

1 Improving the accuracy of cache-aware response 2 time analysis using preemption partitioning

3 Filip Marković 

4 Mälardalen University, Sweden

5 filip.markovic@mdh.se

6 Jan Carlson

7 Mälardalen University, Sweden

8 jan.carlson@mdh.se

9 Sebastian Altmeyer

10 University of Augsburg, Germany

11 altmeyer@informatik.uni-augsburg.de

12 Radu Dobrin

13 Mälardalen University, Sweden

14 radu.dobrin@mdh.se

15 — Abstract —

16 Schedulability analyses for preemptive real-time systems need to take into account cache-related
17 preemption delays (CRPD) caused by preemptions between the tasks. The estimation of the CRPD
18 values must be sound, i.e. it must not be lower than the worst-case CRPD that may occur at runtime,
19 but also should minimise the pessimism of estimation. The existing methods over-approximate
20 the computed CRPD upper bounds by accounting for multiple preemption combinations which
21 cannot occur simultaneously during runtime. This over-approximation may further lead to the
22 over-approximation of the worst-case response times of the tasks, and therefore a false-negative
23 estimation of the system's schedulability. In this paper, we propose a more precise cache-aware
24 response time analysis for sporadic real-time systems under fully-preemptive fixed priority scheduling.
25 The evaluation shows a significant improvement over the existing state of the art approaches.

26 **2012 ACM Subject Classification** Computer systems organization → Real-time system specification;
27 Software and its engineering → Real-time schedulability

28 **Keywords and phrases** Real-time systems, Fixed-Priority Preemptive Scheduling, Preemption delay

29 **Digital Object Identifier** 10.4230/LIPIcs.ECRTS.2020.7

30 **Acknowledgements** We want to thank our colleagues Sebastian Hahn, Jan Reineke, and Darshit
31 Shah, who provided us with the evaluation data derived from the code-level analysis of benchmark
32 programs. Also, we are very grateful to Davor Ćirkinagić who borrowed his powerful computing
33 system for performing schedulability evaluation.

34 **1** Introduction

35 Fully-preemptive scheduling is used in many real-time embedded systems in order to e.g.,
36 overcome the limitations of non-preemptive scheduling which can introduce significant
37 blocking on high priority tasks from lower priority ones. Fully-preemptive scheduling allows
38 for an interruption (preemption) of the task's execution whenever a task with a higher
39 priority is released. However, as shown by Pellizzoni et al. [21], a preemption can introduce a
40 significant preemption related delay, even up to 33% of the task's worst-case execution time.

41 In embedded systems employing a cache-based architecture, one of the major causes of
42 preemption delay is cache-related preemption delay (CRPD), as shown by Bastoni et al. [7].
43 CRPD represents the longest time needed by a resuming task to reload the memory cache
44 blocks which it had loaded prior to the preemption. Since CRPD may significantly increase



© Filip Marković, Jan Carlson, Sebastian Altmeyer, and Radu Dobrin;
licensed under Creative Commons License CC-BY

32nd Euromicro Conference on Real-Time Systems (ECRTS 2020).

Editor: Marcus Völp; Article No. 7; pp. 7:1–7:23



Leibniz International Proceedings in Informatics

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

7:2 Improving the accuracy of cache-aware RTA using preemption partitioning

45 the worst-case execution time of a task, its tight estimation is very important and therefore
46 a new field of timing analysis, called cache-related preemption delay analysis, has emerged in
47 the research of real-time systems.

48 In the context of CRPD analysis for fixed-priority fully-preemptive scheduling, many
49 different approaches have been proposed in the last few decades, of which we describe a
50 selection of the most recent ones. Tomiyama et al. [31], and Busquets-Mataix et al. [11]
51 proposed analyses which are based on the over-approximation that a single preemption
52 causes the CRPD equal to the time needed for reloading all the evicting cache blocks from
53 a preempting task. These analyses neglected the fact that not every eviction results in a
54 cache block reload. Contrary to this, Lee et al. [20] proposed the analysis that bounds the
55 CRPD by accounting for the cache blocks which may be reused at some later point in a
56 task, called the useful cache blocks. However, their analysis did not account for the fact that
57 although useful, some cache block cannot be evicted, and thus cannot result in a cache block
58 reload. These two opposite approaches defined the two main branches in CRPD analysis:
59 ECB-based CRPD and UCB-based CRPD.

60 Later, Tan and Mooney [30] proposed the UCB-union approach, which accounted for the
61 limitations of the above-described approaches. In this approach, CRPD is computed using
62 the information about all possibly affected useful cache blocks along with the evicting cache
63 blocks from the tasks which may evict them. Opposite to that, Altmeyer et al. [3] proposed
64 the ECB-union approach, where the all possibly evicting cache blocks are analysed along
65 with the useful cache blocks from the tasks that may be preempted.

66 The latest and in overall the most precise CRPD approaches are proposed by Altmeyer et
67 al. [4], called ECB-union multiset and UCB-union multiset approaches. Those approaches are
68 improvements over the UCB-union and ECB-union because they account for a more precise
69 estimation of the nested preemptions. The multiset approach was also used by Staschulat et
70 al. [28] for the periodic task systems where they accounted that each additionally accounted
71 preemption of a single preempting task may result in a smaller CRPD value compared to
72 the previous preemptions. In the context of periodic systems, Ramaprasad and Mueller
73 [23, 22] investigated the possibility of tightening the CRPD bounds using preemption patterns.
74 However, in this paper, we consider a sporadic task model which constrains such analysis.

75 Furthermore, in recent years, several cache-aware analysis were proposed in the contexts
76 of: cache partitioning by Altmeyer et al. [26, 5], cache-persistence by Rashid et al. [25, 24]
77 and Stock et al. [29], write-back cache by Davis et al. [12] and Blaß et al. [9].

78 In all of the above-mentioned CRPD analyses, the resulting upper bounds are overly
79 pessimistic mainly because they account for CRPD obtained from preemption combinations
80 which cannot occur simultaneously during runtime.

81 In this paper, we propose a cache-aware response-time analysis that accounts for the
82 above-mentioned source of pessimism, and a few more, in the context of fixed priority
83 fully-preemptive scheduling (FPPS). The evaluation shows a significant improvement over
84 the existing state of the art approaches.

85 In the remainder of the paper, we first define the system model in Section 2. In Section
86 3, we overview the existing SOTA UCB-union and ECB-union based methods, including
87 the multiset variants. In Section 4, we discuss the pessimism in the current state of the
88 art cache-aware analyses. The proposed analysis is defined in Section 5, and the evaluation
89 results are shown in Section 6. The paper is concluded in Section 7.

2 Task Model, Terminology and Notation

In this paper, we consider a sporadic task model, with preassigned fixed task priorities, under fully-preemptive scheduling. A taskset Γ consists of n tasks sorted in a decreasing priority order, where each task τ_i generates an infinite number of jobs, characterised with the following task parameters $\langle P_i, C_i, T_i, D_i \rangle$. Task priority is denoted with P_i and we assume disjunct priorities among the tasks. The worst-case execution time without accounted preemption delays is denoted with C_i . T_i denotes the minimum inter-arrival time between the two consecutive jobs of τ_i , and the relative deadline is denoted with D_i .

We also consider single-core systems with single-level direct-mapped caches, extending the task model by accounting for detailed knowledge about the cache usage. In addition, we describe a possible adjustment in subsection 5.8, thus also considering LRU set-associative caches. For each task τ_i , the information about the accessed cache blocks within the task's execution is assumed as derived, and based on that we define the following cache block sets: ECB_i – a set of *evicting cache blocks* of τ_i , such that cache-set s is in ECB_i if and only if a memory block from s may be accessed during the execution of τ_i .

UCB_i – a set of all *useful cache blocks* throughout the execution of τ_i . As proposed by Lee et al. [20], and superseded by Altmeyer et al. [2], a cache-set s is in UCB_i , if and only if τ_i accesses a memory block m in s such that: a) m must be cached at some program point \mathcal{P} in the execution of τ_i , and b) m may be reused on at least one control flow path starting from \mathcal{P} without the eviction of m on this path.

In the remainder of the paper, we use the following notations for different sets of tasks:

- ◇ $hp(i)$ A set of tasks with priorities higher than τ_i
- ◇ $hpe(i)$ A set of tasks with priorities higher than τ_i , including τ_i , i.e. $hpe(i) = hp(i) \cup \{\tau_i\}$
- ◇ $lp(i)$ A set of tasks with priorities lower than τ_i
- ◇ $aff(i, h)$ A set of tasks with priorities higher than or equal to τ_i and lower than τ_h , i.e. $aff(i, h) = hpe(i) \cup lp(h)$

3 Background

Cache-related preemption delay is computed as the upper bound on the number of cache block reloads that can be caused due to preemptions and potential evictions of memory contents that are used by the preempted tasks. In this paper, CRPD is denoted with γ and is computed as the multiplication of the upper bound on cache block reloads with the constant BRT , which is the longest time needed for a single memory block to be reloaded into cache memory, i.e. block reload time. The general formula is $\gamma = \#reloads \times BRT$.

In this section, we briefly describe the most relevant CRPD-aware analyses for understanding the contributions of this paper. We describe *UCB-union* approach [30], *ECB-union* approach [3], and their multiset variants [4].

UCB-union and *ECB-union* approaches are computed as the least-fixed points of the following equation for the worst-case response time:

$$R_i^{(l+1)} = C_i + \sum_{\tau_h \in hp(i)} \lceil R_i^{(l)} / T_h \rceil \times (C_h + \gamma_{i,h}) \quad (1)$$

In Equation 1, $\gamma_{i,h}$ represents the CRPD due to a single job of a higher priority task τ_h executing within the worst-case response time of task τ_i . This term is computed differently for each of the CRPD approaches.

7:4 Improving the accuracy of cache-aware RTA using preemption partitioning

132 *UCB-Union approach* computes $\gamma_{i,h}^{ucbu}$ with the following equation, accounting that a job
 133 of τ_h causes a reload of each cache block which it may access and which is useful during the
 134 execution of at least one of the tasks from the range $[\tau_{h+1}, \tau_{h+2}, \dots, \tau_i]$.

$$135 \quad \gamma_{i,h}^{ucbu} = \left| \left(\bigcup_{\tau_k \in \text{aff}(i,h)} UCB_k \right) \cap ECB_h \right| \times BRT \quad (2)$$

136 *ECB-Union approach* computes $\gamma_{i,h}^{ecbu}$ with the following equation, accounting that a job
 137 of τ_h is preempted by all of the tasks with higher priority than τ_h , after the job directly
 138 preempted one of the tasks from the range $[\tau_{h+1}, \tau_{h+2}, \dots, \tau_i]$. In this case, the preemption
 139 resulting in the highest number of evicted useful cache blocks is considered.

$$140 \quad \gamma_{i,h}^{ecbu} = \max_{\tau_k \in \text{aff}(i,h)} \left\{ \left| \left(\bigcup_{\tau_{h'} \in \text{hpe}(h)} ECB_{h'} \right) \cap UCB_k \right| \right\} \times BRT \quad (3)$$

141 Improving the above two approaches, Altmeyer et al. [4] introduced *UCB-Union multiset* and
 142 *ECB-Union multiset* which are computed as the least-fixed points of the following equation:

$$143 \quad R_i^{(l+1)} = C_i + \sum_{\tau_h \in \text{hp}(i)} \left(\lceil R_i^{(l)} / T_h \rceil \times C_h + \gamma_{i,h} \right) \quad (4)$$

144 where $\gamma_{i,h}$ represents the CRPD due to each job of a higher priority task τ_h executing within
 145 the the worst-case response time of task τ_i .

146 *ECB-Multiset approach* computes $\gamma_{i,h}^{ecbum}$ accounting that τ_h can preempt each task
 147 $\tau_k \mid h < k \leq i$ the maximum number of times a single job of τ_k can be preempted by jobs of
 148 τ_h , for each job of τ_k that can be released within R_i , i.e. $\lceil R_k / T_h \rceil \times \lceil R_i / T_k \rceil$ times. This is
 149 accounted by the multiset $M_{i,h}$, which consists of the maximum CRPDs from jobs of τ_h on
 150 each preemptable job which can be released within R_i (\uplus represents multiset union):

$$151 \quad M_{i,h} = \biguplus_{\tau_k \in \text{aff}(i,h)} \left(\biguplus_{\lceil R_k / T_h \rceil \times \lceil R_i / T_k \rceil} \left| UCB_k \cap \left(\bigcup_{\tau_{h'} \in \text{hpe}(h)} ECB_{h'} \right) \right| \right) \quad (5)$$

152 Based on the above multiset, $\gamma_{i,h}^{ecbum}$ is computed as the sum of the maximum $\lceil R_i / T_h \rceil$ values
 153 from $M_{i,h}$, accounting that only $\lceil R_i / T_h \rceil$ jobs of τ_h can directly preempt and cause CRPD
 154 on the preemptable jobs accounted in $M_{i,h}$.

155 *UCB-Multiset approach* computes $\gamma_{i,h}^{ucbum}$ by first computing the multiset $M_{i,h}^{ucb}$ which
 156 consists of all possibly useful cache blocks from jobs which can be released within R_i , and
 157 have priority higher than or equal to τ_i , and lower than τ_h .

$$158 \quad M_{i,h}^{ucb} = \biguplus_{\tau_k \in \text{aff}(i,h)} \left(\biguplus_{\lceil R_k / T_h \rceil \times \lceil R_i / T_k \rceil} UCB_k \right) \quad (6)$$

159 Next, this approach computes the multiset $M_{i,h}^{ecb}$ which consists of all possibly evicting
 160 cache blocks within jobs of τ_h that can be released within R_i . The following equation includes
 161 an instance of evicting cache block from τ_h for each job of τ_h that can be released within R_i :

$$162 \quad M_{i,h}^{ecb} = \biguplus_{\lceil R_i / T_h \rceil} \left(ECB_h \right) \quad (7)$$

163 The upper bound on CRPD from jobs of τ_h preempting all jobs from τ_{h+1} to τ_i is equal
 164 to the size of intersection of those multisets, with accounted block reload time:

$$165 \quad \gamma_{i,h}^{ucbum} = |M_{i,h}^{ucb} \cap M_{i,h}^{ecb}| \times BRT \quad (8)$$

166 The Combined-Multiset approach first computes the worst-case response time R_i^{ecbum}
 167 using Equation 4 and $\gamma_{i,h}^{ecbum}$, and similarly does with UCB-Union multiset, using Equation 4
 168 and $\gamma_{i,h}^{ucbum}$ thus deriving R_i^{ucbum} . Then, the final result is computed as $\min(R_i^{ecbum}, R_i^{ucbum})$.




Figure 1 Example of the pessimistic CRPD estimation in both, UCB- and ECB-union based approaches. Notice that the worst-case execution time (black rectangles) is in reality significantly larger than CRPD, but the focus of the figure is rather on preemptions and CRPD depiction.

4 Pessimism in CRPD analyses based on UCB- and ECB-union approaches

In this section, we present the identified problems considering CRPD over-approximation when using UCB- and ECB-based approaches, including the multiset variants.

Problem 1: Combined approach over-approximates the CRPD bounds because all preemptions that may occur within a response time R_i are treated the same, with at most one method at a time. However, within all preemptions that may occur within R_i , different preemption sub-groups may be analysed with different analyses, thus the CRPD may be further reduced by computing the bounds for different preemption sub-groups individually instead of computing it with one method at a time for a single group of all preemptions.

Problem 2: Combined approach accounts for CRPD from many different preemption combinations, which cannot occur together. This is presented with the following example. In Figure 1, we present three tasks τ_1 , τ_2 , and τ_3 with their respective sets of evicting and useful cache blocks. In the example, it is assumed that tasks τ_1 and τ_2 can be released at most once during the execution of τ_3 and that block reload time is equal to 1. Based on this, only two preemption combinations which result in the worst-case CRPD bound are possible: 1) A job of τ_2 directly preempts a job of τ_3 , and a job of τ_1 directly preempts a job of τ_2 (nested preemptions), 2) A job of τ_2 directly preempts a job of τ_3 , and a job of τ_1 directly preempts a job of τ_3 (direct preemptions).

For each task, black rectangles in the figures represent the worst-case execution time, grey rectangles represent CRPD, whereas the sets of integer values above the grey rectangles represent the cache sets whose reloads must be accounted.

Considering the given cache block sets, the actual worst-case CRPD, based on the separately analysed preemption combinations, is:

- ◇ 8 (nested preemption): This is the case because τ_2 evicts cache blocks 3, 4, 7, and 8 which are then reloaded during the post-preemption execution of τ_3 . After that, τ_1 evicts blocks 1 and 2 when preempting τ_2 , which are reloaded during the post-preemption execution of τ_2 , and also τ_1 evicts cache blocks 5 and 6 which are reloaded during the post-preemption execution of τ_3 . Notice that although τ_1 also potentially evicts blocks 3, and 4 from τ_3 , they are accounted as reloads only once within τ_3 , because τ_3 is interrupted once and thus only one reload of each useful cache block within remaining execution of τ_3 is possible.
- ◇ 8 (direct preemptions): This is the case because τ_2 evicts cache blocks 3, 4, 7, and 8 from τ_3 , and τ_1 evicts cache blocks 3, 4, 5, and 6 from τ_3 .

Since any other preemption combination can be derived only by removing one of the preemptions accounted in the two above, the worst-case CRPD is equal to 8. However,

7:6 Improving the accuracy of cache-aware RTA using preemption partitioning

204 UCB-union based approaches (including the multiset variant) compute the following CRPD:

$$\begin{aligned} \gamma_{i,h}^{ucbu} &= \left| \left(\bigcup_{\tau_k \in \text{aff}(i,h)} UCB_k \right) \cap ECB_h \right| \\ \gamma_{3,1}^{ucbu} &= \left| (UCB_3 \cup UCB_2) \cap ECB_1 \right| = 6, \quad \gamma_{3,2}^{ucbu} = \left| UCB_3 \cap ECB_2 \right| = 4 \\ 205 \quad \gamma_{3,1}^{ucbu} + \gamma_{3,2}^{ucbu} &= 4 + 6 = 10, \quad \text{accounted reloads for blocks: } 1, 2, 3, 3, 4, 4, 5, 6, 7, 8 \end{aligned}$$

206 UCB-union based approaches compute CRPD upper bound of 10 reloads, thus approximating
207 two block reloads over the safe upper bound (8 reloads) illustrated in Figure 1. Compared
208 to the leftmost case from Figure 1, the accounted infeasible reloads are for blocks 3 and 4.
209 Compared to the rightmost case, the accounted infeasible reloads are for blocks 1 and 2.

210 ECB-union based approaches (including the multiset variant) compute the CRPD upper-
211 bound as follows:

$$\begin{aligned} \gamma_{i,h}^{ecbu} &= \max_{\tau_k \in \text{aff}(i,h)} \left\{ \left| \left(\bigcup_{\tau_{h'} \in \text{hpe}(h)} ECB_{h'} \right) \cap UCB_k \right| \right\} \\ \gamma_{3,1}^{ecbu} &= \max_{\tau_k \in \{2,3\}} \left\{ \left| (ECB_1) \cap UCB_k \right| \right\} = \max \left\{ \left| ECB_1 \cap UCB_2 \right|, \left| ECB_1 \cap UCB_3 \right| \right\} = 4 \\ \gamma_{3,2}^{ecbu} &= \max \left\{ \left| (ECB_1 \cup ECB_2) \cap UCB_3 \right| \right\} = 6 \\ 212 \quad \gamma_{3,1}^{ecbu} + \gamma_{3,2}^{ecbu} &= 4 + 6 = 10 \end{aligned}$$

213 Similarly to UCB-union based approaches, ECB-union based approaches compute CRPD
214 upper bound of 10 reloads, thus approximating two cache block reloads over the safe bound.

215 Even when the lowest bound of the two is selected, CRPD bound is over-approximated
216 by accounting for two cache block reloads which cannot occur in a single combination of
217 preemptions during runtime. CRPD over-approximation is further increased when multiple
218 jobs of each task are introduced. In this paper, we propose a novel method for computing
219 the CRPD and the worst-case response time, accounting for the above-described problems.

220 5 CRPD-aware Response-Time Analysis

221 In the remainder of the paper, when we refer to the term *preemption* we consider both,
222 indirect (nested) and direct preemptions. We start with defining a cache-aware worst-case
223 response time equation, slightly different than the existing ones. Formally, the response time
224 analysis is defined as the least fixed-point of the following equation:

$$225 \quad R_i^{(l+1)} = C_i + \gamma(i, R_i^{(l)}) + \sum_{\tau_h \in \text{hp}(i)} \left\lceil \frac{R_i^{(l)}}{T_h} \right\rceil C_h \quad (9)$$

226 Notice that unlike in the existing approaches, Equation 9 computes the CRPD upper bound
227 $\gamma(i, t)$, which is a function that implicitly accounts for all preemptions that can occur within
228 duration t , between the first i tasks of Γ . A CRPD upper bound on all preemptions that can
229 occur within duration t can be computed more accurately by applying the following four
230 steps that we describe in more detail in the remainder of this section:

- 231 1. Derive all possible preemptions which can occur within duration t , between the jobs of
232 the first i tasks of Γ , (described in Subsection 5.1).
- 233 2. Divide the possible preemptions into partitions such that each partition accounts for
234 single-job preemptions between the tasks, (described in Subsection 5.2)
- 235 3. Compute the CRPD bounds for each partition individually, (described in Subsection 5.3).
- 236 4. Sum the CRPD bounds of all partitions to obtain the cumulative CRPD bound on all
237 possible preemptions within duration t , (described in Subsection 5.4).




Figure 2 Worst-case preemptions for τ_3 during the time duration $t = 46$.

To show an overview of how the proposed analysis works, we provide the running example from Figure 2, and we compute the upper bound on CRPD within the time duration $t = 46$. The analysis computes the bound as follows:

- Derive all possible preemptions which can occur within duration t , between the jobs of the first i tasks of Γ .

Example: Given the tasks from Figure 2, during the 46 time units, task τ_1 can preempt τ_3 at most two times, and it can preempt τ_2 at most once. Also, τ_2 can preempt τ_3 at most once. More formally, a single preemption from a job of τ_h on a job of τ_j is represented with an ordered pair (τ_h, τ_j) . Thus, all possible preemptions, within 46 time units, can be represented by the following multiset of ordered preemption pairs:

$$\{(\tau_1, \tau_3), (\tau_1, \tau_3), (\tau_1, \tau_2), (\tau_2, \tau_3)\} \quad (10)$$

- Divide the possible preemptions into partitions such that each partition accounts for single-job preemptions between the tasks.

Example: To represent the partitions, we generate the multiset Λ which consists of all possible preemptions, divided into partitions that account for single-job preemptions between the tasks. Given the possible preemptions from the multiset derived in the previous step, the multiset Λ of all partitions is:

$$\Lambda = \left\{ \begin{array}{cc} \textit{Partition 1} & \textit{Partition 2} \\ \downarrow & \downarrow \\ \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\} & , \quad \{(\tau_1, \tau_3)\} \end{array} \right\}$$

The multiset Λ consists of two partitions (each represented as a set of preemptions), such that the first partition consists of the following preemptions $\{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\}$, meaning that it is possible that τ_1 preempts τ_2 , that τ_1 preempts τ_3 , and that τ_2 preempts τ_3 . Jointly, the preemptions consist of all possible preemptions among the three tasks within a duration of 46 time units.

- Compute the CRPD bounds for each partition individually.

Example: As we showed in the previous section, when a single job of each task may preempt the other jobs with lower priority, the upper bound on CRPD is 8 time units. This is the upper bound on all preemptions accounted in *Partition 1*. Considering *Partition 2*, it consists of a single preemption, from a job of τ_1 on τ_3 , and in this case the upper bound is 4 time units since the preemption may lead to the reloads of cache blocks 3, 4, 5 and 6.

- Sum the CRPD bounds of all partitions to obtain the cumulative CRPD bound on all possible preemptions within duration t .

Example: The sum of upper bounds for *Partition 1* and *Partition 2* is $8 + 4 = 12$, which is the upper bound on all preemptions within 46 time units of the three shown tasks.

271 In the remainder of this section, we formally define the introduced terms, and prove that
 272 the proposed analysis results in a safe CRPD upper bound. The running example remains
 273 and it serves for better understanding on how the above values are computed and what they
 274 formally represent. This section is divided into the following subsections: 5.1 – describes the
 275 computation of upper bounds on the number of preemptions, 5.2 – describes the preemption
 276 partitioning, 5.3 – computation of CRPD bound for single partition, 5.4 – computation
 277 of CRPD bound for all preemptions, 5.5 – correctness proof for the computation of the
 278 worst-case response time, 5.6 – time complexity, and 5.7 – the additional computation for
 279 CRPD bound for single partition, based on finding the worst-case preemption combination.

280 5.1 Computing the upper bounds on the number of preemptions

281 ► **Definition 1.** An upper bound $E_j^h(t)$ on times a task τ_h can preempt τ_j ($h < j$) within
 282 duration t is defined with the following equation:

$$283 \quad E_j^h(t) = \begin{cases} \left\lceil \frac{t}{T_h} \right\rceil & , \left\lceil \frac{t}{T_h} \right\rceil \leq \left\lceil \frac{t}{T_j} \right\rceil \\ \left\lceil \frac{t}{T_j} \right\rceil \times \left\lceil \frac{R_j}{T_h} \right\rceil & , \left\lceil \frac{t}{T_h} \right\rceil > \left\lceil \frac{t}{T_j} \right\rceil \end{cases} \quad (11)$$

284 ► **Proposition 2.** $E_j^h(t)$ is an upper bound on number of possible preemptions from τ_h on τ_j
 285 within duration t .

286 **Proof.** Let us consider the following cases:

287 $\left\lceil \frac{t}{T_h} \right\rceil \leq \left\lceil \frac{t}{T_j} \right\rceil$: Each job of τ_h can preempt τ_j at most once, therefore the number of τ_h jobs
 288 which can be released within duration t is a safe bound on the number of preemptions from
 289 τ_h on τ_j within t .

290 $\left\lceil \frac{t}{T_h} \right\rceil > \left\lceil \frac{t}{T_j} \right\rceil$: An upper bound on preemptions from jobs of τ_h on a single job of τ_j is equal
 291 to $\left\lceil \frac{R_j}{T_h} \right\rceil$ since it is also an upper bound on number of times that jobs of τ_h can be released
 292 within the worst-case response time R_j of a single job. Since Equation 11 applies the bound
 293 $\left\lceil \frac{R_j}{T_h} \right\rceil$ on each job of τ_j which can be released within t , the proposition holds. ◀

294 5.2 Preemption partitioning

295 Once the all possible preemptions which can occur within duration t are identified, we divide
 296 them into partitions, such that no partition accounts for the same preemption pair of the
 297 first i tasks in Γ , and such that all partitions jointly account for all possible preemptions.

298 ► **Definition 3.** A multiset $\Lambda_{i,t}$ of partitions consisting of all possible preemptions that can
 299 occur within duration t , between the jobs of the first i tasks of Γ .

$$300 \quad \Lambda_{i,t} = \{\lambda_1, \lambda_2, \dots, \lambda_z\} \text{ such that } \lambda_r = \{(\tau_h, \tau_j) \mid r \leq E_j^h(t)\} \quad (12)$$

301 In Equation 12, $\Lambda_{i,t}$ is defined as a multiset of of sets (partitions). Each set λ_r consists of
 302 possible preemptions and each preemption is represented as an ordered pair (τ_h, τ_j) where
 303 the first element represents the preempting, and the second element represents the preempted
 304 task. The multiset Λ is formed of exactly z partitions, where $z = \max\{E_j^h(t) \mid 1 \leq h < j \leq i\}$.
 305 Each partition consists of disjunct preemptions, meaning that no partition contains two same
 306 preemption pairs.

307 *Example:* Given the taskset from the running example in Figure 2, the multiset $\Lambda_{3,46}$ is
 308 computed as follows.

$$309 \quad \Lambda_{3,46} = \{\lambda_1, \lambda_2\} \text{ where } \lambda_1 = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\} \text{ and } \lambda_2 = \{(\tau_1, \tau_3)\} \quad (13)$$

310 It is important to notice that the multiset union (\uplus) of all partitions in $\Lambda_{3,46}$ results in the
311 multiset of all possible preemptions, e.g.,

$$312 \quad \lambda_1 \uplus \lambda_2 = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\} \uplus \{(\tau_1, \tau_3)\} = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3), (\tau_1, \tau_3)\}$$

313 **► Proposition 4.** *Multiset $\Lambda_{i,t}$ consists of all possible preemptions that may occur within*
314 *duration t , between the jobs of the first i tasks of Γ .*

316 **Proof.** Directly follows from Proposition 2 and Equation 12 since Equation 12 includes each
317 possible preemption, occurable within duration t between the first i tasks of Γ , in one of the
318 partitions of $\Lambda_{i,t}$. ◀

319 5.3 CRPD bound on preemptions from a single partition

320 As suggested in Section 3, considering the *Problem 1* of CRPD over-approximation, computing
321 a bound for different preemption partitions individually, instead of computing it for all
322 preemptions at once, may result in more precise CRPD estimations.

323 To achieve this, once the multiset $\Lambda_{i,t}$ of preemption partitions is computed, an upper
324 bound on CRPD resulting from preemptions of a single partition $\lambda_r \in \Lambda_{i,t}$ can be computed
325 by selecting the minimum CRPD bound among the results from UCB-Union and ECB-Union
326 approaches. Here, we describe the improvements and adjustments on those approaches to
327 compute CRPD bound from preemptions contained within a partition.

328 In the following equations, $\text{aff}(i, h, \lambda_r)$ represents a set of tasks with priorities higher than
329 or equal to τ_i and lower than τ_h which can be preempted by τ_h according to preemptions
330 represented in λ_r . Formally: $\text{aff}(i, h, \lambda_r) = \{\tau_k \mid (\tau_h, \tau_k) \in \lambda_r \wedge \tau_k \in \text{hpe}(i)\}$. Also, with
331 $\text{hp}(i, \lambda_r)$ we denote a set of tasks with priorities higher than τ_i such that for each $\tau_h \in \text{hp}(i, \lambda_r)$
332 there is $(\tau_h, \tau_i) \in \lambda_r$.

333 First, we improve and adjust the ECB-Union approach, proposed by Altmeyer et al. [4].
334 In that approach, for a job of τ_h , it is accounted that it directly preempts one of the tasks
335 from $\text{aff}(i, h, \lambda_r)$ set such that the maximum possible number of UCBs are evicted. In order
336 for this approach to be correct, it is also accounted that tasks that can preempt τ_h also
337 contribute to the evictions of useful cache blocks of the preempted task. This scenario is
338 represented by a CRPD bound $\gamma_{i,h}^{\text{ecbp}}$. We further improve this formulation by accounting
339 that a preemption from a single job of τ_h on any job of τ_k from $\text{aff}(i, h, \lambda_r)$ cannot cause
340 more cache-block reloads than the maximum number of UCBs that can be evicted at a single
341 preemption point of τ_k . The maximum number of UCBs at a single preemption point of τ_k
342 is represented by $\text{ucb}_k^{\text{max}}$. The above translates to Equation 14.

$$343 \quad \gamma_{i,h}^{\text{ecbp}}(\lambda_r) = \max_{\tau_k \in \text{aff}(i,h,\lambda_r)} \left\{ \min \left(\left| \left(\bigcup_{\tau_{h'} \in \text{hp}(h,\lambda_r) \cup \{\tau_h\}} \text{ECB}_{h'} \right) \cap \text{UCB}_k \right|, \text{ucb}_k^{\text{max}} \right) \right\} \quad (14)$$

344 We build the correctness of the proposed computation on the correctness of the standard
345 ECB-Union method [4].

346 **► Proposition 5.** $\gamma_{i,h}^{\text{ecbp}}(\lambda_r)$ is an upper bound on number of reloads that may be imposed by
347 a direct preemption from τ_h on one of the tasks within $\text{aff}(i, h, \lambda_r)$ set.

348 **Proof.** A direct preemption from τ_h on one of the tasks within $\text{aff}(i, h, \lambda_r)$ set cannot cause
349 more reloads than the maximum number of UCBs of a preemptable task, which can be
350 evicted by τ_h and all the tasks that may preempt τ_h . Also, such bound cannot be greater
351 than the maximum number of useful cache blocks $\text{ucb}_k^{\text{max}}$ that may be present at a single
352 preemption point within a preemptable task, which concludes the proof. ◀

7:10 Improving the accuracy of cache-aware RTA using preemption partitioning

353 The ECB-Union based upper bound on all preemptions from λ_r is computed by summing all
 354 $\gamma_{i,h}^{ecbp}$ terms for each possibly preempting task, from τ_1 to τ_{i-1} :

$$355 \quad \gamma_i^{ecbp}(\lambda_r) = \sum_{h=1}^{i-1} \gamma_{i,h}^{ecbp}(\lambda_r) \quad (15)$$

356 **► Proposition 6.** $\gamma_i^{ecbp}(\lambda_r)$ is an upper bound on number of reloads that can be caused by
 358 preemptions from the partition λ_r .

359 **Proof.** For each direct preemption from the preempting jobs of tasks from τ_1 to τ_{i-1} in λ_r ,
 360 Equation 15 accounts that the upper-bounded number of cache-blocks is reloaded in one of
 361 the preemptable jobs, as follows from Proposition 5. Therefore, the proposition holds. ◀

362 Next, we adjust the UCB-Union approach, proposed by Tan and Mooney [30]. In this
 363 approach, for a job of τ_h , it is assumed that it can evict useful cache blocks from each task τ_k
 364 from the $\text{aff}(i, h, \lambda_r)$ set. However, since a single job of τ_h can directly or indirectly preempt
 365 each τ_k at only one of its preemption points, this cost can at most be equal to the sum of
 366 the number of maximum useful cache blocks at single preemption point of each task τ_k from
 367 $\text{aff}(i, h, \lambda_r)$. The above is formally represented with $\gamma_{i,h}^{ucbp}$ in the following equation:

$$368 \quad \gamma_{i,h}^{ucbp}(\lambda_r) = \min \left(\left| \left(\bigcup_{\tau_k \in \text{aff}(i,h,\lambda_r)} UCB_k \right) \cap ECB_h \right|, \sum_{\tau_k \in \text{aff}(i,h,\lambda_r)} \text{ucb}_k^{\max} \right) \quad (16)$$

369 We build the correctness of the proposed computation on the correctness of the standard
 370 UCB-Union method [30].

371 **► Proposition 7.** $\gamma_{i,h}^{ucbp}(\lambda_r)$ is an upper bound on number of reloads within all tasks from
 372 $\text{aff}(i, h, \lambda_r)$, that may be imposed because of the cache-block accesses from a single job of τ_h .

373 **Proof.** A job of τ_h cannot impose more than one cache block reload per cache-memory block
 374 m , such that $m \in ECB_h$, and $m \in UCB_k$ for any τ_k such that $\tau_k \in \text{aff}(i, h, \lambda_r)$, as follows
 375 from UCB-Union [30]. Also, since each task τ_k from $\text{aff}(i, h, \lambda_r)$ can be preempted by τ_h
 376 at only one of its preemption points, the maximum number of reloads from τ_h cannot be
 377 greater than the sum of the maximum numbers of useful cache blocks that may be present
 378 at a preemption point within each such task. This concludes the proof. ◀

379 The UCB-Union based upper bound on all preemptions from λ_r is also computed by
 380 summing all $\gamma_{i,h}^{ucbp}$ terms for each possibly preempting task from τ_1 to τ_{i-1} :

$$381 \quad \gamma_i^{ucbp}(\lambda_r) = \sum_{h=1}^{i-1} \gamma_{i,h}^{ucbp}(\lambda_r) \quad (17)$$

382 **► Proposition 8.** $\gamma_i^{ucbp}(\lambda_r)$ is an upper bound on number of reloads that can be caused by
 384 preemptions from the partition λ_r .

385 **Proof.** For each possibly preempting job from τ_1 to τ_{i-1} in λ_r , Equation 17 accounts that
 386 the job leads to upper-bounded number of cache-block reloads in its possibly preemptable
 387 jobs, as follows from Proposition 7. Therefore, the proposition holds. ◀

388 The final upper bound $\gamma_i(\lambda_r)$ on CRPD from possible preemptions given in λ_r , between
 389 single jobs of the first i tasks from Γ , is defined as the least bound of the two.

$$390 \quad \gamma_i(\lambda_r) = \min \left(\gamma_i^{ecbp}(\lambda_r), \gamma_i^{ucbp}(\lambda_r) \right) \times BRT \quad (18)$$

392 **► Proposition 9.** $\gamma_i(\lambda_r)$ is an upper bound on CRPD from possible preemptions given in the
 393 partition λ_r .

394 **Proof.** Follows from Propositions 6 and 8 since Equation 18 results in the least bound. ◀

5.4 CRPD bound on all preemptions within a time interval

Now, we define a computation for the CRPD bound on all preemptions which can occur within duration t , between the first i tasks of Γ .

► **Definition 10.** An upper bound $\gamma(i, t)$ on CRPD of all preemptions, which can occur within duration t between the first i tasks of Γ , is defined with the following equation:

$$\gamma(i, t) = \sum_{k=1}^{|\Lambda_{i,t}|} \gamma_i(\lambda_r) \quad (19)$$

► **Proposition 11.** $\gamma(i, t)$ is an upper bound on CRPD of all preemptions which can occur within duration t between the first i tasks of Γ .

Proof. Directly follows from Propositions 4 and 9, since Equation 19 is a sum of CRPD upper bounds of preemption partitions that jointly consist of all preemptions within t . ◀

5.5 Worst-case response time

In this subsection, we prove that the computed worst-case response time is an upper bound.

► **Theorem 12.** R_i is an upper bound on worst-case response time of τ_i .

Proof. By induction, over the tasks in Γ in a decreasing priority order.

Base case: $R_1 = C_1$, because $hp(i) = \emptyset$. Since C_i is the worst-case execution time of τ_1 the proposition holds.

Inductive hypothesis: Assume that for all $\tau_h \in hp(i)$, R_h is an upper bound on worst-case response time of τ_h .

Inductive step: We show that Equation 9 computes the worst-case response time of τ_i . Consider the least fixed point of Equation 9, for which $R_i = R_i^{(l)} = R_i^{(l+1)}$. At this point, the equation accounts for the following upper bounds and worst-case execution times:

- ◊ C_i , which is the worst-case execution time of τ_i , assumed by the system model.
- ◊ $\gamma(i, R_i)$, which is proved by Proposition 11 to be an upper bound on CRPD of all jobs which can be released within duration R_i , and have higher than or equal priority to τ_i .
- ◊ $\sum_{\forall \tau_h \in hp(i)} \lceil R_i / T_h \rceil C_h$, which is the worst-case interference caused by execution of all jobs of tasks with higher priority than τ_i without CRPD. Since we proved for all the factors which can prolong the response time of τ_i that they are accounted as the respective execution and CRPD upper bounds in Equation 9, then their sum results in an upper bound. ◀

5.6 Time complexity

The time complexity of the proposed analysis can be improved since in its current form Equation 12 explicitly creates all partitions which can lead to re-computation of CRPD bounds for many identical partitions. For this reason, we first define the matrix $A_{i,t}$, from which it is possible to identify how many repeated partitions there are, and compute CRPD bound for each distinct partition only once, as introduced in Algorithm 1.

► **Definition 13.** A matrix $A_{i,t}$ of upper bounds on number of preemptions between each pair of tasks with higher than or equal priority to P_i which can occur within duration of t , is defined with the following equation:

$$A_{i,t} = (a_{j,h}) \in \mathbb{N}^{i \times i} \mid a_{j,h} = \begin{cases} 0 & , j \leq h \\ E_j^h(t) & , j > h \end{cases} \quad (20)$$

7:12 Improving the accuracy of cache-aware RTA using preemption partitioning

433 ► **Proposition 14.** $A_{i,t}$ stores an upper-bounded number of preemptions, which can occur
 434 within t , between each pair of tasks with higher than or equal priority to P_i .

435 **Proof.** Proposition 14 follows directly from Proposition 2 and the fact that τ_j cannot preempt
 436 any task τ_h of higher priority, or τ_j itself ($j \leq h$). ◀

437 Equation 20 defines a square matrix $A_{i,t}$ such that the number of rows and columns is
 438 equal to i , and each entry of the matrix represents the maximum number of preemptions
 439 from a task τ_h on τ_j within duration t .

440 *Example:* Given the taskset from Figure 2, a matrix of preemptions during 46 time units
 441 looks as follows:

$$442 \quad A_{3,46} = \begin{array}{ccc|c} & \tau_1 & \tau_2 & \tau_3 \\ \begin{array}{c} 0 \\ 1 \\ 2 \end{array} & \begin{array}{c} 0 \\ 1 \\ 2 \end{array} & \begin{array}{c} 0 \\ 0 \\ 1 \end{array} & \begin{array}{c} \tau_1 \\ \tau_2 \\ \tau_3 \end{array} \end{array} \quad (21)$$

443

444 The element $a(2,1) = 1$ represents the maximum number of preemptions from τ_1 on τ_2
 445 during 46 time units (note $R_2 = 14$).

446 *Algorithm explanation:* In a matrix $A_{i,t}$, there are at most $\frac{n*(n-1)}{2}$ values representing
 447 different numbers of possible preemptions among the tasks. Therefore, there are at most
 448 $\frac{n*(n-1)}{2}$ distinct partitions to be generated. In Algorithm 1, we define the procedure that
 449 first generates $A_{i,t}$ (line 3), and then generates distinct partitions one by one (lines 4 – 10).
 450 For each distinct partition, we compute the number sp of times a partition is repeated, then
 451 compute the partition (line 6), and account for its CRPD bound sp times in the cumulative
 452 CRPD bound ξ that is updated for each distinct partition (lines 7 and 8). After this, the
 453 partitioned preemptions are removed (line 9), and the next distinct partition is computed
 454 until no more preemptions are left to be partitioned. Formally, termination criteria is satisfied
 455 when $A_{i,t}$ equals to the zero matrix $0_{i \times i}$. At the end, the algorithm results in the same
 456 CRPD bound as Equation 19. Using this algorithm, the time complexity is $\mathcal{O}(n^3 * x)$, where
 457 the complexity of computation at line 6 is $\mathcal{O}(x)$.

Data: Time duration t , task index i , Taskset Γ

Result: CRPD upper bound ξ on all jobs with priority higher than or equal to P_i ,
 which can be released within duration t .

```

1  fn  $\gamma(i, t)$ 
2  |    $\xi \leftarrow 0$ 
3  |    $A_{i,t} \leftarrow$  generate the matrix of maximum preemption counts between the tasks
   |   (Equation 20)
4  |   while  $A_{i,t} \neq 0_{i \times i}$  do
5  |   |    $sp \leftarrow$  minimum value from  $A_{i,t}$ , greater than zero
6  |   |    $\lambda \leftarrow \{(\tau_h, \tau_j) \mid sp \leq a_{j,h}\}$ 
7  |   |    $\gamma_i(\lambda) \leftarrow$  compute the CRPD upper bound from the preemptions in  $\lambda$ 
8  |   |    $\xi \leftarrow \xi + sp \times \gamma_i(\lambda)$ 
9  |   |    $A_{i,t} \leftarrow$  decrease all values, greater than zero, by  $sp$ 
10 |   end
11 |   return  $\xi$ 

```

■ **Algorithm 1** Algorithm for computing the cumulative CRPD during a time interval of length t .

5.7 CRPD computation using preemption scenarios

In this section, we propose an alternative computation for the upper bound $\gamma_i(\lambda)$ on CRPD of preemptions in the partition λ . The goal is to compute the CRPD bound from a single worst-case preemption combination among the preemptions from λ , addressing *Problem 2* from Section 4. To achieve this, we first formally define the following terms:

- ◇ *Preemption scenario* (τ_i, PT) , and its CRPD upper bound $\gamma(\tau_i, PT)$,
- ◇ *Preemption combination* Π_λ^c , and its CRPD upper bound $\gamma(\Pi_\lambda^c)$.

Informally, we define a preemption combination as a set of feasible preemptions where only one job of each task is involved, such that all accounted preemptions are present in a partition λ . Before being able to formally define a preemption combination, we first formally define a preemption scenario.

► **Definition 15** (Preemption scenario (τ_i, PT)). *A preemption scenario represents a single interruption due to preemption of a task and it is defined as an ordered pair (τ_i, PT) , where τ_i represents the preempted (interrupted) task, and PT is a set of preempting tasks which execute after the interruption at τ_i and before the immediate resumption of τ_i . Formally, for each preemption scenario (τ_i, PT) it holds that $PT \subseteq hp(i)$.*

Example: Given the example from Fig. 2, the first preemption scenario in τ_3 is $(\tau_3, \{\tau_1, \tau_2\})$, and the second preemption scenario is $(\tau_3, \{\tau_1\})$. Also, in the same figure, τ_2 is preempted once and this preemption scenario is $(\tau_2, \{\tau_1\})$.

In order to compute the upper bound on CRPD on τ_i , resulting from one interruption scenario, the ordering of the preempting tasks is not important. All of them are equally capable of evicting cache blocks of τ_i between its preemption and resumption, regardless of their ordering.

► **Definition 16** (CRPD of a preemption scenario $\gamma(\tau_i, PT)$). *An upper bound $\gamma(\tau_i, PT)$ on the CRPD of a preempted task τ_i resulting from a preemption scenario (τ_i, PT) is:*

$$\gamma(\tau_i, PT) = |UCB_i \cap \bigcup_{\tau_h \in PT} ECB_h| \times BRT \quad (22)$$

► **Proposition 17.** *$\gamma(\tau_i, PT)$ is an upper bound on the CRPD of a preempted task τ_i resulting from a preemption scenario (τ_i, PT) .*

Proof. Since Equation 22 accounts that each UCB from τ_i is definitely reloaded with the worst-case block reload time if there is a corresponding evicting cache block from any of the preempting tasks within a scenario, the proposition holds. ◀

Example: Given the preemption scenario $(\tau_3, \{\tau_1, \tau_2\})$, the upper bound on CRPD of τ_3 is:

$$\gamma(\tau_3, \{\tau_1, \tau_2\}) = |UCB_3 \cap \{ECB_1 \cup ECB_2\}| = 6$$

A safe upper bound on CRPD of τ_i resulting from a preemption scenario (τ_i, PT) can be tightened even more if the low-level task analysis can provide more detailed information on UCBs and ECBs at different program points. In such case, the bound is computed as the maximum intersection of UCBs from a single point within τ_i and the evicting cache blocks of tasks in PT . This is the case because Equation 22 considers a single preempted point and tasks which may evict cache blocks before preempted task resumes to execute. On the other hand, many existing approaches consider multiple preemption scenarios at once, and therefore this improvement is not applicable in their case, as shown by Shah et al. [27].

7:14 Improving the accuracy of cache-aware RTA using preemption partitioning

497 Given the formal definition of preemption scenarios, now we can define a preemption
 498 combination. With this definition, we need to insure that a preemption combination consists
 499 only of preemption scenarios which account for interactions between a single job of each task.
 500 Therefore, we need to insure that a preemption combination does not include two preemption
 501 scenarios which are mutually exclusive, given the constraint of using only single jobs.

502 ► **Definition 18** (Preemption combination Π^c). *A preemption combination Π^c is defined as a
 503 set of disjoint non-empty preemption scenarios between single jobs of tasks in Γ such that:*

504 1) *If there are two preemption scenarios (τ_j, PT^j) and (τ_l, PT^l) in Π^c such that $\tau_h \in$
 505 $PT^j \cap PT^l$ and $P_l < P_j$, it implies that $\tau_j \in PT^l$.*

506 2) *Each preempting task $\tau_h \in PT$, where $(\tau_j, PT) \in \Pi^c$, can be in at most one preemption
 507 scenario imposed on τ_j .*

508 **The first constraint** refers to a case: If a single job of task τ_h preempts single jobs of tasks
 509 τ_j and τ_l , where $P_j > P_l$, then that job of τ_l is definitely preempted by the τ_j job.

510 **The second constraint** accounts that an additional preemption scenario with τ_h implies
 511 that Π^c accounted for two jobs of τ_h preempting a job of τ_j , while the definition accounts for
 512 at most one job of each task.

513 *Example:* Given a definition of a preemption combination, one possible combination is:
 514 $\{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\}$ which describes the preemption scenario where τ_1 directly
 515 preempts τ_3 at one preemption point, while τ_2 directly preempts τ_3 at another preemption
 516 point. However, the set $\{(\tau_3, \{\tau_1\}), (\tau_2, \{\tau_1\})\}$ is not a preemption combination, because
 517 it describes the case where a job of τ_3 is preempted by a job of τ_1 , and a job of τ_2 is preempted
 518 by a job of τ_1 , while a job of τ_2 does not preempt a job of τ_3 . Since this is the case, more
 519 than two jobs of τ_1 are accounted, which violates Definition 18.

520 ► **Definition 19** (Preemption combination consistent with λ). *We say that Π^c is a preemption
 521 combination consistent with the preemption partition λ iff for any preemption scenario
 522 $(\tau_j, PT) \in \Pi^c$ the preemptions captured by the scenario are possible, i.e. present in λ .
 523 Formally: $\forall (\tau_j, PT) \in \Pi^c : \forall \tau_h \in PT : (\tau_j, \tau_h) \in \lambda$.*

524 *Example:* Given the preemption partition $\lambda = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\}$, a preemption
 525 combination consistent with λ is $\{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\}$ since it describes preemption
 526 scenarios made of preemptions that are possible, i.e. present in λ .

527 ► **Definition 20** (CRPD $\gamma(\Pi^c)$ of a preemption combination). *An upper bound $\gamma(\Pi^c)$ on the
 528 CRPD of a single preemption combination Π^c is defined as the sum of upper bounds of all
 529 preemption scenarios in Π^c . Formally, it is defined as:*

$$530 \quad \gamma(\Pi^c) = \sum_{(\tau_k, PT) \in \Pi^c} \gamma(\tau_k, PT) \quad (23)$$

531 ► **Proposition 21.** *$\gamma(\Pi^c)$ is an upper bound on CRPD of preemptions accounted within Π^c .*

532 **Proof.** A combination consists of a number of preemption scenarios representing the preemp-
 533 tions from preempting tasks on different preemption points. Following from Proposition 17,
 534 a sum of CRPD upper bounds of each task interruption in a combination is an upper bound
 535 on CRPD of all preemptions accounted in a combination, which concludes the proof. ◀

536 *Example:* Given the preemption combination of two direct preemptions from Figure 1,
 537 CRPD upper bound is: $\gamma(\{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\}) = \gamma(\tau_3, \{\tau_1\}) + \gamma(\tau_3, \{\tau_2\}) =$
 538 $4 + 4 = 8$. Given the preemption combination of a nested preemption from the figure, it is:
 539 $\gamma(\{(\tau_3, \{\tau_1, \tau_2\}), (\tau_2, \{\tau_1\})\}) = \gamma(\tau_3, \{\tau_1, \tau_2\}) + \gamma(\tau_2, \{\tau_1\}) = 6 + 2 = 8$.

540 Now, we can imagine a set which consists of all possible preemption combinations which
 541 are consistent with preemptions enlisted in λ . Then, among all the generated preemption
 542 combinations, we find one which results in the worst-case CRPD and declare that as a safe
 543 upper-bound, since that is the maximum obtainable CRPD value among all the possible
 544 preemption combinations.

545 However, to generate such complete set of all possible preemption combinations is
 546 computationally inefficient. A potential solution can come from the fact that it is enough to
 547 compute a subset of the complete set of combinations, as long as we are sure that no greater
 548 CRPD value can be obtained in the remaining, unaccounted combinations. We show this with
 549 the following *example*: Given a set of possible preemptions $\lambda = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\}$
 550 from Figure 1, let us consider the following set of two combinations:

$$551 \quad \Pi_{3,\lambda} = \left\{ \left\{ (\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\}) \right\}, \left\{ (\tau_3, \{\tau_1, \tau_2\}), (\tau_2, \{\tau_1\}) \right\} \right\}$$

552 $\Pi_{3,\lambda}$ consists of: 1) a combination of direct preemption scenarios, and 2) a combination of
 553 nested preemption. Any other possible preemption combination, from preemptions in λ , can
 554 only be derived by omitting at least one preemption from a preemption scenario from one of
 555 the two combinations given in $\Pi_{3,\lambda}$. E.g. preemption combination $\{(\tau_3, \{\tau_1, \tau_2\})\}$ is equal
 556 to and results in the same CRPD as $\{(\tau_3, \{\tau_1, \tau_2\}), (\tau_2, \emptyset)\}$. Also, all of the preemption
 557 scenarios from $\{(\tau_3, \{\tau_1, \tau_2\})\}$ are already included in the second combination. Therefore,
 558 all the other possible combinations cannot result in a greater CRPD value than those in
 559 $\Pi_{3,\lambda}$, meaning that this subset is sufficient to compute a safe CRPD.

560 A preemption combination which is constructed by adding a preemption scenario to
 561 any of the combinations in $\Pi_{3,\lambda}$ cannot be obtained. This is the case because in the first
 562 combination, it is accounted that τ_3 is preempted by both tasks, at two different points,
 563 accounting for all preemptions in λ where τ_3 can be preempted, i.e. (τ_1, τ_3) and (τ_2, τ_3) . In
 564 this case, τ_2 cannot be preempted considering preemption $(\tau_1, \tau_2) \in \lambda$ as shown in Definition
 565 18. In the second combination, τ_3 is interrupted once while both tasks preempt it, and τ_2 is
 566 preempted by τ_1 , meaning that all preemptions from λ are accounted.

567 In order to define a safe set of combinations $\Pi_{i,\lambda}$, such that at least one of those combin-
 568 ations may result in the worst-case CRPD, we first introduce a term of **set partitioning**¹
 569 in order to represent different ways one task may be preempted by the others. *Example*:
 570 Given the task τ_3 , set partitions of a set $\{\tau_1, \tau_2\}$ of its potentially preempting tasks are: 1)
 571 $\{\{\tau_1, \tau_2\}\}$, and 2) $\{\{\tau_1\}, \{\tau_2\}\}$.

572 Given a preemptable task τ_k , and a set of its possibly preempting tasks PT , all set
 573 partitions of PT represent all the ways τ_k may be preempted such that each task from
 574 PT preempts τ_k . This is the case because set partitions represent all the ways a set can
 575 be grouped in non-empty subsets, such that each set element is included in exactly one
 576 subset. Analogically, in this paper, each set partition is transformed into a preemption
 577 combination defining one way how a task (e.g. τ_3 above) can be preempted. Each set
 578 partition consists of subsets, and each subset represents a preemption scenario on the
 579 preempted task. This transformation is formally defined in function *generateCombs*(PT, τ_k)
 580 in Algorithm 2. *Example*: Considering different ways τ_3 can be preempted, set partition
 581 $\{\{\tau_1, \tau_2\}\}$ consists of a single subset, and forms a preemption combination $\{(\tau_3, \{\tau_1, \tau_2\})\}$,
 582 while set partition $\{\{\tau_1\}, \{\tau_2\}\}$ consists of two subsets and forms a preemption combination
 583 $\{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\}$ with two preemption scenarios: $(\tau_3, \{\tau_1\})$, and $(\tau_3, \{\tau_2\})$.

¹ Set partitioning is a mathematical concept sometimes also known as Bell partitioning [1, 18] named after Eric Temple Bell. There are many fast algorithms for generating set partitions, e.g. [13, 14].

584 ► **Proposition 22.** *Given a preemptable task τ_k , and a set of possibly preempting tasks PT ,
585 any combination of preemptions on τ_k will result in a less than or equal CRPD than any
586 combination generated from $generateCombs(PT, \tau_k)$.*

587 **Proof.** By contradiction: Let us assume that there is a preemption combination Π_k^c represent-
588 ing the ways how τ_k can be preempted by tasks from PT , and that Π_k^c can result in a greater
589 CRPD than any combination derived from $generateCombs(PT, \tau_k)$. The combinations gener-
590 ated from $generateCombs(PT, \tau_k)$ represent all the ways τ_k may be preempted such that each
591 task from PT preempts τ_k since set partitions represent all the ways a set can be grouped in
592 non-empty subsets, such that each element is included in exactly one subset. Thus, Π_k^c must
593 omit at least one preemption, compared to at least one preemption combination derived
594 from $generateCombs(PT, \tau_k)$. The initial assumption therefore contradicts Proposition 21
595 because Π_k^c cannot impose larger CRPD than the corresponding preemption combination
596 from $generateCombs(PT, \tau_k)$, which accounts for the same preemptions as in preemption
597 scenarios from Π_k^c and at least one additional preemption compared to Π_k^c . ◀

598 Using the concept of set partitioning to represent the ways a single task may be preempted,
599 we generate a set $\Pi_{i,\lambda}$ of preemption combinations on how all tasks from τ_i to τ_1 can interact
600 among each other:

601 ► **Definition 23.** *By $\Pi_{i,\lambda}$ we denote the result from Algorithm 2, i.e. the set of preemption
602 combinations between the first i tasks of Γ such that each combination is consistent with λ .*

603 We describe Algorithm 2 in more details and we use a running example from Figure 2
604 to show the algorithm walk-through in Figure 3. As stated before, for each τ_k , from τ_i to
605 τ_1 , the algorithm first generates possible combinations on how τ_k can be preempted, using
606 set partitioning (line 3). This process is defined in function $generateCombs$ (line 7) and it
607 translates the set partitions of possibly preempting tasks on τ_k , into different ways τ_k can be
608 preempted, which is represented with a set of preemption combinations (line 16). Then, for
609 each of those combinations, the algorithm performs $extendCombs$ (line 4), which is a function
610 that updates the existing preemption combinations, with further preemption scenarios that
611 are possible on the preempting tasks of τ_k . Take for example the preemption combination
612 Π^c , given in Figure 3.

613 The combination represents the case where τ_4 is preempted at one preemption point, by
614 all of its three possibly preempting tasks (τ_1, τ_2, τ_3). In the figure, this is represented by one
615 arrow (standing for one preempted point of τ_4) and tasks preempting a point (above the
616 arrow). Function $extendCombs()$ further computes possible ways of preempting τ_3 since τ_3 is
617 the lowest-priority preempting task from the preemption scenario ($\tau_4, \{\tau_1, \tau_2, \tau_3\}$). Those
618 ways are represented with a set Π_3^c of preemption combinations. After this, the function
619 updates the preemption combinations with a Cartesian product of the two. Therefore, on the
620 right side of the figure, you may notice that now we have two new combinations, updating the
621 Π_c with different ways τ_3 can be preempted. The topmost combination can be updated further
622 on, since preemption scenario ($\tau_3, \{\tau_1, \tau_2\}$) can be updated with additional scenario on how τ_2
623 can be preempted by τ_1 . This is eventually computed within the algorithm because condition
624 in line 18 insures that all combinations are updated until no new preemption scenario can be
625 added to any of the existing preemption combinations. More formally, this criteria is satisfied
626 when for each preemption scenario (τ_x, PT) within any preemption combination from $\Pi_{i,\lambda}$,
627 function $extended?((\tau_x, PT))$ yields true (\top), meaning that all preemption scenarios are
628 extended.

629 ► **Proposition 24.** $\Pi_{i,\lambda}$ is a safe set of preemption combinations between the single jobs of
 630 the first i tasks in Γ , i.e. there is no preemption combination consistent with λ with a higher
 631 CRPD than the maximum CRPD of the combinations in $\Pi_{i,\lambda}$,

Data: Set of possible preemption pairs λ ,
 task index i
Result: A set $\Pi_{i,\lambda}$ of preemption
 combinations consistent with λ

```

1  $\Pi_{i,\lambda} \leftarrow \emptyset$ 
2 for  $k \leftarrow i$  to 2 by  $-1$  do
3    $\Pi'_{k,\lambda} \leftarrow \text{generateCombs}(hp(k), \tau_k)$ 
4    $\Pi_{i,\lambda} \leftarrow \Pi_{i,\lambda} \cup \text{extendCombs}(\Pi'_{k,\lambda})$ 
5 end
6 return  $\Pi_{i,\lambda}$ 

7 fn  $\text{generateCombs}(PT, \tau_k)$ 
8    $\Pi_{k,\lambda} \leftarrow \emptyset$ 
9    $PT_{k,\lambda} \leftarrow$  remove those tasks from  $PT$ 
    that cannot preempt  $\tau_k$  according to  $\lambda$ 
10   $\text{partitions}(PT) \leftarrow$  generate all possible
    partitions of a set  $PT_{k,\lambda}$ , representing
    ways a job of  $\tau_k$  can be preempted.
11  for each  $\text{partition} \in \text{partitions}(PT)$ 
12     $\Pi_k^c \leftarrow \emptyset$ 
13    for each  $\text{subset} \in \text{partition}$ 
14       $\Pi_k^c \leftarrow \Pi_k^c \cup \{(\tau_k, \text{subset})\}$ 
15     $\Pi_{k,\lambda} \leftarrow \Pi_{k,\lambda} \cup \Pi_k^c$ 
16  return  $\Pi_{k,\lambda}$ 

17 fn  $\text{extendCombs}(\Pi_{q,\lambda})$ 
18  while  $\exists \Pi^c \in \Pi_{q,\lambda} : \exists (\tau_r, PT) \in \Pi^c \mid$ 
     $\text{extended}((\tau_r, PT), \Pi^c) = \perp$  do
19     $\tau_l \leftarrow$  lowest-priority task in  $PT$ 
20     $\Pi'_l \leftarrow \text{generateCombs}(PT \setminus \tau_l, \tau_l)$ 
21    for each  $\Pi^c \in \Pi_{q,\lambda} \mid (\tau_r, PT) \in \Pi^c$ 
22       $\Pi_{q,\lambda} \leftarrow \Pi_{q,\lambda} \cup (\Pi^c \times \Pi'_l)$ 
23       $\Pi_{q,\lambda} \leftarrow \Pi_{q,\lambda} \setminus \Pi^c$ 
24  end
25  return  $\Pi_{q,\lambda}$ 

26 fn  $\text{extended}((\tau_r, PT), \Pi^c)$ 
27   $\tau_l \leftarrow$  lowest-priority task in  $PT$ 
28  if  $\exists (\tau_x, PT') \in \Pi^c \mid \tau_x = \tau_l$  then
29    return  $\perp$ 
30  else return  $\perp$ ;
```

■ **Algorithm 2** Algorithm that generates a
 set $\Pi_{i,\lambda}$ of preemption combinations.

633 **Proof.** By contradiction: Let us assume that there is a preemption combination Π^c_o between
 634 the single jobs of the first i tasks, consistent with λ , which can result in a higher CRPD than
 635 any of the combinations in $\Pi_{i,\lambda}$. For each task τ_k such that $(1 < k \leq i)$, it is accounted that
 636 τ_k experiences the worst-case CRPD at one of the generated combinations, as follows from
 637 Proposition 22 and line 3. Each such a combination is extended in line 4, accounting for
 638 further worst-case preemption scenarios on how all the preempting tasks can be preempted,
 639 and Algorithm 2 stops only when no more preemption scenarios can be generated and added
 640 to a set of preemption combinations $\Pi_{i,\lambda}$. This further implies that Π^c_o must omit at least
 641 one preemption from at least one of its preemption scenarios compared to any combination
 642 from $\Pi_{i,\lambda}$. Moreover, by construction of Algorithm 2, there is a preemption combination
 643 Π^c_w in $\Pi_{i,\lambda}$ which is a superset over the Π^c_o , i.e. there is a mapping of preemption scenarios
 644 between Π^c_w and Π^c_o such that each preemption scenario of Π^c_w includes same preemptions as

Data: $\lambda = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\}$, $i = 3$

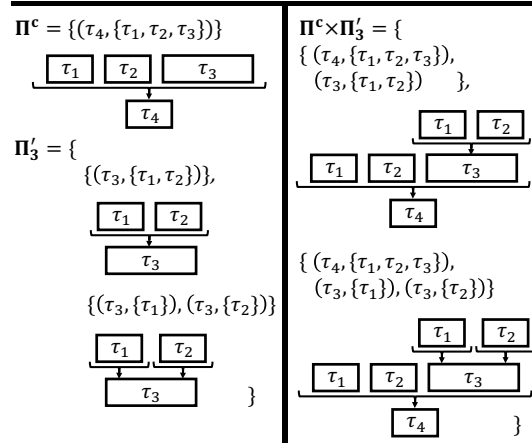
Algorithm run:

```

 $\Pi_{3,\lambda} \leftarrow \emptyset$ 
for  $k = 3$ 
   $\Pi'_{3,\lambda} \leftarrow \text{generateCombs}(hp(3), \tau_3)$ 
   $\leftarrow \{ \{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\},$ 
     $\{(\tau_3, \{\tau_1, \tau_2\})\} \}$ 
   $\Pi_{3,\lambda} \leftarrow \emptyset \cup \text{extendCombs}(\Pi'_{3,\lambda})$ 
   $\leftarrow \{ \{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\},$ 
     $\{(\tau_3, \{\tau_1, \tau_2\}), (\tau_2, \{\tau_1\})\} \}$ 

for  $k = 2$ 
   $\Pi'_{2,\lambda} \leftarrow \text{generateCombs}(hp(2), \tau_2)$ 
   $\leftarrow \{ \{(\tau_2, \{\tau_1\})\} \}$ 
   $\Pi_{3,\lambda} \leftarrow \{ \{(\tau_3, \{\tau_1\}), (\tau_3, \{\tau_2\})\},$ 
     $\{(\tau_3, \{\tau_1, \tau_2\}), (\tau_2, \{\tau_1\})\} \}$ 
     $\cup \{ \{(\tau_2, \{\tau_1\})\} \}$ 

return  $\Pi_{3,\lambda}$ 
```



■ **Figure 3** Top: Algorithm walkthrough with an example from Figure 2. Bottom: Example for extending the combination Π^c of four tasks.

645 the respective scenario in Π_o^c , but may also include additional ones not accounted by Π_o^c . As
 646 follows from Propositions 21 and 22, Π_o^c can only result in CRPD less than or equal to the
 647 one from Π_w^c . This contradicts the initial assumption. ◀

648 Finally, we can compute an upper bound on CRPD resulting from the worst-case pre-
 649 emptation combination consisting of the preemptions in λ , with the following equation:

$$650 \quad \gamma_i(\lambda) = \max_{\Pi^c \in \Pi_{i,\lambda}} \gamma(\Pi^c) \quad (24)$$

652 ▶ **Proposition 25.** $\gamma_i(\lambda)$ is an upper bound on CRPD from preemptions given in the partition
 653 λ , between the single jobs of the first i tasks from Γ .
 654

655 **Proof.** Equation 24 computes the maximum upper bound from all preemption combinations
 656 accounted by $\Pi_{i,\lambda}$. Then, following from Propositions 21 and 24, the proposition holds. ◀

657 *Example:* Given a set of possible preemptions $\lambda = \{(\tau_1, \tau_2), (\tau_1, \tau_3), (\tau_2, \tau_3)\}$ from Figure 1
 658 and continuing from the example after Proposition 21, the upper bound on CRPD resulting
 659 from preemptions in λ is computed as $\gamma(\lambda) = \max(\{8, 8\}) = 8$.

660 5.8 Adjustment for LRU caches

661 The proposed methods can also be used for set-associative LRU caches with a single modifi-
 662 cation, as shown by Altmeyer et al. [4, 6]. In case of LRU set-associative cache, a cache-set
 663 may contain several useful cache blocks, e.g., $UCB_2 = \{1, 2, 2, 2\}$ means that τ_2 contains
 664 three cache blocks in cache-set 2, and one UCB in cache set 1. Upon preemption, one ECB
 665 of the pre-empting task may suffice to evict all UCBs of the same cache-set, meaning that
 666 $ECB_1 = \{1, 2\}$ may evict all cache blocks of τ_2 . Therefore, the current notion of the ECBs
 667 and UCBs of a task may remain unchanged if a bound on CRPD due to preemption from τ_h
 668 on τ_i is defined as: $UCB_i \cap' ECB_h$ where the result is a multiset that contains each element
 669 from UCB_i if it is also in ECB_h , e.g. $UCB_2 \cap' ECB_1 = \{1, 2, 2, 2\} \cap' \{1, 2\} = \{1, 2, 2, 2\}$. In
 670 case of FIFO and PLRU cache replacement policies, the concepts of useful and evicting cache
 671 blocks cannot be applied, as shown by Burguiere et al. [10].

672 6 Evaluation

673 In this section, we show the evaluation results. The goal of the evaluation was to investigate
 674 to what extent the proposed method is able to identify schedulable tasksets upon the analysis
 675 of the cache-related preemption delays. We compared the state-of-the-art analyses for CRPD:
 676 (*ECB-Union Multiset*), (*UCB-Union Multiset*) methods, and (*Combined Multiset*), with two
 677 versions of the proposed method, i.e. (*Partitioning-ver1*) which computes CRPD according to
 678 Section 5.3, and the version (*Partitioning-ver2*) which computes CRPD from the worst-case
 679 preemption combination, presented in Section 5.7.

680 As shown by Shah et al. [27], an evaluation of the CRPD-aware methods should consider
 681 task parameters derived by using the existing low-level analysis tools. Therefore, in this paper
 682 we use the suggested task parameters that are derived with LLVMTA analysis tool [17], used
 683 on Mälardalen [16] and TACLe [15] benchmark tasks. The derived task characteristics are
 684 shown in Table 1, and they are: worst-case execution time, expressed in terms of wall-clock
 685 time, set of evicting cache blocks, set of definitely useful cache blocks (shown in the table as
 686 the size of the respective sets – ECB and DC-UCB), and the maximum number (Max DC-
 687 UCB) of definitely useful cache blocks per any program point of a task. The characteristics

Task (TACLe Bench.)	WCET	ECB	UCB	Max	...continuation	WCET	ECB	UCB	Max
app/lift	13592762	250	125	23	sequential/petrinet	39951	256	92	2
app/powerwindow	55842069	256	120	25	sequential/ri_dec	1811372648	256	173	44
kernel/binarysearch	2860	43	19	18	sequential/ri_enc	39467989	256	181	44
kernel/bsort	3332496	42	30	29	sequential/statemate	1949343	256	91	1
kernel/complex_update	8190	36	28	27	sequential/susan	2051176771	256	255	79
kernel/countnegative	260303	78	45	45					
kernel/fft	493123975	103	87	52	Task (Mälardalen Bench.)				
kernel/filterbank	38302875	164	151	66	adpcm	82492494	256	230	103
kernel/fir2dim	86737	212	197	116	bs	3052	43	23	20
kernel/iir	3307	41	32	31	bsort100	3146185	57	40	30
kernel/insertsort	16148	50	35	28	cut	127558	123	58	44
kernel/jfdctint	9043	115	107	54	compress	1090099	247	150	63
kernel/lms	1758977	82	56	23	cover	74509	256	38	15
kernel/ludcmp	97908	173	137	44	crc	1376054	121	62	30
kernel/matrix1	248058	48	43	42	edn	739866	256	222	123
kernel/md5	367421931	256	149	72	expint	2161270	117	47	29
kernel/minver	67700	254	173	46	fdct	10258	126	113	62
kernel/pm	141189221	256	247	45	fft1	271733	222	154	63
kernel/prime	386343	80	54	41	fibcall	8406	28	16	16
kernel/sha	28380272	253	185	31	fir	12413071	94	42	21
kernel/st	1763900	161	80	43	insertsort	11291	29	16	15
sequential/adpcm_dec	52530	233	145	59	jamne_complex	33778	39	28	27
sequential/adpcm_enc	58861	236	158	75	jfdctint	21742	132	122	54
sequential/andibeam	6434692	256	212	46	lcdnum	6100	51	11	9
sequential/cjpeg_transupp	535718162	256	256	103	lms	10178805	242	134	38
sequential/cjpeg_wrbmp	1610145	138	80	38	ludcmp	116312	210	168	44
sequential/dijkstra	39781181581	151	80	46	matmult	1447379	85	51	31
sequential/epic	7423276281	256	256	107	minver	67157	256	178	47
sequential/g723_enc	22919200	256	154	81	ndes	1050163	253	176	38
sequential/gsm_dec	3744323	256	236	69	ns	126865	55	37	34
sequential/gsm_encode	2115350	256	256	118	nsichneu	201969	256	183	2
sequential/h264_dec	24979237	256	166	29	prime	7782800	75	47	33
sequential/huff_dec	9360435	254	144	44	qsort	163089	142	83	39
sequential/mpeg2	130756234186	256	256	154	qurt	71655	130	40	26
sequential/ndes	996427	253	167	39	select	6306	159	73	55
					sqrt	22436	53	21	12
					st	3701746	192	95	52
					statemate	41579	256	105	1
					ud	355318	194	151	39

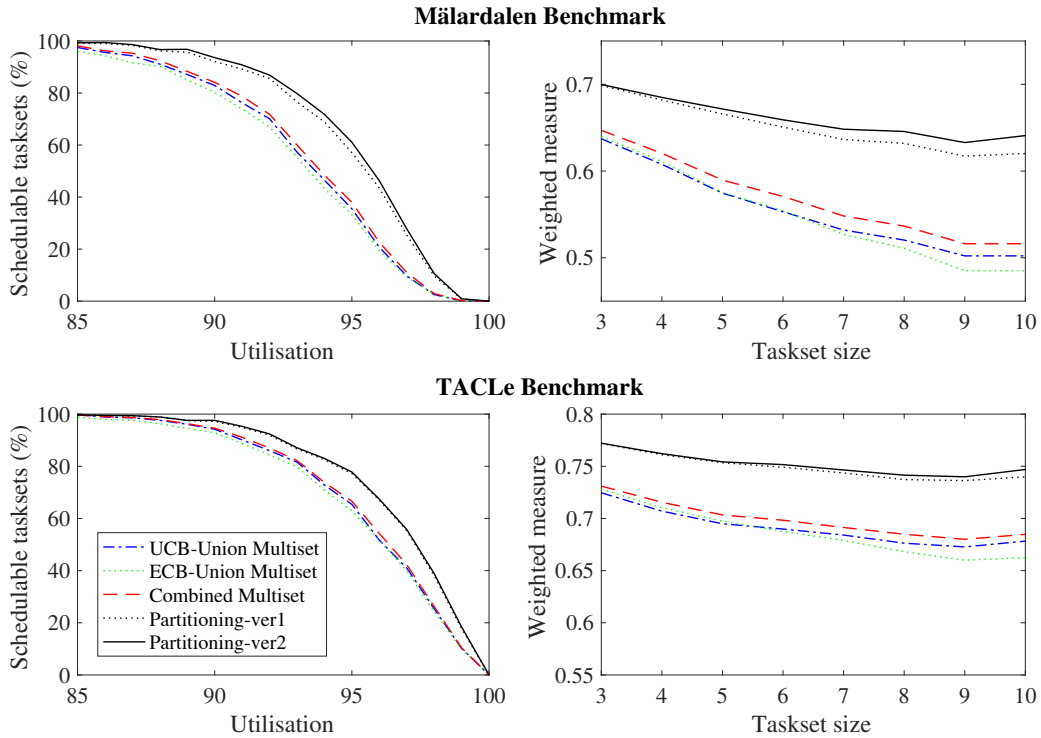
■ **Table 1** Task characteristics obtained with LLVMTA [17] analysis tool used on Mälardalen [16] and TACLe [15] benchmark tasks.

688 were derived with assumed direct-mapped instruction cache and a data scratchpad. The
689 assumed cache memory consists of 256 sets with line size equal to 8 bytes, while block reload
690 time is equal to 22 cycles. For more details about the low-level analysis refer to [27].

691 Tasksets are generated by randomly selecting a subset of tasks from one of the two
692 benchmarks, Mälardalen or TACLe, specified in each figure. We generated 1000 tasksets
693 for each pair of selected utilisation and taskset size. Since the task binaries were analysed
694 individually, they all start at the same address (mapping to cache set 0). In a multi-task
695 scheduling situation this can hardly be a case because the ECB and UCB placement is
696 determined by their respective locations in memory. We took this into account by randomly
697 shifting the cache set indices, e.g. the ECB in cache set i is shifted to the cache line equal to
698 $(i + \text{random}(256)) \bmod 256$. Task utilisations were generated using U-Unifast algorithm, as
699 proposed by Bini et al. [8]. Minimum inter-arrival times were then computed using equation
700 $T_i = C_i/U_i$, while the deadlines are assumed to be implicit, i.e. $D_i = T_i$. Task priorities
701 were assigned using deadline-monotonic order.

702 In Figure 4, on the leftmost plot, we show the schedulability results of an experiment
703 where we generated tasksets of size 9, from the Mälardalen tasks (top left), and TACLe tasks
704 (bottom left). For each generated taskset, its utilisation was varied from 0.5 to 1, by step
705 of 0.01. The results show that *Partitioning-ver2* and *Partitioning-ver1* identify the highest
706 number of schedulable tasksets, even up to 23% more for *Partitioning-ver2*, and 20% more
707 for *Partitioning-ver1*, compared to *Combined multisets*.

708 To increase the exhaustiveness of the performed evaluation and the respective results, for
709 the rightmost plots from Figure 4 we used the weighted schedulability measure in order to
710 show a 2-dimensional plot which would otherwise be a 3-dimensional plot, as proposed by
711 Bastoni et al. [7]. In those figures, we show the weighted schedulability measure $W_y(|\Gamma|)$, for
712 schedulability test y as a function of taskset size $|\Gamma|$. For each taskset size (in range from
713 3 to 10), this measure combines data for all of the tasksets generated for each utilisation
714 level from 0.85 to 1, with step of 0.1, since for utilisation levels from 0 to 0.85 all of the
715 compared methods deem almost all tasksets to be schedulable. For each taskset size $|\Gamma|$, the
716 schedulability measure $W_y(|\Gamma|)$ is equal to $W_y(|\Gamma|) = \sum_{\forall \Gamma} (U_\Gamma \times B_y(\Gamma, |\Gamma|)) / \sum_{\forall \Gamma} U_\Gamma$, where



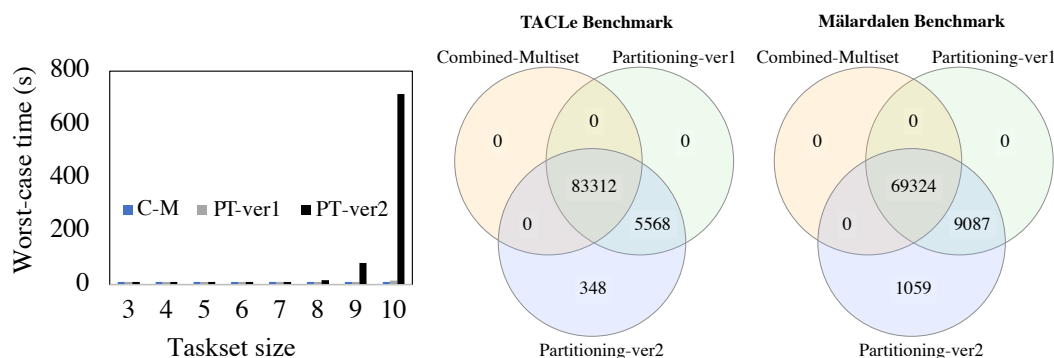
■ **Figure 4** *Left*: Schedulability ratio at different taskset utilisation. *Right*: Weighted measure at different taskset size.

717 $B_y(\Gamma, |\Gamma|)$ is the binary result (1 if schedulable, 0 otherwise) of a schedulability test y for
 718 a taskset Γ and taskset size $|\Gamma|$. Weighting the schedulability results by taskset utilisation
 719 means that the method which succeeds to produce a higher weighted measure, compared to
 720 the others, is more prone to identify tasksets with higher utilisation as schedulable.

721 The results show that *Partitioning-ver2* is able to identify more schedulable tasksets
 722 compared to the others for any given taskset size, immediately followed by *Partitioning-ver1*.
 723 Also, as the taskset size increases, the multiset-based methods deteriorate more in identifying
 724 schedulable tasksets compared to the proposed methods. This means that partitioning-based
 725 methods are able to identify more tasksets as schedulable with an increase of the taskset size
 726 and utilisation.

727 Next, we report the worst-case computation time results since *Partitioning-ver2* uses set
 728 partitioning which is known to be a computation with quadratic/exponential complexity,
 729 depending on the algorithm type. The results, reported in the left-most plot in Figure 5,
 730 were computed on MacBook Pro (Retina, 13-inch, Early 2015) version, with Intel Core
 731 i5 processor of 2,9 GHz, and DDR3 RAM memory of 8 GB, and 1867 MHz. We used a
 732 sequential set partitioning algorithm, and as shown in the graph, in this case exponential
 733 complexity is evident for *Partitioning-ver2*. However, the proposed method is intended to be
 734 used offline, and its performance can be improved using the algorithm from [14], and even
 735 more with parallel computing, e.g. set partitioning algorithms proposed by Djokic et al. [13].
 736 In contrast, *Partitioning-ver1* has a low worst-case time measured for each experiment for
 737 different taskset sizes, similar to the the *Combined-multiset* approach.

738 Finally, we show the relations between the results in Figure 5 (central and rightmost fig-
 739 ures). In the central figure, it is evident that all tasksets from TACLe benchmark, that are iden-
 740 tified as schedulable by *Combined-multiset*, are also identified as schedulable by *Partitioning-*



■ **Figure 5** *Leftmost:* The worst-case measured analysis time per taskset, at different taskset size. *Center and rightmost:* Venn Diagrams[19] representing schedulability result relations between different methods, over 120000 analysed tasksets per each – TACLe Benchmark and Mälardalen Benchmark.

741 *ver1* and *Partitioning-ver2*. However, partitioning-based approaches identify 5568 (and
 742 9087) additional schedulable tasksets depending on the benchmark, while *Partitioning-ver2*
 743 identifies additional 348 (and 1059) schedulable tasksets compared to *Partitioning-ver1*. In
 744 conclusion of the evaluation, we notice that the proposed partitioning-based algorithms
 745 outperform existing state of the art *Combined-Multiset* approach. Also *Partitioning-ver2*
 746 outperforms *Partitioning-ver1*, however this comes with the expense of time complexity.
 747 The complexity of *Partitioning-ver2* can be further decreased by narrowing down the task
 748 interactions for which the preemption combinations should be generated. This remains as
 749 a part of the future work as well as the formal proof of the dominance relations between
 750 the methods. Finally, the proposed approaches allow for a hybrid, joint use of the two
 751 proposed algorithms, while *Partitioning-ver1* significantly outperforms the existing multiset
 752 approaches without the expense of time complexity.

753 7 Conclusions

754 In this paper, we proposed a partitioning based cache-aware schedulability analysis for precise
 755 and safe estimation of cache-related preemption delays in the context of fully-preemptive
 756 scheduling of real-time systems with sporadic tasks with fixed priorities. The proposed
 757 methods are based on a precise analysis of: 1) different preemption subgroups, and 2)
 758 different preemption combinations that may occur within a system, and therefore they are
 759 able to compute more precise cache-related preemption delay estimations compared to the
 760 state of the art approaches. The evaluation was performed using the realistic task parameters
 761 from well-established benchmarks, obtained with a low-level analysis tool, and it showed
 762 that the proposed approaches manage to identify significantly more schedulable tasksets
 763 compared to the other preemption-cost aware approaches.

764 In future work, we will apply the proposed method in the context of limited preemptive
 765 scheduling since for such task model partitioning-based consideration of preemptions can
 766 lead to a more precise computation of cache-related preemption delay. We will also apply
 767 the proposed methods to other cache architectures and replacement protocols since many
 768 existing analyses inherit the overly pessimistic estimations which are identified in this paper.
 769 Finally, we will define a more precise static analysis on number of cache block reloads that
 770 are possible during the execution of a task since the existing useful cache block concept
 771 significantly over-approximates cache-block reloadability.

772 — References

- 773 1 Martin Aigner. A characterization of the bell numbers. *Discrete mathematics*, 205(1-3):207–210,
774 1999.
- 775 2 Sebastian Altmeyer and Claire Burguiere. A new notion of useful cache block to improve
776 the bounds of cache-related preemption delay. In *Real-Time Systems, 2009. ECRTS'09. 21st*
777 *Euromicro Conference on*, pages 109–118. IEEE, 2009.
- 778 3 Sebastian Altmeyer, Robert I Davis, and Claire Maiza. Cache related pre-emption delay aware
779 response time analysis for fixed priority pre-emptive systems. In *2011 IEEE 32nd Real-Time*
780 *Systems Symposium*, pages 261–271. IEEE, 2011.
- 781 4 Sebastian Altmeyer, Robert I Davis, and Claire Maiza. Improved cache related pre-emption
782 delay aware response time analysis for fixed priority pre-emptive systems. *Real-Time Systems*,
783 48(5):499–526, 2012.
- 784 5 Sebastian Altmeyer, Roeland Douma, Will Lunniss, and Robert I Davis. On the effectiveness
785 of cache partitioning in hard real-time systems. *Real-Time Systems*, 52(5):598–643, 2016.
- 786 6 Sebastian Altmeyer, Claire Maiza, and Jan Reineke. Resilience analysis: tightening the CRPD
787 bound for set-associative caches. In *ACM Sigplan Notices*, volume 45, pages 153–162. ACM,
788 2010.
- 789 7 Andrea Bastoni, Björn Brandenburg, and James Anderson. Cache-related preemption and
790 migration delays: Empirical approximation and impact on schedulability. *Proceedings of*
791 *OSPERT*, pages 33–44, 2010.
- 792 8 Enrico Bini and Giorgio C Buttazzo. Measuring the performance of schedulability tests.
793 *Real-Time Systems*, 30(1-2):129–154, 2005.
- 794 9 Tobias Blaß, Sebastian Hahn, and Jan Reineke. Write-back caches in wcet analysis. In
795 *29th Euromicro Conference on Real-Time Systems (ECRTS 2017)*. Schloss Dagstuhl-Leibniz-
796 Zentrum fuer Informatik, 2017.
- 797 10 Claire Burguière, Jan Reineke, and Sebastian Altmeyer. Cache-related preemption delay
798 computation for set-associative caches-pitfalls and solutions. In *9th International Workshop*
799 *on Worst-Case Execution Time Analysis (WCET'09)*. Schloss Dagstuhl-Leibniz-Zentrum für
800 Informatik, 2009.
- 801 11 José V Busquets-Mataix, Juan José Serrano, Rafael Ors, Pedro Gil, and Andy Wellings.
802 Adding instruction cache effect to schedulability analysis of preemptive real-time systems. In
803 *Real-Time Technology and Applications Symposium, 1996. Proceedings., 1996 IEEE*, pages
804 204–212. IEEE, 1996.
- 805 12 Robert I Davis, Sebastian Altmeyer, and Jan Reineke. Response-time analysis for fixed-priority
806 systems with a write-back cache. *Real-Time Systems*, 54(4):912–963, 2018.
- 807 13 Borivoje Djokić, Masahiro Miyakawa, Satoshi Sekiguchi, Ichiro Semba, and Ivan Stojmenović.
808 Parallel algorithms for generating subsets and set partitions. In *International Symposium on*
809 *Algorithms*, pages 76–85. Springer, 1990.
- 810 14 MC Er. A fast algorithm for generating set partitions. *The Computer Journal*, 31(3):283–284,
811 1988.
- 812 15 Heiko Falk, Sebastian Altmeyer, Peter Hellinckx, Björn Lisper, Wolfgang Puffitsch, Christine
813 Rochange, Martin Schoeberl, Rasmus Bo Sørensen, Peter Wägemann, and Simon Wegener.
814 Taclebench: A benchmark collection to support worst-case execution time research. In *16th*
815 *International Workshop on Worst-Case Execution Time Analysis (WCET 2016)*. Schloss
816 Dagstuhl-Leibniz-Zentrum für Informatik, 2016.
- 817 16 Jan Gustafsson, Adam Betts, Andreas Ermedahl, and Björn Lisper. The mälardalen wcet
818 benchmarks: Past, present and future. In *10th International Workshop on Worst-Case*
819 *Execution Time Analysis (WCET 2010)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik,
820 2010.
- 821 17 Sebastian Hahn, Michael Jacobs, and Jan Reineke. Enabling compositionality for multicore
822 timing analysis. In *Proceedings of the 24th international conference on real-time networks and*
823 *systems*, pages 299–308. ACM, 2016.

- 824 18 Paul R Halmos. *Naïve set theory*. Courier Dover Publications, 2017.
- 825 19 Henry Heberle, Gabriela Vaz Meirelles, Felipe R da Silva, Guilherme P Telles, and Rosane
826 Minghim. Interactivenn: a web-based tool for the analysis of sets through venn diagrams.
827 *BMC bioinformatics*, 16(1):169, 2015.
- 828 20 Chang-Gun Lee, Joosun Han, Yang-Min Seo, Sang Luyi Min, Rhan Ha, Seongsoo Hong,
829 Chang Yun Park, Minsuk Lee, and Chong Sam Kim. Analysis of cache-related preemption
830 delay in fixed-priority preemptive scheduling. *IEEE Transactions on Computers*, 47(6):700–713,
831 1998.
- 832 21 Rodolfo Pellizzoni, Bach D Bui, Marco Caccamo, and Lui Sha. Coscheduling of CPU and
833 I/O transactions in COTS-based embedded systems. In *Real-Time Systems Symposium, 2008*,
834 pages 221–231. IEEE, 2008.
- 835 22 Harini Ramaprasad and Frank Mueller. Bounding worst-case response time for tasks with
836 non-preemptive regions. In *2008 IEEE Real-Time and Embedded Technology and Applications
837 Symposium*, pages 58–67. IEEE, 2008.
- 838 23 Harini Ramaprasad and Frank Mueller. Tightening the bounds on feasible preemptions. *ACM
839 Transactions on Embedded Computing Systems (TECS)*, 10(2):27, 2010.
- 840 24 Syed Aftab Rashid, Geoffrey Nelissen, Sebastian Altmeyer, Robert I Davis, and Eduardo
841 Tovar. Integrated analysis of cache related preemption delays and cache persistence reload
842 overheads. In *2017 IEEE Real-Time Systems Symposium (RTSS)*, pages 188–198. IEEE, 2017.
- 843 25 Syed Aftab Rashid, Geoffrey Nelissen, Damien Hardy, Benny Akesson, Isabelle Puaut, and
844 Eduardo Tovar. Cache-persistence-aware response-time analysis for fixed-priority preemptive
845 systems. In *2016 28th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 262–272.
846 IEEE, 2016.
- 847 26 Altmeyer Sebastian, Douma Roeland, Lunniss Will, and I Davis Robert. Evaluation of cache
848 partitioning for hard real-time systems. In *proceedings Euromicro Conference on Real-Time
849 Systems (ECRTS)*, pages 15–26, 2014.
- 850 27 Darshit Shah, Sebastian Hahn, and Jan Reineke. Experimental evaluation of cache-related
851 preemption delay aware timing analysis. In *18th International Workshop on Worst-Case
852 Execution Time Analysis (WCET 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik,
853 2018.
- 854 28 Jan Staschulat, Simon Schliecker, and Rolf Ernst. Scheduling analysis of real-time systems
855 with precise modeling of cache related preemption delay. In *Real-Time Systems, 2005.(ECRTS
856 2005). Proceedings. 17th Euromicro Conference on*, pages 41–48. IEEE, 2005.
- 857 29 Gregory Stock, Sebastian Hahn, and Jan Reineke. Cache persistence analysis: Finally exact.
858 In *Real-Time Systems Symposium (RTSS)*, Dec 2019.
- 859 30 Yudong Tan and Vincent Mooney. Timing analysis for preemptive multitasking real-time
860 systems with caches. *ACM Transactions on Embedded Computing Systems (TECS)*, 6(1):7,
861 2007.
- 862 31 Hiroyuki Tomiyama and Nikil D Dutt. Program path analysis to bound cache-related preemp-
863 tion delay in preemptive real-time systems. In *Proceedings of the eighth international workshop
864 on Hardware/software codesign*, pages 67–71. ACM, 2000.