge is that this value, ΔC_c^{m2} ,
 499 may not be constant. I.e., when we increase the computation by $\Delta C^e = Q1 \cdot m$, ΔC_c^{m2} may
 500 have different values depending on other parameter values.

501 Our solution is to compute the maximum and minimum values of ΔC_c^{m2} , $\Delta C_c^{m2}(max)$
 502 and $\Delta C_c^{m2}(min)$, respectively, and then finding $\Delta \rho^*$ by case analysis, as described below.

503 When we increase the computation of the contention-only phase by ΔC^e , the total
 504 execution of that phase, including any contention, will increase at least by that much. This
 505 execution can replace memory accesses via controller 2 that did not have any contention, i.e.

506 memory accesses in excess of $RBS2$ accesses per period, which can then be shifted towards
 507 the end of the execution. ΔC_c^{m2} will be maximum if the shifted memory accesses are added
 508 to a regulation period with no memory accesses via controller 2, up to a limit of $RBS2$
 509 memory accesses per regulation period, as shown in Figure 4 i). Thus, in this case, as a
 510 result of adding ΔC^e memory accesses we get:

$$\begin{aligned}
 \Delta C_c^{m2}(max) &= RBS2 \cdot \left\lfloor \frac{\Delta C^e}{RBS2 + P - RBS2 \cdot m} \right\rfloor \\
 &\quad + \min(RBS2, \Delta C^e \bmod (RBS2 + P - RBS2 \cdot m)) \\
 511 &= RBS2 \cdot \left\lfloor \frac{\Delta C^e}{Q2} \right\rfloor + \min(RBS2, \Delta C^e \bmod Q2) \tag{16}
 \end{aligned}$$

512 The first term corresponds to the number of additional periods with $RBS2$ memory accesses.
 513 (Note that ΔC^e is used both to shift memory accesses via controller 2, and to fill the "hole"
 514 in the remaining of the period, $P - RBS2 \cdot m$.) The second term corresponds to the number
 515 of memory accesses in the last incomplete regulation period, if any: essentially, the memory
 516 accesses that can be replaced with the remaining of ΔC^e that was not used for the additional
 517 full periods, upper-bounded by $RBS2$.

518 On the other hand, ΔC_c^{m2} will be minimum, if, before adding ΔC^e , the execution ended
 519 immediately after the $RBS2$ accesses with contention. This is shown in Figure 4 ii). In this
 520 case, the analysis is similar to the one above, and therefore we can also use (16), except
 521 that rather than using ΔC^e , we need to use $\max(\Delta C^e - (P - RBS2 \cdot m), 0)$, because the
 522 remainder of the period at which the execution ended needs to be filled with "computation"
 523 before an earlier memory access via controller 2 without contention stall can experience the
 524 maximum contention stall by shifting it towards the end of the execution.

525 We can now distinguish there sub-cases, depending on the relative values of ΔC_c^{m2*} ,
 526 $\Delta C_c^{m2}(max)$ and $\Delta C_c^{m2}(min)$.

527 **Sub-case $\Delta C_c^{m2*} \geq C_c^{m2}(max)$:** In this case, the increase in the number of memory
 528 accesses with contention cannot make up for the eliminated regulation stall, so $\Delta \rho^* = 0$.

529 **Sub-case $\Delta C_c^{m2*} < C_c^{m2}(min)$:** In this case, the increase in the number of memory
 530 accesses with contention suffices to make up for the eliminated regulation stall. Therefore, the
 531 worst-case stall increases as we reduce the number of regulation stalls until one of the following
 532 3 cases occurs: 1) there are no more regulation stalls; 2) there are not enough memory
 533 accesses via controller 2, ΔC_c^{m2*} , without the maximum contention stall, to compensate
 534 for the loss in the regulation stall; or 3) the number of memory accesses via controller 1
 535 in at least one period of the second phase exceeds $Q1 - 1$, in which case we would have a
 536 regulation stall, and therefore there would be no reduction in the number of regulation stalls.

537 Because ΔC_c^{m2} varies, we do not know a closed form expression for the number of
 538 regulation periods to reduce. Thus, we use the iterative procedure shown in Algorithm 3.
 539 We hence start with $\Delta \rho^* = 0$ and keep increasing it by one until one of the above 3 stop
 540 conditions is satisfied. Specifically, while there are still enough memory accesses via controller
 541 2 without maximum contention stall, C_c^{m2} , and there is still one regulation stall (line 15),
 542 $\Delta \rho^*$ is tentatively increased by one. In each iteration, we tentatively compute the total stall
 543 using Yao's analysis with the appropriate parameters (line 18) and the number of memory
 544 accesses via controller 2 that suffer no contention (line 20), for the tentative value of $\Delta \rho^*$. If
 545 the number of memory accesses via controller 1 in all periods of the contention-only phase
 546 (line 21) does not exceed $Q1 - 1$, then the tentative values become definitive (line 22), and
 547 the algorithm loops again, otherwise it exits the loop and terminates.

548 **All other cases, i.e. $C_c^{m2}(min) \leq \Delta C_c^{m2*} < C_c^{m2}(max)$:** In this case, the total stall

Algorithm 3 Compute $\Delta\rho^*$

Input: Parameters: C^{m1} , C^{m2} , m , C^e , $Q1$, $Q2$ and P (omitting task index for simplicity)

Output: $\Delta\rho^*$

- 1: $RBS1 = \frac{P-Q1}{m-1}$, $RBS2 = \frac{P-Q2}{m-1}$ and $C = C^e + C^{m1} + C^{m2}$
- 2: $\Delta C^e = m \cdot Q1$
- 3: $\Delta C_c^{m2}(max) = \text{Equation 16}$
- 4: $\Delta C_c^{m2}(min) = \text{Equation 16, but replacing } \Delta C^e \text{ with } \max(\Delta C^e - (P - m \cdot RBS2), 0)$
- 5: $\Delta C_c^{m2*} = \lfloor \frac{P-m \cdot Q1}{m-1} \rfloor$
- 6: **if** $(C^e \geq \lfloor \frac{C^{m2}}{RBS2} \rfloor \cdot (P - m \cdot RBS2) - (C^{m1} \bmod Q1) \cdot m)$ **then**
- 7: $\Delta\rho^* = 0$ ▷ There is enough "computation"
- 8: **else if** $(\Delta C_c^{m2}(max) \leq \Delta C_c^{m2*})$ **then** ▷ Which implies $\Delta C_c^{m2}(min) \leq \Delta C_c^{m2*}$
- 9: $\Delta\rho^* = 0$ ▷ Maximize regulation stalls on Controller one
- 10: **else if** $(\Delta C_c^{m2}(min) > \Delta C_c^{m2*})$ **then** ▷ Which implies $\Delta C_c^{m2}(max) > \Delta C_c^{m2*}$
- 11: $\Delta\rho^* = 0$
- 12: $stall = single(Q = Q2, C^m = C^{m2}, C^e = C^e + (C^{m1} \bmod Q1) \cdot m)$
- 13: $R = C^{m2} + C^e + (C^{m1} \bmod Q1) \cdot m + stall$
- 14: $C_c^{m2} = C^{m2} - \lfloor \frac{R}{P} \rfloor \cdot RBS2 - min(\lfloor \frac{R \bmod P}{m} \rfloor, RBS2)$
- 15: **while** $(C_c^{m2} > \Delta C_c^{m2*} \wedge \Delta\rho^* < \lfloor \frac{C^{m1}}{Q1} \rfloor)$ **do**
- 16: $\Delta\rho_t^* = \Delta\rho^* + 1$
- 17: $\hat{C}^{m1} = C^{m1} \bmod Q1 + \Delta\rho_t^* \cdot Q1$ ▷ Accesses via controller 1 in second phase
- 18: $stall = single(Q = Q2, C^m = C^{m2}, C^e = C^e + \hat{C}^{m1} \cdot m)$
- 19: $R = C^{m2} + C^e + \hat{C}^{m1} \cdot m + stall$
- 20: $C_{ct}^{m2} = max(C^{m2} - \lfloor \frac{R}{P} \rfloor \cdot RBS2 - min(\lfloor \frac{R \bmod P}{m} \rfloor, RBS2), 0)$
- 21: **if** $(\hat{C}^{m1} - min(Q1 - 1, max(0, \lfloor \frac{R \bmod P}{m} \rfloor - RBS2))) \leq (Q1 - 1) \cdot \lfloor \frac{R}{P} \rfloor$ **then** ▷ Enough reg. periods to ensure that there is no reg. stall in periods with accesses via both controllers.
- 22: $\Delta\rho^* = \Delta\rho_t^*$, $C_c^{m2} = C_{ct}^{m2}$
- 23: **else break**
- 24: **end if**
- 25: **end while**
- 26: **else** ▷ $\Delta n_c^2(min) \leq \Delta n_c^{2*} < \Delta n_c^2(max)$
- 27: $\Delta\rho(max) = 0$, $stall(max) = 0$ ▷ Variables for maximum stall
- 28: **for** $\Delta\rho^* = 0$ **to** $\lfloor \frac{C^{m1}}{Q1} \rfloor$ **do** ▷ Do exhaustive search
- 29: $\hat{C}^{m1} = C^{m1} \bmod Q1 + \Delta\rho_t^* \cdot Q1$
- 30: $stall = single(Q = Q2, C^m = C^{m2}, C^e = C^e + \hat{C}^{m1} \cdot m)$ ▷ Cont. stall on both controllers
- 31: $R = C^{m2} + C^e + \hat{C}^{m1} \cdot m + stall$ ▷ Duration of contention-only phase
- 32: **if** $stall + (\lfloor \frac{C^{m1}}{Q1} \rfloor - \Delta\rho^*) \cdot (P - Q1) > stall(max)$
- 33: $\wedge (\hat{C}^{m1} - min(Q1 - 1, max(0, \lfloor \frac{R \bmod P}{m} \rfloor - RBS2))) \leq (Q1 - 1) \cdot \lfloor \frac{R}{P} \rfloor$ **then**
- 33: $stall(max) = stall + (\lfloor \frac{C^{m1}}{Q1} \rfloor - \Delta\rho^*) \cdot (P - Q1)$
- 34: $\Delta\rho^*(max) = \Delta\rho^*$
- 35: **end if**
- 36: **end for**
- 37: $\Delta\rho^* = \Delta\rho^*(max)$
- 38: **end if**
- 39: **return** $\Delta\rho^*$

549 sometimes increases when the number of regulations stalls decreases by one and sometimes it
550 does not. Thus in this case, our approach to find the value of $\Delta\rho^*$ is to compute the stall for
551 every possible value of $\Delta\rho^*$ and pick the one that leads to the maximum stall. Algorithm 3,
552 lines 27-37, details the computation of $\Delta\rho^*$ in this case.

Algorithm 4 Sensitivity analysis to reclaim memory bandwidth from both controllers

Input: $b1, b2, m, \Delta$ (threshold to stop the algorithm) and τ
Output: Minimum required memory bandwidth of both controllers

```

1:  $b_{min}^1 = 0, b_{max}^1 = b1, b_{min}^2 = 0, b_{max}^2 = b2$ 
2: while ( $b_{max}^1 - b_{min}^1 > \Delta \vee b_{max}^2 - b_{min}^2 > \Delta$ ) do
3:   for each controller  $j \in \{1, 2\}$  do
4:     if ( $b_{max}^j - b_{min}^j > \Delta$ ) then
5:        $X^j = \lfloor \frac{b_{min}^j + b_{max}^j}{2} \rfloor$ 
6:       if ( $j == 1$ ) then
7:          $Schedulability = \text{MultiControllerSchedulabilityAnalysis}(X^j, b_{max}^2, m, \tau)$ 
8:       else
9:          $Schedulability = \text{MultiControllerSchedulabilityAnalysis}(b_{max}^1, X^j, m, \tau)$ 
10:      end if
11:      if ( $Schedulability == \text{true}$ ) then  $b_{max}^j = X^j$ 
12:      else  $b_{min}^j = X^j$ 
13:      end if
14:    end if
15:  end for
16: end while
17: return  $\{b_{max}^1$  and  $b_{max}^2\}$ 

```

553 **5.3 Schedulability analysis**

554 Until now, we assumed one task per core. When many tasks are assigned to a core, the task
555 in consideration and those of higher priority can be modelled by one synthetic task, using the
556 approach in [20], and schedulability analysis can be performed as summarized in Section 4.

557 **6 Bandwidth Allocation and Task-to-core Assignment Heuristics**

558 We propose 5 heuristics for allocating tasks and memory bandwidth of both controllers to the
559 cores. They are evaluated in terms of system schedulability. We use Audsley’s algorithm [1]
560 to assign task priorities, even if it is no longer necessarily optimal in the presence of stalls.

561 **Even:** The total memory bandwidth of both controllers is equally distributed among all
562 cores. Subject to this even share, the task-to-core assignment is performed using first-fit.

563 **Uneven:** Initially, this heuristic also distributes both controller’s bandwidth evenly
564 among cores and employs the first-fit for task-to-core assignment. However, instead of
565 declaring failure whenever a task does not fit on any core, it sets that task aside, and moves
566 on to consider the next task. Any tasks that remain unassigned after considering all tasks,
567 are handled in-order as follows. Each core’s memory bandwidth from both controllers is
568 “trimmed” to the minimum value that preserves schedulability, via the sensitivity analysis
569 of Algorithm 4, explained later in this section. Let the total reclaimed bandwidth from all
570 cores be $B1$ and $B2$ from controllers 1 and 2, respectively. A second round of first-fit tries to
571 assign the remaining tasks, assuming that the bandwidth of the target core i is increased by
572 $B1$ and $B2$ for controllers 1 and 2, respectively. Upon successfully assigning such a task, we
573 trim anew the target cores’s memory budgets via sensitivity analysis, adjust the available
574 reclaimed budgets and move on to the next task in a similar manner.

575 **Greedy-fit:** Initially, the total memory bandwidth of both controllers is assigned to the
576 first core and the task-set is iterated over once (in a given order) to assign the maximum
577 number of tasks to this core; if a task does not fit, we try the next one. Afterwards, the
578 spare bandwidth from each controllers on this core is reclaimed via sensitivity analysis, and

579 is fully assigned to next core. And so on, until all tasks are assigned or the cores run out.

580 **Humble-fit:** Similar to greedy-fit, except that when a task assignment fails, we move to
581 the next core (attempting no more task assignments on the current one).

582 **Memory-fit:** Initially, $b1_i = b2_i = 0$, for every core i , where bx_i is the allocated memory
583 bandwidth of controller x on core i . Each task is assigned to the core i that requires the
584 least increase to $b1_i + b2_i$ to accommodate it, subject to existing task assignments.

585 “Uneven” explores a larger solution space than “Even.” “Greedy-fit” and “Humble-fit”
586 aggressively optimise for processing capacity use foremost. Conversely, “Memory-fit” optimises
587 for bandwidth instead. Hence, all heuristics sample the solution space in different ways.

588 **Sensitivity analysis:** Algorithm 4 presents the sensitivity analysis that trims the unused
589 memory bandwidth from both controllers and outputs the least required memory bandwidth
590 from each controller. This sensitivity analysis, used for bandwidth optimisation, is an
591 adaptation of binary interval search ([19, 2]). It gives both controllers an equal chance to
592 preserve their bandwidth in a round-robin fashion. By comparison, completely optimizing one
593 controller followed by the second one, may lead to an imbalanced approach, hence avoided.

594 **7 Evaluation**

595 **Experimental Setup** We developed a Java tool for our experiments. Its first module
596 generates the synthetic task sets and sets up a platform with the given input parameters. A
597 second module performs task-to-core allocation and feasibility analysis with two controllers.

598 We generate the task-set with a given target $U = x \cdot m : x \in (0, 1]$ using UUnifast-discard
599 algorithm [6, 9] for unbiased distribution of task utilisations. The task-set size is given as
600 input. Task periods are log-uniform-distributed, in the range 10-100 ms. We assume implicit
601 deadlines, even if our analysis also holds for constrained deadlines. The WCET of a task is
602 derived as $C_i = U_i \cdot T_i$. The total memory accesses of each task are randomly selected in
603 the range $[0, \Gamma \cdot C_i]$, with memory intensity factor $\Gamma \in (0, 1]$ user-defined. The total memory
604 accesses are randomly divided between the two memory controllers. By default the task-set
605 is sorted in descending order of utilisation. For each set of input parameters, we generate
606 1000 task-sets. We use independent pseudo-random number generators for the utilisations,
607 minimum inter-arrival times/deadlines, memory accesses and reuse their seeds [12]. Table 2
608 summarises all parameters, with default values underlined. We observed that size of the
609 regulation period has no effect on the schedulability ratio.

610 To avoid having hundreds of plots, in each experiment we vary only one parameter, with
611 others conforming to the defaults from Table 2 and present the results as plots of *weighted*
612 *schedulability*. This performance metric, adopted from [4], condenses what would have been
613 three-dimensional plots into two dimensions. It is a weighted average that gives more weight to
614 task-sets with higher utilisation, which are supposedly harder to schedule. Specifically, using
615 notation from [7], let $S_y(\tau, p)$ represent the result (0 or 1) of the schedulability test y for a
616 given task-set τ with an input parameter p . Then $W_y(p)$, the weighted schedulability for that
617 test y as a function p , is $W_y(p) = \sum_{\forall \tau} (\bar{U}(\tau) \cdot S_y(\tau, p)) / \sum_{\forall \tau} \bar{U}(\tau)$. Here, $\bar{U}(\tau) \stackrel{\text{def}}{=} U(\tau)/m$
618 is the system utilisation, normalised by the number of cores m .

619 No other stall analysis with two controllers exists in the literature to compare with. We
620 therefore compare our approach against a system where the two controllers are partitioned
621 among cores that can only make requests to their assigned controller. The benefit of such
622 partitioning is that it roughly cuts contention in half. On the other hand, tasks assigned to
623 one controller cannot access data addressable by the other controller.

624 For the comparison, half the cores are assigned to each controller. Since each core

■ **Table 2** Overview of Parameters

Parameters	Values	Default
Number of cores (m)	{4, 8, 12, 16}	4
Task-set size (n)	{8, 16, 24, 32, 40, 48}	16
Regulation period (P)	{1us, 10us, 100us, 1ms}	100us
Inter-arrival time (T_i)	10ms to 100ms	N/A
Nominal utilisation ($\bar{U} = \frac{U}{m}$)	{0.1 : 0.01 : 1}	N/A
Memory intensity (Γ)	{0.1 : 0.1 : 1}	0.5

625 accesses only one controller, the feasibility of the tasks assigned to it can be tested with Yao's
 626 analysis [20]. We adapt the task-to-core assignment heuristics and bandwidth allocation
 627 schemes presented in Section 6 for the partitioned case: The even heuristic equally divides a
 628 controller's bandwidth among its associated cores. Similarly, in the uneven heuristic, the
 629 readjustment of the controllers bandwidth is performed only among the controller's associated
 630 cores. In the greedy-fit/humble-fit, all bandwidth of a given controller is only assigned to its
 631 first associated core with an objective to maximise the number of tasks assigned to it. The
 632 trimmed-off bandwidth from this controller is assigned to its remaining associated cores. If
 633 the task is not feasible on the cores associated to the first controller, its feasibility is next
 634 checked on the set of cores associated with the second controller. In the memory-fit, a task
 635 is assigned to the core with the lowest bandwidth requirement of its controller. We use Yao-
 636 and MC- prefixes to denote the partitioned and our approach, respectively, followed by the
 637 name of the heuristic (even, uneven, greedy-fit, humble-fit and memory-fit).

638 **Results** Figure 5 presents the weighted schedulability for different number of cores for
 639 both systems with partitioned and shared controllers (our approach) using the proposed
 640 heuristics. The first important result is that all heuristics under partitioning perform better
 641 than their corresponding heuristic under shared controllers, which is due to the stall being
 642 roughly cut in half in the former approach. This difference ranges around 10% – 30% in
 643 absolute terms of weighted schedulability. Of course, this expected result applies only when
 644 there are no dependencies across partitions. However, in many systems, there is always
 645 some sharing/communication of data among tasks and this might make such partitioning
 646 impossible. In other cases, a single controller cannot deliver enough bandwidth. This may
 647 become more frequent in the future, as applications getting more demanding. Therefore safe
 648 analysis for predictable access to both controllers, like the one proposed here, is needed.

649 In terms of heuristics, memory-fit, uneven, even, humble-fit and greedy-fit is the descending
 650 ordered list w.r.t. weighted schedulability ratio. The memory-fit heuristic, which optimises
 651 the use of memory bandwidth, performing best, implies that memory bandwidth is typically
 652 the scarce resource for the given set of parameters. The uneven and even heuristics are more
 653 balanced in terms of bandwidth and processing capacity distribution and hence, perform
 654 close to memory-fit. Humble-fit and greedy-fit are too aggressive in construction to optimise
 655 the use of processing capacity at the cost of memory resources and hence underperform
 656 the other heuristics in a memory-scarce setup. Greedy-fit manages the memory resources
 657 comparatively better than humble-fit and hence, outperforms it. Yet, if the applications are
 658 compute-intensive and the system is not scarce w.r.t. memory resource, the heuristics that
 659 optimise for processing resources may become handy and outperform their counterparts.

660 With more cores, the contention from other cores increases and hence, the schedulability
 661 of the system decreases. Figure 6 presents the effect of memory intensity over the proposed




Figure 5




Figure 6




Figure 7




Figure 8

662 heuristics. Obviously, higher memory intensity increases the contention on the shared
 663 controllers, consequently decreasing the schedulability. We also compared the effect of
 664 the task indexing over the different heuristics as shown in Figure 7. The numbers 0, 1, 2
 665 and 3 on the X-axis correspond to task-set ordering w.r.t. descending order of deadlines,
 666 utilisation, total memory requests and memory density (i.e. total memory requests divided
 667 by the T_i), respectively. Task-set indexing w.r.t. utilisation benefits the memory-fit, even
 668 and uneven heuristics. Figure 8 shows that task-set size has very limited effect on the
 669 memory-fit, uneven and even approaches and they scale well when that increases. Conversely,
 670 the performance of humble-fit and greedy-fit degrade with greater task-set sizes due to their
 671 aggressive optimisation of processor usage at the expense of memory bandwidth.

672 8 Conclusion

673 This paper demonstrated that worst-case memory stall analyses for single-memory-controller
 674 multicores with memory regulation are unsafe if applied to multicores with multiple memory
 675 controllers. We overcome this limitation by proposing a new memory stall analysis for
 676 multicore platforms with two memory controllers that captures the cases where all cores can
 677 access both controllers. We also proposed five memory allocation heuristics, each specialising
 678 in optimising processing capacity and/or memory bandwidth. The experimentally quantified
 679 cost of allowing all cores to flexibly access the memory space of two controllers is 10 – 30%
 680 in terms of weighted schedulability. Results further show that the proposed memory-fit
 681 heuristic performs well if bandwidth is scarce. The even and uneven heuristics are suitable for
 682 balanced systems, while greedy-fit and humble-fit are handy for compute-intensive systems.

References

- 683
684 1 N. C. Audsley. On priority assignment in fixed priority scheduling. *Information Processing*
685 *Letters*, 79(1):39–44, 2001.
- 686 2 M. A. Awan, K. Bletsas, P. F. Souto, and E. Tovar. Semi-partitioned mixed-criticality
687 scheduling. In *Proceedings of the 30th International Conference on the Architecture of Com-*
688 *puting Systems (ARCS 2017)*, pages 205–218, 2017. doi:10.1007/978-3-319-54999-6_
689 16.
- 690 3 Muhammad Ali Awan, Pedro Souto, Konstantinos Bletsas, Benny Akesson, and Eduardo
691 Tovar. Mixed-criticality scheduling with memory bandwidth regulation. In *Proceedings of*
692 *the 55th IEEE/ACM Conference on Design Automation and Test in Europe (DATE 2018)*,
693 March 2018.
- 694 4 A. Bastoni, B. B. Brandenburg, and J. H. Anderson. Cache-related preemption and mi-
695 gration delays: Empirical approximation and impact on schedulability. *Proceedings of the*
696 *OSPERT*, pages 33–44, 2010.
- 697 5 M. Behnam, R. Inam, T. Nolte, and M. Sjödin. Multi-core composability in the face of
698 memory-bus contention. *ACM SIGBED Review*, 10(3):35–42, 2013. doi:10.1145/2544350.
699 2544354.
- 700 6 E. Bini and G. C. Buttazzo. Measuring the performance of schedulability tests. *Journal*
701 *of Real-Time Systems*, 30(1-2):129–154, May 2005. doi:10.1007/s11241-005-0507-9.
- 702 7 A. Burns and R. I. Davis. Adaptive mixed criticality scheduling with deferred preemption.
703 In *Proceedings of the 35th IEEE Real-Time Systems Symposium (RTSS 2014)*, pages 21–30,
704 Dec 2014. doi:10.1109/RTSS.2014.12.
- 705 8 D. Dasari, B. Akesson, V. Nélis, M. A. Awan, and S. M. Petters. Identifying the sources
706 of unpredictability in cots-based multicore systems. In *Proceedings of the 8th IEEE In-*
707 *ternational Symposium on Industrial Embedded Systems (SIES 2013)*, pages 39–48, June
708 2013.
- 709 9 R. I. Davis and A. Burns. Priority assignment for global fixed priority pre-emptive schedul-
710 ing in multiprocessor real-time systems. In *Proceedings of the 30th IEEE Real-Time Systems*
711 *Symposium (RTSS 2009)*, pages 398–409, Dec 2009. doi:10.1109/RTSS.2009.31.
- 712 10 J. Flodin, K. Lampka, and W. Yi. Dynamic budgeting for settling dram contention of
713 co-running hard and soft real-time tasks. In *Proceedings of the 9th IEEE International*
714 *Symposium on Industrial Embedded Systems (SIES 2014)*, pages 151–159, June 2014. doi:
715 10.1109/SIES.2014.6871199.
- 716 11 Rafia Inam, Nesredin Mahmud, Moris Behnam, Thomas Nolte, and Mikael Sjodin. Multi-
717 core composability in the face of memory-bus contention. In *Proceedings of the 20th IEEE*
718 *Real-Time Technology and Applications Symposium (RTAS 2014)*, 2014.
- 719 12 Raj Jain. *The art of computer systems performance analysis - techniques for experimental*
720 *design, measurement, simulation, and modeling*. Wiley professional computing. Wiley, 1991.
- 721 13 R. Mancuso, R. Pellizzoni, M. Caccamo, L. Sha, and H. Yun. WCET(m) estimation in
722 multi-core systems using single core equivalence. In *Proceedings of the 27th Euromicro*
723 *Conference on Real-Time Systems (ECRTS 2015)*, pages 174–183, July 2015. doi:10.
724 1109/ECRTS.2015.23.
- 725 14 Renato Mancuso, Rodolfo Pellizzoni, Neriman Tokcan, and Marco Caccamo. WCET Deriv-
726 ation under Single Core Equivalence with Explicit Memory Budget Assignment. In *Proceed-*
727 *ings of the 29th Euromicro Conference on Real-Time Systems (ECRTS 2017)*, volume 76 of
728 *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:23, Dagstuhl, Ger-
729 many, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- 730 15 J. Nowotsch, M. Paulitsch, D. Bühler, H. Theiling, S. Wegener, and M. Schmidt. Multi-core
731 interference-sensitive WCET analysis leveraging runtime resource capacity enforcement. In

- 732 *Proceedings of the 26th Euromicro Conference on Real-Time Systems (ECRTS 2014)*, pages
733 109–118, 2014. doi:10.1109/ECRTS.2014.20.
- 734 **16** NXP. QorIQ Layerscape Processors Based on Arm Technology, 2018. [www.nxp.com/products/processors-and-microcontrollers/applications-processors/
735 qorIQ-platforms/p-series](http://www.nxp.com/products/processors-and-microcontrollers/applications-processors/qorIQ-platforms/p-series).
- 736
- 737 **17** R. Pellizzoni and H. Yun. Memory servers for multicore systems. In *Proceedings of the 22nd
738 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2016)*,
739 pages 97–108, April 2016. doi:10.1109/RTAS.2016.7461339.
- 740 **18** Lui Sha, Marco Caccamo, Renato Mancuso, Jung-Eun Kim, Man-Ki Yoon, Rodolfo Pel-
741 lizzoni, Heechul Yun, Russel Kegley, Dennis Perlman, Greg Arundale, Bradford Richard,
742 et al. Single core equivalent virtual machines for hard real-time computing on multicore
743 processors. Technical report, Univ. of Illinois at Urbana Champaign, 2014.
- 744 **19** Paulo Baltarejo Sousa, Konstantinos Bletsas, Eduardo Tovar, Pedro Souto, and Benny
745 Åkesson. Unified overhead-aware schedulability analysis for slot-based task-splitting.
746 *Journal of Real-Time Systems*, 50(5-6):680–735, 2014.
- 747 **20** G. Yao, H. Yun, Z. P. Wu, R. Pellizzoni, M. Caccamo, and L. Sha. Schedulability ana-
748 lysis for memory bandwidth regulated multicore real-time systems. *IEEE Transactions on
749 Computers*, 65(2):601–614, Feb 2016.
- 750 **21** H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha. Memory access control in
751 multiprocessor for real-time systems with mixed criticality. In *Proceedings of the 24th
752 Euromicro Conference on Real-Time Systems (ECRTS 2012)*, pages 299–308, July 2012.
753 doi:10.1109/ECRTS.2012.32.
- 754 **22** H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha. Memguard: Memory bandwidth
755 reservation system for efficient performance isolation in multi-core platforms. In *Proceedings
756 of the 19th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS
757 2013)*, pages 55–64, April 2013. doi:10.1109/RTAS.2013.6531079.