

Improved cache related pre-emption delay aware response time analysis for fixed priority pre-emptive systems

Sebastian Altmeyer · Robert I. Davis · Claire Maiza

Abstract Without the use of caches the increasing gap between processor and memory speeds in modern embedded microprocessors would have resulted in memory access times becoming an unacceptable bottleneck. In such systems, cache related pre-emption delays can be a significant proportion of task execution times. To obtain tight bounds on the response times of tasks in pre-emptively scheduled systems, it is necessary to integrate worst-case execution time analysis and schedulability analysis via the use of an appropriate model of pre-emption costs.

In this paper, we introduce a new method of bounding pre-emption costs, called the ECB-Union approach. The ECB-Union approach complements an existing UCB-Union approach. We improve upon both of these approaches via the introduction of Multiset variants which reduce the amount of pessimism in the analysis. Further, we combine these Multiset approaches into a simple composite approach that dominates both. These approaches to bounding pre-emption costs are integrated into response time analysis for fixed priority pre-emptively scheduled systems. Further, we extend this analysis to systems where tasks can access resources in mutual exclusion, in the

This paper forms an extended version of Altmeyer et al. (2011b). The main extensions are as follows. Section 4 introduces the ECB-Union Multiset and UCB-Union Multiset approaches, and includes additional examples showing how pessimism can occur with the ECB-Union approach. Section 5 shows how blocking can be integrated into the Multiset approaches. Sections 6 and 7 provide a case study and empirical evaluation examining the performance of the Multiset approaches.

S. Altmeyer (✉)

Compiler Design Lab, Saarland University, Saarbrücken, Germany

e-mail: altmeyer@cs.uni-saarland.de

R.I. Davis

University of York, York, UK

e-mail: rob.davis@cs.york.ac.uk

C. Maiza

Verimag, INP Grenoble, Grenoble, France

e-mail: Claire.Maiza@imag.fr

process resolving omissions in existing models of pre-emption delays. A case study and empirical evaluation demonstrate the effectiveness of the ECB-Union, Multiset and combined approaches for a wide range of different cache configurations including cache utilization, cache set size, reuse, and block reload times.

Keywords Fixed priority pre-emptive scheduling · Cache related pre-emption delay · Response time analysis

1 Introduction

During the last two decades, applications in aerospace and automotive electronics have progressed from deploying embedded microprocessors clocked in the 10's of MHz range to significantly higher performance devices operating in the high 100's of MHz to GHz range. The use of such high performance embedded processors has meant that memory access times have become a significant bottleneck, necessitating the use of cache to tackle the increasing gap between processor and memory speeds.

In the majority of research papers on fixed priority pre-emptive scheduling an assumption is made that the costs of pre-emption can either be neglected or subsumed into the worst-case execution time of each task. With today's high performance embedded processors, pre-emption costs can make up a significant proportion of each task's execution time. Such costs cannot be neglected nor is it necessarily viable to simply subsume them into worst-case execution times, as this can lead to a pessimistic overestimation of response times.

In this paper, we consider the costs incurred when a pre-empting task evicts useful cache blocks of a pre-empted task. These useful cache blocks subsequently need to be reloaded after the pre-empted task resumes execution, introducing an additional cache related pre-emption delay (CRPD).

Non-pre-emptive scheduling is one way of avoiding such cache related pre-emption costs; however, disabling pre-emption is often not an option. Systems that include tasks or interrupt handlers with short deadlines typically cannot disable pre-emption for the full duration of each task's execution. An alternative approach is co-operative scheduling, with re-scheduling only possible at specific pre-emption points within each task, or after a pre-determined time has elapsed, thus dividing each task into a series of non-pre-emptable sections. Recently, significant progress has been made in this area, with algorithms designed to make an optimal selection of pre-emption points Bertogna et al. (2010, 2011). These algorithms minimise the overall cost of pre-emption for each task while maintaining the schedulability of the taskset as a whole. However, difficulties remain, for example in determining the placement of pre-emption points when the code includes branches and loops.

Exact response time analysis for fixed priority pre-emptive systems was developed during the 1980's and 1990's and subsequently refined into a set of engineering techniques (Joseph and Pandya 1986; Audsley et al. 1993; Davis et al. 2008). However, basic response time analysis does not consider cache related pre-emption costs explicitly. In (Schneider 2000), the pre-emption costs were effectively subsumed into the execution time bound for each task. While this approach accounts for both direct and

indirect or nested pre-emption costs, the cumulative way in which pre-emption costs are considered can introduce significant pessimism in the computed response times. Explicit integration of pre-emption costs has previously been considered in a number of ways: analyzing the effect of the pre-empting task (Busquets-Mataix et al. 1996; Tomiyama and Dutt 2000), the effect on the pre-empted task (Lee et al. 1998), or a combination of both (Staschulat et al. 2005; Tan and Mooney 2007). With later refinements giving an upper bound on the number of pre-emptions (Ramaprasad and Mueller 2006). More recent work has also considered how the pre-emption cost varies with the progress of task execution (Marinho et al. 2012).

In fixed priority pre-emptive systems, there are a number of ways of managing task priorities that can be used to reduce the number of pre-emptions and hence the overall pre-emption costs. These include; non-pre-emption groups (Davis et al. 2000), pre-emption thresholds (Keskin et al. 2010; Regehr 2002; Wang and Saksena 1999), and FP-FIFO scheduling (Martin et al. 2007), which is supported by a large number of real-time operating systems, including the Linux kernel (SCHED_FIFO).

In this paper, we build upon previous work that integrates pre-emption costs into response time analysis for fixed priority pre-emptive scheduling. Section 2 introduces the scheduling model, terminology, and notation used. In Sect. 3, we review existing approaches to integrating pre-emption costs into response time analysis. Building on the insights gained from this review, Sect. 4 introduces the new ECB-Union and Multiset approaches to computing pre-emption costs. The ECB-Union approach complements an existing UCB-Union approach, and the Multiset approaches remove some pessimism present in both. We combine the ECB-Union Multiset and UCB-Union Multiset approaches into a simple composite that dominates both. In Sect. 5, we extend our analysis to systems where tasks can access resources in mutual exclusion, in the process resolving omissions in existing models of pre-emption delays. A case study in Sect. 6 and an empirical evaluation in Sect. 7 demonstrate the effectiveness of the ECB-Union, Multiset and combined approaches for a wide range of different task parameters and cache configurations. Section 8 concludes with a summary of the main contributions of the paper.

The research in this paper focuses on fixed priority pre-emptive scheduling with unique priority levels; however, the approaches derived are also applicable to FP-FIFO scheduling.

2 Task model, terminology, and notation

We are interested in an application executing under a fixed priority pre-emptive scheduler on a single processor. The application is assumed to comprise a static set of n tasks $(\tau_1, \tau_j, \dots, \tau_n)$, each assigned a fixed priority. We use the notation $hp(i)$ (and $lp(i)$) to mean the set of tasks with priorities higher than (lower than) that of τ_i . Similarly, we use the notation $hep(i)$ (and $lep(i)$) to mean the set of tasks with priorities higher than or equal to (lower than or equal to) that of τ_i . We consider systems where each task has a unique priority.

Application tasks may arrive either periodically at fixed intervals of time, or sporadically after some minimum inter-arrival time has elapsed. Each task, is characterized by: its relative deadline D_i , worst-case execution time C_i , minimum inter-arrival time or period T_i and release jitter J_i , defined as the maximum time between the task arriving and it being released (ready to execute). Tasks are assumed to have constrained deadlines, i.e. $D_i \leq T_i$. It is assumed that once a task starts to execute it will never voluntarily suspend itself. The processor utilization U_i of task τ_i is given by C_i/T_i . The total utilization U of a taskset is the sum of the individual task utilizations. The worst-case response time R_i of a task τ_i , is the longest time from it becoming ready to execute to it completing execution. A task is referred to as schedulable if its worst-case response time is less than or equal to its deadline less release jitter ($R_i \leq D_i - J_i$). A taskset is referred to as schedulable if all of its tasks are schedulable. We use $E_j(R_i)$ to denote the maximum number of times that a task τ_j can execute (pre-empt) during the response time R_i of task τ_i .

In Sects. 3 and 4 we assume that tasks are independent. In Sect. 5, we relax this restriction, permitting tasks to access shared resources in mutual exclusion according to the Stack Resource Policy (SRP) (Baker 1991). As a result of the operation of the SRP, a task τ_i may be blocked by lower priority tasks for at most B_i , referred to as the blocking time.

In our analysis of cache related pre-emption delays, we use $\text{aff}(i, j)$ to mean the set of tasks that can not only execute between the release and completion of task τ_i and so affect its response time, but can also be pre-empted by task τ_j . For the basic task model, without shared resources, $\text{aff}(i, j) = \text{hep}(i) \cap \text{lp}(j)$.

With respect to a given system model, a schedulability test is said to be *sufficient* if every taskset it deems to be schedulable is in fact schedulable. Similarly, a schedulability test is said to be *necessary* if every taskset it deems to be unschedulable is in fact unschedulable. Tests that are both sufficient and necessary are referred to as *exact*.

A schedulability test A is said to *dominate* another schedulability test B if all of the tasksets deemed schedulable by test B are also deemed schedulable by test A, and there exist tasksets that are schedulable according to test A but not according to test B. Schedulability tests A and B are said to be *incomparable* if there exists tasksets that are deemed schedulable by test A and not by test B and also tasksets that are deemed schedulable by test B and not by test A.

2.1 Pre-emption costs

We now extend the sporadic task model introduced above to include pre-emption costs. To this end, we need to explain how pre-emption costs can be derived. To simplify the following explanation and examples, we assume direct-mapped caches.

The additional execution time due to pre-emption is mainly caused by cache eviction: the pre-empting task evicts cache blocks of the pre-empted task that have to be reloaded after the pre-empted task resumes. The additional context switch costs due to the scheduler invocation and a possible pipeline-flush can be upper-bounded by a constant. We assume that these *constant* costs are already included in C_i . Hence, from here on, we use *pre-emption cost* to refer only to the cost of additional cache

reloads due to pre-emption. This cache related pre-emption delay (CRPD) is bounded by $g \times \text{BRT}$ where g is an upper bound on the number of cache block reloads due to pre-emption and BRT is an upper-bound on the time necessary to reload a memory block in the cache (block reload time).

To analyse the effect of pre-emption on a pre-empted task, Lee et al. (1998) introduced the concept of a useful cache block: A memory block m is called a useful cache block (UCB) at program point \mathcal{P} , if (i) m may be cached at \mathcal{P} and (ii) m may be reused at program point \mathcal{Q} that may be reached from \mathcal{P} without eviction of m on this path. In the case of pre-emption at program point \mathcal{P} , only the memory blocks that (i) are cached and (ii) will be reused, may cause additional reloads. Hence, the number of UCBs at program point \mathcal{P} gives an upper bound on the number of additional reloads due to a pre-emption at \mathcal{P} . The maximum possible pre-emption cost for a task is determined by the program point with the highest number of UCBs. For each subsequent pre-emption, the program point with next smaller number of UCBs can be considered. Thus, the j -th highest number of UCBs can be counted for the j -th pre-emption. A tighter definition is presented by Altmeyer and Burguière (2009); however, in this paper we need only the basic concept.

The worst-case impact of a pre-empting task is given by the number of cache blocks that the task may evict during its execution. Recall that we consider direct-mapped caches: in this case, loading one block into the cache may result in the eviction of at most one cache block. A memory block accessed during the execution of a pre-empting task is referred to as an evicting cache block (ECB). Accessing an ECB may evict a cache block of a pre-empted task.

In this paper, we represent the sets of ECBs and UCBs as sets of integers with the following meaning:

$$\begin{aligned} s \in \text{UCB}_i &\Leftrightarrow \tau_i \text{ has a useful cache block in cache-set } s \\ s \in \text{ECB}_i &\Leftrightarrow \tau_i \text{ may evict a cache block in cache-set } s \end{aligned}$$

A bound on the pre-emption cost due to task τ_j *directly* pre-empting τ_i is therefore given by $\text{BRT} \cdot |\text{UCB}_i \cap \text{ECB}_j|$. Precise computation is more complex as different program points may exhibit different sets of UCBs. Hence, the worst-case pre-emption delay considering a pre-empting and pre-empted task may not necessarily occur at the pre-emption point with the highest number of UCBs but instead at the point with the largest intersection between UCBs and ECBs—(see Altmeyer and Maiza 2011) for a detailed description of the computation of pre-emption costs. Note that the simplification we apply, using ECB_i and UCB_i , does not impact the correctness of the equations.

Separate computation of the pre-emption cost is restricted to architectures without timing anomalies (Lundqvist and Stenström 1999) but is independent of the type of cache used, i.e. data, instruction or unified cache.

2.2 Set-associative caches

In the case of set-associative LRU caches,¹ a single cache-set may contain several useful cache blocks. For instance, $UCB_1 = \{1, 2, 2, 2, 3, 4\}$ means that task τ_1 contains 3 UCBs in cache-set 2 and one UCB in each of the cache sets 1, 3 and 4. As one ECB suffices to evict all UCBs of the same cache-set, multiple accesses to the same set by the pre-empting task does not need to appear in the set of ECBs. Hence, we keep the set of ECBs as used for direct-mapped caches. A bound on the CRPD in the case of LRU caches due to task τ_i directly pre-empting τ_j is thus given by the intersection $UCB_j \cap' ECB_i = \{m | m \in UCB_j : m \in ECB_i\}$, where the result is also a multiset that contains each element from UCB_j if it is also in ECB_i . A precise computation of the CRPD in the case of LRU caches is given in Altmeyer et al. (2010). In this paper, we assume direct-mapped caches. All equations provided within this paper are for direct-mapped caches, they are also valid for set-associative LRU caches with the above adaptation to the set-intersection.

3 Response time analysis for pre-emptive systems

Response time analysis (Audley et al. 1993; Joseph and Pandya 1986) for fixed priority pre-emptive scheduling calculates the worst-case response time R_i of task τ_i , using the following equation.

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i + J_j}{T_j} \right\rceil (C_j) \quad (1)$$

The worst-case response time appears on both the left-hand side (LHS) and the right-hand side RHS of the equation. As the RHS is a monotonically non-decreasing function of R_i , the equation can be solved using fixed point iteration: Iteration starts with an initial value for the response time, typically $r_i^0 = C_i$, and ends either when $r_i^{n+1} = r_i^n$ in which case the worst-case response time R_i is given by r_i^n or when $r_i > D_i - J_i$ in which case the task is unschedulable. We note that (1) does not explicitly include pre-emption costs.

When pre-emption costs are included, then synchronous release of all higher priority tasks does not necessarily represent the worst-case scenario as shown by Meumeu Yoms and Sorel (2007). The worst-case number of pre-emptions due to each higher priority task can however be bounded, and used to extend (1) forming a sufficient, but not exact schedulability test.

Equation (1) can be extended by $\gamma_{i,j}$ representing the pre-emption cost due to each job of a higher priority *pre-empting* task τ_j executing within the worst-case response time of task τ_i (Busquets-Mataix et al. 1996):

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i + J_j}{T_j} \right\rceil (C_j + \gamma_{i,j}) \quad (2)$$

¹The concept of UCBs and ECBs cannot be applied to FIFO or PLRU replacement policies as shown by Burguère et al. (2009).

Task τ_j does not necessarily pre-empt task τ_i directly; a nested pre-emption is also possible. Any pre-emption by task τ_j of a task τ_k that executes while τ_i is pre-empted may also increase the response time of task τ_i . The problem of obtaining a valid yet tight upper bound on the pre-emption costs is made difficult by the effects of nested pre-emption, as a pre-empting task may evict useful cache-blocks belonging to a number of pre-empted tasks.

The precise meaning of $\gamma_{i,j}$ and its computation depends on the approach used. Below, we review a number of existing approaches and discuss their advantages and disadvantages.

3.1 ECB-Only

Busquets-Mataix et al. (1996) and later Tomiyama and Dutt (2000), used the ECBs of the pre-empting task to bound the pre-emption costs:

$$\gamma_{i,j}^{\text{ecb}} = \text{BRT} \cdot |\text{ECB}_j| \quad (3)$$

In this case, $\gamma_{i,j}$ represents the worst-case effect of task τ_j on any arbitrary lower priority task, independent of such a task's actual UCBs.

3.2 UCB-Only

By contrast, Lee et al. (1998) used the number of UCBs to bound the pre-emption costs. Here, however one has to correctly account for nested pre-emptions. The cost of τ_j pre-empting some task τ_k of intermediate priority may be higher than that of τ_j pre-empting τ_i . Thus, the pre-emption cost due to a job of task τ_j executing during the response time of task τ_i is only bounded by the maximum number of UCBs over all tasks that may be pre-empted by τ_j and have at least the priority of τ_i (i.e. tasks from the set $\text{aff}(i, j) = \text{hep}(i) \cap \text{lp}(j)$).

$$\gamma_{i,j}^{\text{ucb}} = \text{BRT} \cdot \max_{\forall k \in \text{aff}(i, j)} \{|\text{UCB}_k|\} \quad (4)$$

The disadvantage of the ECB-Only and UCB-Only approaches is clear: considering only the pre-empted tasks *or* alternatively only the pre-empting tasks leads to an over-approximation. Not every UCB may be evicted during pre-emption, and not every ECB may evict a UCB. This is illustrated in Fig. 1.






Fig. 1 Taskset $\{\tau_1, \tau_2\}$ with $C_1 = 1$, $C_2 = 2$ and block reload time 1. Response time analysis of task τ_2 : only counting the number of possibly evicted UCBs (4) or possibly evicting ECBs (3) leads to a pre-emption cost of 2, whereas the actual pre-emption cost is 0

Fig. 2 Two tasksets $\{\tau_1, \tau_2, \tau_3\}$ with $C_1 = 1, C_2 = 2, C_3 = 3$, and a block reload time of 1



(a) τ_1 pre-empting τ_2 causes higher costs than τ_1 pre-empting τ_3 .



(b) Nested pre-emption: τ_1 pre-empting τ_2 pre-empting τ_3 , causes higher costs than any non-nested pre-emption.

Figure 1 shows an example taskset that leads to an overestimation when the pre-emption cost is estimated using (3) or (4). Task τ_1 accesses blocks in cache sets 1 and 2. Task τ_2 accesses blocks in cache sets 1, 2, 3 and 4. However, only sets 3 and 4 may contain useful cache blocks, hence a pre-emption of task τ_2 by task τ_1 never evicts any useful cache blocks; and so there are no cache reloads due to pre-emption. However, (4) and (3) account for 2 additional reloads; an overestimation of the pre-emption cost.

Since both (3) and (4) can over-estimate the actual pre-emption cost, combining both UCBs and ECBs might be expected to result in precise bounds. However, the naive computation $\gamma_{i,j} = \text{BRT} \cdot |\text{UCB}_i \cap \text{ECB}_j|$ is optimistic and thus cannot be used. It may lead to underestimation in two cases: when the cost of task τ_j pre-empting a task τ_k of intermediate priority is higher than that of τ_j pre-empting τ_i (see Fig. 2(a)) and when the execution of τ_j may evict useful cache blocks of both task τ_i and of task τ_k (see Fig. 2(b)).

3.3 UCB-Union

Tan and Mooney (2007) considered both the pre-empted and the pre-empting task. They take the union of all possible affected useful cache blocks and combine this with the set of ECBs of the pre-empting task.


$$\gamma_{i,j}^{\text{ucb-u}} = \text{BRT} \cdot \left| \left(\bigcup_{\forall k \in \text{aff}(i,j)} \text{UCB}_k \right) \cap \text{ECB}_j \right| \quad (5)$$

This UCB-Union approach dominates the ECB-Only approach since:

$$\gamma_{i,j}^{\text{ecb}} \geq \gamma_{i,j}^{\text{ucb-u}}$$

but may be worse than the UCB-Only approach in some cases. Consider the taskset shown in Fig. 3, the values of $\gamma_{i,j}$ for the response time analysis of task τ_3 are as

Fig. 3 Taskset $\{\tau_1, \tau_2, \tau_3\}$ with $C_1 = 1, C_2 = C_3 = 2$, and a block reload time of 1. Equation (5) computes total pre-emption costs of 6, whereas the actual cost is only 4



follows:

$$\gamma_{3,1}^{\text{ucb-u}} = |\{(\text{UCB}_2 \cup \text{UCB}_3) \cap \text{ECB}_1\}| = |\{1, 2, 3, 4\} \cap \{1, 2, 3, 4\}| = 4$$

$$\gamma_{3,2}^{\text{ucb-u}} = |\{(\text{UCB}_3) \cap \text{ECB}_2\}| = |\{3, 4\} \cap \{1, 2, 3, 4\}| = 2$$

Given that each task is executed at most once, the total computed pre-emption cost is 6. However, the actual pre-emption cost is only 4: either UCBs in cache sets $\{1, 2, 3, 4\}$ have to be reloaded (in the case of nested pre-emption) or UCBs in cache sets $\{3, 4\}$ are reloaded twice (in the case of consecutive pre-emption of τ_3 by τ_1 and then by τ_2).

In the case of set-associative caches, Tan and Mooney (2007) account only for those cache blocks that are actually evicted due to pre-emption. We note that this can be optimistic, as shown by Burguère et al. (2009).

3.4 Staschulat's formula

Staschulat et al. (2005) also combine information about the pre-empting and the pre-empted task; however, their approach is somewhat more complex than the methods described so far. Below we give a concise description of their method, for further details and a more complete description see Staschulat et al. (2005).

The analysis of Staschulat et al. is extended to account for the fact that each additional pre-emption of task τ_i may result in a smaller pre-emption cost than the last. (Their approach is an improvement over that of Petters and Farber 2001.) The basic response time analysis used differs from (2): $\gamma_{i,j}$ does not refer to the cost of a single pre-emption, but instead to the total cost of all pre-emptions due to jobs of task τ_j executing within the response time of task τ_i .

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left(\left\lceil \frac{R_i + J_j}{T_j} \right\rceil C_j + \gamma_{i,j}^{\text{sta}} \right) \quad (6)$$

Staschulat et al. (2005) compute the maximum number of pre-emptions q , including nested pre-emptions, which may impact the response time of task τ_i due to cache blocks evicted by task τ_j . Thus q is given by the sum of the maximum number of jobs of task τ_j and tasks of lower priority than τ_j but higher priority than τ_i that can execute during the response time R_i of task τ_i

$$q = \sum_{\forall k \in \text{hp}(i) \cap (\text{lp}(j) \cup \{j\})} E_k(R_i) \quad (7)$$

where $E_k(R_i)$ is used to denote the maximum number of jobs of task τ_k that can execute during response time R_i . For our task model, $E_k(R_i) = \lceil (R_i + J_k) / T_k \rceil$. The total pre-emption cost $\gamma_{i,j}^{\text{sta}}$ due to jobs of task τ_j pre-empting during the response time of task τ_i is then bounded by the q largest costs of task τ_j pre-empting jobs of any lower priority task $\tau_k \in \text{hep}(i) \cap \text{lp}(j)$ that can execute during the response time of task τ_i . As each job of such a task τ_k may execute up to $E_k(R_i)$ times during R_i , and each of those jobs could potentially be pre-empted at most $E_j(R_k)$ times by task τ_j , the $E_j(R_k)$ highest pre-emption costs of τ_j directly pre-empting τ_k must be considered $E_k(R_i)$ times:

$$\gamma_{i,j}^{\text{sta}} = \text{BRT} \cdot \sum_{l=1}^q |M^l| \quad (8)$$

where M^l is the l -th largest element from the multiset M

$$M = \bigcup_{k \in \text{hep}(i) \cap \text{lp}(j)} \left(\bigcup_{E_k(R_i)} \{(\text{UCB}_k \cap \text{ECB}_j)^n \mid n \in [1; E_j(R_k)]\} \right) \quad (9)$$

and $(\text{UCB}_k \cap \text{ECB}_j)^n$ gives the n -th highest pre-emption cost for task τ_j pre-empting task τ_k . Note that M is a multiset and the union over $E_k(R_i)$ means that the set of values for τ_k are repeated $E_k(R_i)$ times.

The drawback of this approach is that the number of pre-emptions taken into account strongly over-estimates the number of pre-emptions that have an actual influence on the response time; particularly when there are a large number of tasks. In addition, the reduction in the pre-emption costs for a sequence of pre-emptions is typically rather limited (Bertogna et al. 2011 show that the maximal pre-emption cost can occur at various program points within a task's execution). The program point \mathcal{P} in a task which exhibits the highest number of UCBs often occurs within a loop, thus, it has to be taken into account as often as the loop iterates. In addition, program points close to \mathcal{P} will often have a similar number of UCBs. We note that Staschulat et al. also present an improvement to their analysis in Staschulat et al. (2005); however, the problem of strongly over-estimating the number of pre-emptions remains.

4 ECB-Union and multiset approaches

This section presents new approaches to pre-emption cost aware response time analysis. The ECB-Union approach complements the UCB-Union approach and the Multiset extensions remove some pessimism present in both these methods.

4.1 ECB-Union approach

With the ECB-Union approach, to account for nested pre-emptions, we compute the union of all ECBs that may affect a pre-empted task. The intuition here is that direct pre-emption by task τ_j is represented by the pessimistic assumption that task τ_j has

itself already been pre-empted by all of the tasks of higher priority and hence may result in eviction by the set $\bigcup_{h \in \text{hp}(j) \cup \{j\}} \text{ECB}_h$

$$\gamma_{i,j}^{\text{ecb-u}} = \text{BRT} \cdot \max_{\forall k \in \text{aff}(i,j)} \left\{ \left| \text{UCB}_k \cap \left(\bigcup_{h \in \text{hp}(j) \cup \{j\}} \text{ECB}_h \right) \right| \right\} \quad (10)$$

Task τ_j may directly pre-empt any task $\tau_k \in \text{aff}(i, j)$ impacting the response time of task τ_i . Thus taking the maximum over all of the tasks in $\text{aff}(i, j)$ ensures that the pre-emption cost for the highest number of evicted useful cache blocks is considered. Note we use $\text{hp}(j) \cup \{j\}$ to mean task τ_j and all tasks of higher priority than task τ_j , rather than $\text{hep}(j)$. This is because the two sets are different in the more general case where tasks can share priority levels. Note that (10) is combined with (2) to determine task response times. The ECB-Union approach (10) dominates the UCB-Only approach, since:


$$\gamma_{i,j}^{\text{ucb}} \geq \gamma_{i,j}^{\text{ecb-u}}$$

The ECB-Union approach and the UCB-Union approach (Tan and Mooney 2007) are incomparable. Figure 3 provides an example where the ECB-Union approach outperforms the UCB-Union approach: here the ECB-Union approach covers both a nested pre-emption (τ_3 pre-empted by τ_2 which is pre-empted by τ_1) and consecutive pre-emption (of τ_3 by τ_1 and τ_2), obtaining for each pre-emption a cost of 2 and thus, a total cost of 4. In contrast, the UCB-Union approach gives a total cost of 6.

$$\begin{aligned} \gamma_{3,1}^{\text{ecb-u}} &= \max_{\forall k \in \{2,3\}} \{ |\text{UCB}_k \cap \text{ECB}_1| \} \\ &= \max \{ |\text{UCB}_2 \cap \text{ECB}_1|, |\text{UCB}_3 \cap \text{ECB}_1| \} \\ &= \max \{ |\{1, 2\}|, |\{3, 4\}| \} = 2 \\ \gamma_{3,2}^{\text{ecb-u}} &= \max_{\forall k \in \{3\}} \{ |\text{UCB}_k \cap (\text{ECB}_1 \cap \text{ECB}_2)| \} \\ &= |\text{UCB}_3 \cap (\text{ECB}_1, \text{ECB}_2)| \\ &= |\{3, 4\} \cap \{1, 2, 3, 4\}| = |\{3, 4\}| = 2 \end{aligned}$$

Figure 4 provides an example from Altmeyer et al. (2011a) where the UCB-Union approach outperforms the ECB-Union approach. Here, for the latter approach, the

Fig. 4 Taskset $\{\tau_1, \tau_2, \tau_3\}$ with $C_1 = 1, C_2 = C_3 = 2$, and block reload time 1. Equation (10) computes a total pre-emption cost of 6, whereas the actual cost is only 4






Fig. 5 Possible pessimism of the ECB-Union approach. Pre-emption cost of task τ_1 pre-empting task τ_2 contributes six times to the response time of task τ_3 ; although it may happen at most three times

pre-emption costs increasing the response time R_3 of task τ_3 are computed as follows:

$$\begin{aligned}
 \gamma_{3,1}^{\text{ecb-u}} &= \max_{\forall k \in \{2,3\}} \{|\text{UCB}_k \cap \text{ECB}_1|\} \\
 &= \max\{|\text{UCB}_2 \cap \text{ECB}_1|, |\text{UCB}_3 \cap \text{ECB}_1|\} \\
 &= \max\{|\emptyset|, |\{1, 2\}|\} = 2 \\
 \gamma_{3,2}^{\text{ecb-u}} &= \max_{\forall k \in \{3\}} \{|\text{UCB}_k \cap (\text{ECB}_1 \cap \text{ECB}_2)|\} \\
 &= |\text{UCB}_3 \cap (\text{ECB}_1, \text{ECB}_2)| \\
 &= |\{1, 2, 3, 4\} \cap \{1, 2, 3, 4\}| = |\{1, 2, 3, 4\}| = 4
 \end{aligned}$$

With the ECB-Union approach, the eviction of UCBs of task τ_3 ($\{1, 2\}$) are considered twice, even though they must be reloaded at most once, leading to an over-estimation of the total pre-emption costs of 6. The UCB-Union approach, in this case, computes the precise total of 4.

The ECB-Union approach removes some of the pessimism in accounting for nested pre-emptions, by including the direct effects of $E_j(R_i)$ pre-emptions by task τ_j , during the response time R_i of task τ_i . The indirect effect of pre-emptions by task τ_j are then accounted for via the ECB-Union term in (10). Thus the indirect effects of any cache block evictions by task τ_j are included when the pre-emption costs of an intermediate task $\tau_k \in \text{hp}(i) \cap \text{lp}(j)$ are accounted for. This avoids some over-counting of indirect pre-emption costs as illustrated in Fig. 5. Figure 5 shows the worst case scenario for task τ_3 . During the response time of task τ_3 , task τ_1 can pre-empt a total of 6 times. Task τ_1 can have a direct effect on both task τ_2 and τ_3 as shown by the solid arrows leading to pre-emption delays and an indirect effect on task τ_3 as shown by the dashed arrow. Task τ_2 can also have a direct effect on τ_3 . The ECBUnion approach, accounts for the indirect effect of task τ_1 on task τ_3 in the pre-emption cost term for task τ_2 pre-empting task τ_3 via the inclusion of ECB_1 in the ECB-union.

4.2 ECB-Union Multiset approach

The ECB-Union approach is pessimistic in that it assumes that task τ_j can pre-empt any task $\tau_k \in \text{aff}(i, j)$ up to $E_j(R_i)$ times. While this is potentially true when $\tau_k = \tau_i$, it can be a pessimistic assumption in the case of intermediate tasks. This is illustrated by the taskset in Fig. 5. Here, the direct pre-emption cost of task τ_1 pre-empting task

τ_2 is higher than that of task τ_1 pre-empting task τ_3 . Thus the pre-emption cost $\gamma_{3,1}$ comes from $|\text{UCB}_2 \cap \text{ECB}_1|$ which is included $E_1(R_3) = 6$ times. However, this is pessimistic, task τ_1 can only actually pre-empt task τ_2 at most $E_1(R_2)E_2(R_3) = 3$ times during the response time of task τ_3 . The ECB-Union Multi-set approach removes this source of pessimism by combining Staschulat's method of counting the maximum number of pre-emptions of each intermediate task with the ECB-Union approach. Similar to Staschulat's approach, the ECB-Union Multi-set approach uses $\gamma_{i,j}^{\text{ecb-m}}$ to represent the pre-emption costs due to all jobs of task τ_j executing during the response time of task τ_i , and therefore effectively makes use of (6) as the response time analysis equation. The pre-emption cost $\gamma_{i,j}^{\text{ecb-m}}$ is computed as follows, recognising the fact that task τ_j can pre-empt each intermediate task τ_k at most $E_j(R_k)E_k(R_i)$ times during the response time of task τ_i . First, we form a multiset M that contains the cost

$$\left| \text{UCB}_k \cap \left(\bigcup_{h \in \text{hp}(j) \cup \{j\}} \text{ECB}_h \right) \right| \quad (11)$$

of τ_j pre-empting task τ_k $E_j(R_k)E_k(R_i)$ times, for each task $\tau_k \in \text{aff}(i, j)$. Hence:

$$M = \bigcup_{k \in \text{aff}(i, j)} \left(\bigcup_{E_j(R_k)E_k(R_i)} \left| \text{UCB}_k \cap \left(\bigcup_{h \in \text{hp}(j) \cup \{j\}} \text{ECB}_h \right) \right| \right) \quad (12)$$

$\gamma_{i,j}^{\text{ecb-m}}$ is then given by the $E_j(R_i)$ largest values in M .

$$\gamma_{i,j}^{\text{ecb-m}} = \text{BRT} \cdot \sum_{l=1}^{E_j(R_i)} |M^l| \quad (13)$$

where M^l is the l -th largest value in M . We note that by construction, the ECB-Union Multiset approach dominates the ECB-Union approach.

4.3 UCB-Union Multiset approach

The UCB-Union approach is pessimistic in that it assumes that task τ_j can, directly or indirectly, pre-empt any task $\tau_k \in \text{aff}(i, j)$ up to $E_j(R_i)$ times during the response time R_i of task τ_i . While this is potentially true when $\tau_k = \tau_i$, it can be a pessimistic assumption in the case of intermediate tasks. This is illustrated by the taskset in Fig. 6. Here, task τ_1 can only pre-empt task τ_2 once during the response time of task τ_3 , yet




Fig. 6 Possible pessimism of the UCB-Union approach. Useful cache blocks of task τ_2 are assumed to contribute up to three times to the response time of task τ_3 . Task τ_2 , however, is executed at most once during the response time of τ_3

the UCB-Union approach considers that UCBs of task τ_2 could be evicted by all three jobs of task τ_1 executing within this interval: $\gamma_{3,1} = (\text{UCB}_2 \cup \text{UCB}_3) \cap \text{ECB}_1$ with $\gamma_{3,1}$ included $E_1(R_3) = 3$ times by the response time analysis. The UCB-Union Multi-set approach removes this source of pessimism by combining Staschulat's method of counting the maximum number of pre-emptions of each intermediate task with the UCB-Union approach. Similar to Staschulat's approach, the UCB-Union Multi-set approach uses $\gamma_{3,1}^{\text{ucb-m}}$ to represent the pre-emption costs due to all jobs of task τ_j executing during the response time of task τ_i , and therefore effectively makes use of (6) as the response time analysis equation. The pre-emption cost $\gamma_{3,1}^{\text{ucb-m}}$ is computed as follows, recognising the fact that task τ_j can pre-empt each intermediate task τ_k directly or indirectly at most $E_j(R_k)E_k(R_i)$ times during the response time of task τ_i . First, we form a multi-set $M_{i,j}^{\text{ucb}}$ containing $E_j(R_k)E_k(R_i)$ copies of the UCB_k of each task $k \in \text{aff}(i, j)$. This multi-set reflects the fact that during the response time R_i of task τ_i , task τ_j cannot evict a UCB of task τ_k more than $E_j(R_k)E_k(R_i)$ times. Hence:

$$M_{i,j}^{\text{ucb}} = \bigcup_{k \in \text{aff}(i,j)} \left(\bigcup_{E_j(R_k)E_k(R_i)} \text{UCB}_k \right) \quad (14)$$

Next, we form a multi-set M_j^{ecb} containing $E_j(R_i)$ copies of the ECB_j of task τ_j . This multi-set reflects the fact that during the response time R_i of task τ_i , task τ_j can evict ECBs in the set ECB_j at most $E_j(R_i)$ times.

$$M_j^{\text{ecb}} = \bigcup_{E_j(R_i)} (\text{ECB}_j) \quad (15)$$

$\gamma_{i,j}^{\text{ucb-m}}$ is then given by the size of the multi-set intersection of M_j^{ecb} and $M_{i,j}^{\text{ucb}}$

$$\gamma_{i,j}^{\text{ucb-m}} = \text{BRT} \cdot |M_{i,j}^{\text{ucb}} \cap M_j^{\text{ecb}}| \quad (16)$$

We note that by construction, the UCB-Union Multiset approach dominates the UCB-Union approach.

Example Let $\text{ECB}_1 = \{1, 2, 3, 4\}$, $\text{UCB}_2 = \{2, 3, 4\}$, $\text{UCB}_3 = \{1, 2\}$, and $E_1(R_3) = 3$, $E_1(R_2) = 1$, $E_2(R_3) = 1$. We then have:


$$M_{3,1}^{\text{ucb}} = \{1, 2\} \cup \{1, 2\} \cup \{1, 2\} \cup \{2, 3, 4\} = \{1, 2, 1, 2, 1, 2, 2, 3, 4\}$$

$$M_1^{\text{ecb}} = \{1, 2, 3, 4\} \cup \{1, 2, 3, 4\} \cup \{1, 2, 3, 4\} = \{1, 2, 3, 4, 1, 2, 3, 4, 1, 2, 3, 4\}$$

$$M_1^{\text{ecb}} \cap M_{3,1}^{\text{ucb}} = \{1, 2, 1, 2, 1, 2, 3, 4\}$$

Hence the total pre-emption cost given by the UCB-Union Multi-set approach $\gamma_{3,1}^{\text{ucb-m}} = 8$ in this case. This compares with a total pre-emption cost of 12 computed by the UCB-Union approach ($\gamma_{3,1}^{\text{ucb}} = 4$, included 3 times).

Fig. 7 Venn Diagramm illustrating the relation between the different pre-emption cost aware schedulability tests. The larger the area, the more tasksets deemed schedulable



4.4 Combined approach

The UCB-Union Multiset approach dominates the UCB-Union approach, which in turn dominates the ECB-Only approach. Similarly the ECB-Union Multiset approach dominates the ECB-Union approach, which in turn dominates the UCB-Only approach. Given that the UCB-Union Multiset approach (16) and the ECB-Union Multiset approach (13) are incomparable, we can combine both to deliver a more precise bound on task response times that, by construction, dominates the use of either approach alone:

$$R_i = \min(R_i^{\text{ucb-m}}, R_i^{\text{ecb-m}}) \quad (17)$$

where $R_i^{\text{ucb-m}}$ is the response time of task τ_i computed using (16) and $R_i^{\text{ecb-m}}$ is the response time of task τ_i computed using (13). Figure 7 illustrates these relationships. Note that some task-sets are deemed schedulable by the combined approach, but not by any of the other approaches in isolation.

5 Blocking time

The discussion in Sects. 3 and 4 assumes non-blocking execution, i.e. no shared resources. In this section, we relax this restriction, permitting tasks to access mutually exclusive shared resources according to the Stack Resource Policy (SRP) introduced by Baker (1991), extending the Priority Ceiling Protocol of Sha et al. (1990).

The SRP associates a *ceiling* priority with each resource. This ceiling priority is equal to the highest priority of any task that can access the resource. At run-time, when a task accesses a resource, its priority is immediately increased to the ceiling priority of the resource. Thus SRP bounds the amount of blocking B_i which task τ_i is subject to, to the maximum time for which any lower priority task holds a resource that is shared with task τ_i or any other task of equal or higher priority. SRP ensures that a task can only ever be blocked prior to actually starting to execute.




Fig. 8 Tasks τ_2 and τ_3 share a common resource x , τ_3 starts to execute, blocks τ_2 , which is released at time 1, and is pre-empted by τ_1 . Thus, the finishing time of τ_2 is delayed not only by the time for which τ_3 accesses the resource, but also by the additional pre-emption delay, reloading UCBs $\{1, 2\}$ after the resource access of task τ_3 is pre-empted by task τ_1

When a lower priority task τ_k locks a resource and so blocks task τ_i , it can still be pre-empted by tasks with priorities higher than that of the ceiling priority of the resource. B_i does not account for the additional pre-emption cost due to such pre-emptions.

When pre-emption costs are considered, the synchronous release of task τ_i and all higher priority tasks immediately after a lower priority task has locked a resource may not necessarily represent the worst-case scenario for task τ_i . Nevertheless, the worst-case number of pre-emptions can be bounded and used to form a sufficient test. Previous work integrating pre-emption costs into response time analysis (Busquets-Mataix et al. 1996; Lee et al. 1998; Staschulat et al. 2005; Tan and Mooney 2007) extends (2) in this way to include blocking via the simple addition of the blocking factor B_i :

$$R_i = C_i + B_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i + J_j}{T_j} \right\rceil (C_j + \gamma_{i,j}) \quad (18)$$

In the case of the analysis by Busquets-Mataix et al. (1996), this is correct, as the pre-emption cost is accounted for only via the ECBs of the pre-empting tasks and is therefore unaltered by the addition of resource accesses that could potentially also be pre-empted. In contrast, Lee et al. (1998), Staschulat et al. (2005), Tan and Mooney (2007) make use of the UCBs of pre-empted tasks. If, as is the case with the SRP, pre-emption can still occur during resource access, then these analyses are optimistic and need to be modified to correctly account for the additional pre-emption costs that can occur.² The key point is that the blocking factor B_i does not represent execution of task τ_i , but instead represents execution of some resource access within a lower priority task. Such a resource access may be pre-empted, during the response time of task τ_i and therefore its UCBs need to be taken into account, as illustrated by the example in Fig. 8.

We now extend the ECB-Union (Multiset) and UCB-Union (Multiset) approaches to take account of blocking. Specifically, we extend the pre-emption cost equations (10) and (13), and (5) and (16) to include the UCBs of tasks in the set $b(i, j)$, where

²If all resource accesses are non-pre-emptive, then there are no additional pre-emption costs to be accounted for.

$b(i, j)$ is defined as the set of tasks with priorities lower than that of task τ_i that lock a resource with a ceiling priority higher than or equal to the priority of task τ_i but lower than that of task τ_j . These tasks can block task τ_i , but can also be pre-empted by task τ_j . Hence they need to be included in the set of tasks $\text{aff}(i, j)$ whose UCBs are considered when determining the pre-emption cost $\gamma_{i,j}$ due to task τ_j :

$$\text{aff}(i, j) = (\text{hep}(i) \cap \text{lp}(j)) \cup b(i, j) \quad (19)$$

Tasks in $b(i, j)$ have lower priorities than task τ_i and so cannot pre-empt during the response time of task τ_i , hence their ECBs do not need to be considered when computing $\gamma_{i,j}$. Using (19) extends the ECB-Union approach (10) and the UCB-Union approach (5) to correctly account for pre-emption costs when tasks share resources according to the SRP. For the Multiset approaches (13) and (16), we also need to bound the number of jobs of the blocking tasks during the response time of τ_i . We do this by assuming that each blocking task can be pre-empted at most $E_j(R_i)$ times by task τ_j (i.e., assuming that $E_j(R_k)E_k(R_i)$ is given by $E_j(R_i)$ in this case).

Revisiting the example given in Fig. 8, we observe that as the set of affected tasks $\text{aff}(2, 1)$ now includes task τ_3 as well as task τ_2 , (5) correctly accounts for the overall pre-emption cost of 2 due to the resource access of task τ_3 being pre-empted by task τ_1 during the response time of task τ_2 .

We note that in the simplest case of the SRP where tasks share resources that are accessed non-pre-emptively (i.e. with ceiling priorities equal to that of the highest priority task), then the set of tasks $b(i, j)$ is empty (since no task can pre-empt during a resource access) and hence the pre-emption cost $\gamma_{i,j}$ is the same as for the basic task model, with no increase in pre-emption costs due to blocking.

Although providing valid upper bounds on the pre-emption costs, the above extension can be pessimistic. This is because it includes the UCBs of each lower priority task in $\text{aff}(i, j)$, rather than just the UCBs of each resource access within those tasks. More precise analysis could potentially be obtained by considering each resource access as a sub-task with its own UCBs. Such a detailed approach is however beyond the scope of this paper.

When determining the blocking factor B_i we cannot use the resource access execution times as they occur within the non-pre-emptive execution of each containing task τ_k . This is because we must assume that task τ_k could be pre-empted immediately before a resource access and any useful cache blocks evicted. Instead, the execution time of each resource access must be determined assuming execution of that section of code with no pre-emption, and starting from the worst-case initial state.

6 Case study

In this section, we evaluate the effectiveness of the different approaches based on a case study. The worst-case execution times and the set of useful cache blocks and evicting cache blocks have been derived from the Mälardalen benchmark suite (Gustafsson et al. 2010), see Table 1, where the values are taken from Altmeyer and Maiza (2011). The target architecture is an ARM7 processor³ with direct-mapped

³<http://www.arm.coms/products/CPUs/families/ARM7Family.html>.

Table 1 Execution times and number of UCBs and ECBs for a selection of benchmarks from the Mälardalen Benchmark Suite

	WCET	UCBs	ECBs
bs	445	5	35
minmax	504	9	79
fac	1252	4	24
fibcall	1351	5	24
insertsort	6573	10	41
loop3	13449	4	817
select	17088	15	151
qsort-exam	22146	15	170
fir	29160	9	105
sqrt	39962	14	477
ns	43319	13	64
qurt	214076	14	484
crc	290782	14	144
matmult	742585	23	100
bsort100	1567222	35	62

instruction cache of size 2 kB with a line size of 8 Bytes (and thus, 256 cache sets) and a block reload time of 8 μ s. The ARM7 features an instruction size of 4 Bytes.

We note that although the case study tasks do not represent a set of tasks scheduled on an embedded real-time system, they do represent typical components of real-time applications and thus deliver meaningful values. We created a taskset from the above data by assigning periods and implicit deadlines such that all 15 tasks had equal utilization. This is an entirely arbitrary choice. Evaluation with randomly generated taskset parameters is reported in Sect. 7. The periods were generated by multiplying each execution time by a constant c ($\forall_i : T_i = c \cdot C_i$). We varied c from 15 upwards hence varying the utilization of the taskset from 1.0 downwards. The tasks were assigned priorities in deadline monotonic priority order. Note that deadline monotonic priority order is optimal in this case only when pre-emption costs are zero.

Table 2 lists the breakdown utilization; the maximum utilization at which a scaled version of the case study taskset was deemed schedulable by each approach.

Staschulat's approach performs worst, with a breakdown utilization of 0.508. Equation (7) computes the number of pre-emptions taken into account. For the effect of task τ_1 (bs) to task τ_5 (insertsort), only the 8 highest costs of τ_1 pre-empting any task from τ_2 to τ_5 need to be considered. However, for the effect of task τ_1 (bs) to task τ_{15} (bsort100), the 47362 highest costs need to be considered. Although the single pre-emption costs (for τ_i pre-empted by τ_j) are much smaller, the total cost is very pessimistic.

The ECB-Union Multiset approach performs best with breakdown utilization of 0.789, followed by the ECB-Union and UCB-Union Multiset approaches, with breakdown utilization of 0.767. As the cache contention is high (3 out of the 15 tasks fill the whole cache), a single pre-emption often evicts all of the UCBs of the pre-empted task(s). In addition, the total number of ECBs is much higher than the total number of UCBs hence the ECB-Only approach (3) is much more pessimistic than the

Table 2 Case study taskset: breakdown utilization for different approaches

Analysis	Breakdown utilization:
No Pre-emption Cost	0.95
Combined	0.789
ECB-Union Multiset	0.789
UCB-Union Multiset	0.767
ECB-Union	0.767
UCB-Union	0.698
UCB-Only	0.75
ECB-Only	0.612
Staschulat	0.508

UCB-Only approach (4) and so has a lower breakdown utilization of 0.612. As a consequence, the ECB-Union approach (10) outperforms the UCB-Union approach (5) which has a breakdown utilization of 0.698. The combination the two Multiset approaches (17) does not improve upon the ECB-Union Multiset approach. Finally, (1) deems the case study taskset schedulable up to a utilization of 0.95 ignoring pre-emption costs.

7 Evaluation

In this section, we evaluate the effectiveness of the different approaches to pre-emption cost computation on a large number of tasksets with varying cache configurations and varying taskset parameters. The task parameters used in our experiments were randomly generated as follows:

- The default taskset size was 10.
- Task utilizations were generated using the UUnifast (Bini and Buttazzo 2005) algorithm.
- Task periods were generated according to a log-uniform distribution with a factor of 100 difference between the minimum and maximum possible task period and a minimum period of 5 ms. This represents a spread of task periods from 5 ms to 500 ms, as found in most automotive and aerospace hard real-time applications.
- Task execution times were set based on the utilization and period selected: $C_i = U_i \cdot T_i$.
- Task deadlines were implicit,⁴ i.e., $D_i = T_i$.
- Priorities were assigned in deadline monotonic order.

The following parameters affecting pre-emption costs were also varied, with default values given in parentheses:

- The number of cache-sets ($CS = 256$).
- The block-reload time ($BRT = 8 \mu s$)

⁴Evaluation for constrained deadlines, i.e., $D_i \in [C_i; T_i]$ gives broadly similar results although fewer tasksets are deemed schedulable by all approaches.

- The cache usage of each task, and thus, the number of ECBs, were generated using the UUnifast (Bini and Buttazzo 2005) algorithm (for a total cache utilization $CU = 10$). UUnifast may produce values larger than 1 which means a task fills the whole cache. We assumed the ECBs of each task to be consecutively arranged starting at a random cache set $S \in [0; CS - 1]$, i.e. from S to $S + |ECB| \bmod CS$.
- For each task, the UCBs were generated according to a uniform distribution ranging from 0 to the number of ECBs⁵ times a reuse factor: $[0, RF \cdot |ECB|]$. The factor RF was used to adapt the assumed reuse of cache-sets to account for different types of real-time applications, for example, from data processing applications with little reuse up to control-based applications with heavy reuse.

Staschulat's approach exploits the fact that for the i -th pre-emption only the i -th highest number of UCBs has to be considered. As our case study and other measurements (Bertogna et al. 2011) have shown, a significant reduction typically only occurs at a high number of pre-emptions. For the purposes of evaluation, for Staschulat's approach, we simulated what in practice is likely to be an optimistic reduction: reducing the number of UCBs per pre-emption by one each time.

In each experiment the taskset utilization not including pre-emption cost was varied from 0.025 to 0.975 in steps of 0.025. For each utilization value, 1000 tasksets were generated and the schedulability of those tasksets determined using the appropriate pre-emption cost computation integrated into response time analysis.

7.1 Base configuration

We conducted experiments varying the number of tasks, the cache-size (i.e. number of cache-sets (CS)), the block reload time (BRT), the cache utilization (CU) and the reuse factor (RF). As a base configuration we used the default values of 10 tasks, a cache of 256 cache-sets, a block-reload time of 8 μ s, a reuse factor of 30 % and a cache-utilization of 10. The latter two parameters were chosen according to the actual values observed in the case-study. Figure 9 illustrates the performance of the different approaches for this base configuration. The graph also shows a line marked Simulation-UB. This refers to the use of simulation to form a necessary schedulability test. We simulated execution and pre-emption of the tasks starting from near simultaneous release. The tasks were released in order, lowest priority first, to increase the number of pre-emptions considered. We assumed that any partial execution of a task uses all its ECBs and UCBs. If any task missed its deadline, then the taskset was proven to be unschedulable w.r.t. the pre-emption cost model used, thus providing a valid upper bound on taskset schedulability including pre-emption costs. Note that the lines on the graphs appear in the same order as they are described in the legend. The graphs are best viewed online in colour. For each approach, we determined the average breakdown utilization for the tasksets generated for the base configuration, see Table 3. These results show that the ECB-Union, and the Multiset approaches significantly improve upon the performance of previous methods.

⁵The set of ECBs generated may exceed the number of cache sets. In this case, the set of ECBs used within the response time analysis is limited to the number of cache sets. However, for the generation of the UCBs, the original set of ECBs is used.




Fig. 9 Evaluation of base configuration. Number of tasksets deemed schedulable at the different total utilizations

Table 3 Average breakdown utilization and weighted schedulability measure of base configuration tasksets for each approach

Analysis	Average breakdown utilization	Weighted schedulability
No Pre-emption Cost	0.93	0.86
Simulation	0.81	0.62
Combined	0.70	0.50
ECB-Union Multiset	0.66	0.46
UCB-Union Multiset	0.67	0.46
ECB-Union	0.62	0.42
UCB-Union	0.59	0.37
UCB-Only	0.57	0.36
ECB-Only	0.40	0.20
Staschulat	0.33	0.32

Exhaustive evaluation of all combinations of cache and taskset configuration parameters is not possible. We therefore fixed all parameters except one and varied the remaining parameter in order to see how performance depends on this value. The graphs below show the weighted schedulability measure $W_y(p)$ (Bastoni et al. 2010) for schedulability test y as a function of parameter p . For each value of p , this measure combines data for all of the tasksets τ generated for all of a set of equally spaced utilization levels. Let $S_y(\tau, p)$ be the binary result (1 if schedulable, 0 otherwise) of schedulability test y for a taskset τ and parameter value p then:

$$W_y(p) = \left(\sum_{\forall \tau} u(\tau) \cdot S_y(\tau, p) \right) / \sum_{\forall \tau} u(\tau) \quad (20)$$

where $u(\tau)$ is the utilization of taskset τ . This weighted schedulability measure reduces what would otherwise be a 3-dimensional plot to 2 dimensions (Bastoni et al.




Fig. 10 Weighted schedulability measure; varying cache utilization from 0 to 20, in steps of 2

2010). Weighting the individual schedulability results by taskset utilization reflects the higher value placed on being able to schedule higher utilization tasksets.

7.2 Cache utilization & cache-reuse

Cache utilization and cache-reuse are the most important factors for pre-emptively scheduled systems. If all tasks fit into the cache, i.e. the cache utilization is less than one or there is no cache-reuse at all, then no additional cache related pre-emption delays occur. The other extreme is when each task completely fills the cache. In this case, each UCB must be assumed to be evicted, and hence the overall pre-emption delay depends solely on the number of UCBs. Figure 10 shows the weighted schedulability measure for each approach as a function of the cache utilization. At a low cache utilization, only a few UCBs are actually evicted. The set of ECBs per task is low, and often smaller than the number of UCBs of all possibly pre-empted tasks. Thus, an upper bound on the possibly evicted UCBs per pre-empting task (as computed by the UCB-Union approach) is slightly pessimistic, while the ECB-Union approach is in this case more pessimistic. The situation changes with increased cache utilization. As each task uses a larger proportion of the whole cache on average, the UCB-Union approach becomes significantly more pessimistic than the ECB-Union approach. The same relation can be found in the case of the Multiset approaches, which—due to the reduced pessimism—perform significantly better than the basic union-approaches. Figure 11, shows the weighted schedulability measure for each approach as a function of the reuse factor. At low values of the reuse factor, the set of UCBs per task is low compared to the ECBs, and so the UCB-Union method is more pessimistic than the ECB-Union method, while at high values of the reuse factor, the opposite applies as the set of UCBs for each task becomes similar to its set of ECBs. The same behaviour is observed in the case of the Multiset approaches, in Fig. 11 these two lines cross at a medium level of reuse, while the Combined approach outperforms both, providing the best performance in all cases. Since the reuse factor only affects the number of UCBs, the performance of the ECB-Only approach is in-




Fig. 11 Weighted schedulability measure; varying reuse factor from 0 % to 100 %, in steps of 10 %




Fig. 12 Weighted schedulability measure; varying number of tasks from $2 = 2^1$ to $2^6 = 64$

dependent of the reuse factor. As expected, performance of the ECB-Only approach is relatively poor at low levels of reuse, but competitive at high levels.

7.3 Number of tasks

In this experiment, we varied the number of tasks with the other parameters fixed at their default values. Figure 12 shows that the more tasks there are, the less likely a taskset of a given utilization is to be schedulable. This is because with an increased number of tasks the number of pre-emptions and hence the overall pre-emption costs increase, reducing the schedulability of the taskset. Note that the pre-emption costs are not included in the taskset utilization. This reduction in schedulability with increasing taskset size holds for all of the approaches, with a greater reduction observed with Staschulat's approach for the reasons explained in Sect. 6.

Note that the upper bound derived by simulation shows a much smaller reduction. This is because, as the number of tasks increases, the number of possible execution




Fig. 13 Weighted schedulability measure; varying number of cache sets from $2^5 = 64$ to $2^{10} = 1024$




Fig. 14 Weighted schedulability measure; varying block reload time from $2^0 = 1 \mu s$ to $2^4 = 32 \mu s$

scenarios increases rapidly, thus it becomes less likely that the simulation will deliver the worst-case scenario.

7.4 Cache-size

The number of cache-sets also has an influence on the overall performance of the different approaches. The more cache-sets there are, the higher the impact of a pre-emption may be, given the same cache utilization and block reload time. Hence as the number of cache sets is increased, all of the approaches show a similar decrease in schedulability with the exception of the basic response time analysis which does not include pre-emption costs, see Fig. 13. Varying the block reload time results in similar behaviour, see Fig. 14. We note that when increasing the cache size, the execution time of each task might also be expected to decrease. In this experiment; however, we keep WCETs constant and examine only the effect on schedulability of changing the cache size.




Fig. 15 Weighted schedulability measure; varying the range of task periods $[5, 50]$ to $[5, 5 \times 10^4]$

7.5 Range of task periods

The range of task periods may also influence the performance of the different approaches. We therefore conducted experiments varying the task period generation. Our base configuration used task periods in the range 5 ms to 500 ms, typical of many real-time systems. In Fig. 15, we varied the number of orders of magnitude v spanning the minimum to the maximum task period and hence the range of task periods and deadlines given by $5[1, 10^v]$ ms. In Fig. 15 we see that as the range of task periods is increased, by making the maximum period larger, schedulability improves for all of the approaches. This is because increasing the maximum period has the effect of reducing the proportion of task that are generated with smaller periods (e.g. in the range 1–10 ms). Given that the block reload time is constant in this experiment, the ratio of pre-emption costs to taskset utilization reduces for increasing ranges of task periods, thus improving schedulability (recall that pre-emption costs are not included in taskset utilization). Note the smaller improvement with no pre-emption costs is a property of fixed priority scheduling; which on average can schedule higher utilization tasksets when there is a wide disparity in task periods.

In Fig. 16, we varied the scaling factor w from 1 to 10 and hence the range of task periods given by $w[1, 100]$ ms. Given that the block reload time is again constant in this experiment, the ratio of pre-emption costs to taskset utilization decreases as the task periods, deadlines and execution times are all scaled up, thus increasing schedulability for all of the approaches that include pre-emption costs. In fact these results are similar to the ones for varying block reload times, but with the results for larger values of the scaling factor w corresponding to those for smaller block reload times.

8 Conclusions

The major contribution of this paper is the introduction of new methods of bounding pre-emption costs; the ECB-Union approach and the ECB-Union Multiset and




Fig. 16 Weighted schedulability measure; varying the scale of task periods $w[1, 100]$ from $w = 2$ to $w = 10$

UCB-Union Multiset approaches. The ECB-Union Multiset approach dominates the ECB-Union approach, which in turn dominates the UCB-Only approach of Lee et al. (1998). Similarly, the UCB-Union Multiset approach dominates the UCB-Union approach of Tan and Mooney (2007), which in turn dominates the ECB-Only approach of Busquets-Mataix et al. (1996) and Tomiyama and Dutt (2000). The Multiset approaches removes pessimism present in both the ECB-Union and UCB-Union approaches. The ECB-Union Multiset and UCB-Union Multiset approaches are incomparable and so we combined them into a composite response time test that dominates both.

We extended the ECB-Union (Multiset) and UCB-Union (Multiset) approaches to systems that permit tasks to access shared resources in mutual exclusion according to the Stack Resource Policy. Our work in this area revealed that previous approaches to computing pre-emption delays, although including blocking factors in their schedulability analyses, did not account for the pre-emption of blocking tasks during a resource access. This omission can lead to optimistic (unsound) response times, an issue that we corrected.

Finally, we examined the performance of the various approaches to computing pre-emption costs via a case study and an empirical evaluation of taskset schedulability. The latter showed that a combined response time analysis test using both the new ECB-Union Multiset and UCB-Union Multiset approaches derived in this paper provides an effective method of determining task schedulability. This combined approach offers a significant improvement in performance over previous approaches for a wide range of different task and cache configurations, including cache utilization level, amount of reuse, cache size, and block reload times.

Acknowledgements This research and collaboration came about as a result of the 1st Real-Time Scheduling Open Problems Seminar (RTSOPS 2010) (see Altmeyer and Burguière 2010). This work was partially funded by the UK EPSRC funded Tempo project (EP/G055548/1), the Transregional Collaborative Research Center AVACS of the German Research Council (DFG) and the EU funded ArtistDesign Network of Excellence. We also like to thank Alan Burns, Jack Whitham and the anonymous reviewers for their comments on an earlier draft of this paper.

References

- Altmeyer S, Burguière C (2009) A new notion of useful cache block to improve the bounds of cache-related preemption delay. In: Proceedings ECRTS, pp 109–118
- Altmeyer S, Burguière C (2010) Influence of the task model on the precision of scheduling analysis for preemptive systems. In: Proceedings RTSOPS, pp 5–6
- Altmeyer S, Maiza C (2011) Cache-related preemption delay via useful cache blocks: survey and redefinition. *J Syst Archit* 57:707–719
- Altmeyer S, Maiza C, Reineke J (2010) Resilience analysis: tightening the crpd bound for set-associative caches. In: Proceedings LCTES, New York, NY, USA. ACM Press, New York, pp 153–162
- Altmeyer S, Davis R, Maiza C (2011a) Pre-emption cost aware response time analysis for fixed-priority pre-emptive systems. Tech rep. Available from <http://www-users.cs.york.ac.uk/~robdavis/>
- Altmeyer S, Davis I, Maiza C (2011b) Cache related pre-emption aware response time analysis for fixed priority pre-emptive systems. In: Davis I, Fisher N (eds) Proceedings of the 32nd IEEE real-time systems symposium (RTSS'11), pp 261–271
- Audsley N, Burns A, Richardson M, Tindell K, Wellings AJ (1993) Applying new scheduling theory to static priority pre-emptive scheduling. *Softw Eng J* 8:284–292
- Baker TP (1991) Stack-based scheduling for realtime processes. *Real-Time Syst* 3:67–99
- Bastoni A, Brandenburg B, Anderson J (2010) Cache-related preemption and migration delays: empirical approximation and impact on schedulability. In: Proceedings OSPERT, pp 33–44
- Bertogna M, Buttazzo G, Marinoni M, Yao G, Esposito F, Caccamo M (2010) Preemption points placement for sporadic task sets. In: Proceedings ECRTS, pp 251–260
- Bertogna M, Xhani O, Marinoni M, Esposito F, Buttazzo G (2011) Optimal selection of preemption points to minimize preemption overhead. In: Proceedings of the 2011 23rd Euromicro conference on real-time systems. ECRTS, vol 11. IEEE Computer Society, Washington, pp 217–227
- Bini E, Buttazzo G (2005) Measuring the performance of schedulability tests. *Real-Time Syst* 30:129–154
- Burguière C, Reineke J, Altmeyer S (2009) Cache-related preemption delay computation for set-associative caches—pitfalls and solutions. In: Proceedings WCET
- Busquets-Mataix JV, Serrano JJ, Ors R, Gil P, Wellings A (1996) Adding instruction cache effect to schedulability analysis of preemptive real-time systems. In: Proceedings RTAS, pp 204–212
- Davis R, Merriam N, Tracey N (2000) How embedded applications using an rtos can stay within on-chip memory limits. In: Work progress session RTSS
- Davis R, Zabus A, Burns A (2008) Efficient exact schedulability tests for fixed priority real-time systems. *IEEE Trans Comput* 57:1261–1276
- Gustafsson J, Betts A, Ermedahl A, Lisper B (2010) The Mälardalen WCET benchmarks—past, present and future OCG, Brussels, pp 137–147
- Joseph M, Pandya P (1986) Finding response times in a real-time system. *Comput J* 29(5):390–395
- Keskin U, Bril R, Lukkien J (2010) Exact response-time analysis for fixed-priority preemption-threshold scheduling. In: Proceedings work-in-progress session ETFA
- Lee CG, Hahn J, Seo YM, Min S, Ha R, Hong S, Park Y, Lee M, Kim CS (1998) Analysis of cache-related preemption delay in fixed-priority preemptive scheduling. *IEEE Trans Comput* 47(6):700–713
- Lundqvist T, Stenström P (1999) Timing anomalies in dynamically scheduled microprocessors. In: Proceedings RTSS, p 12
- Marinho J, Nélis V, Petters SM, Puaut I (2012) Preemption delay analysis for floating non-preemptive region scheduling. In: DATE 2012, pp 497–502
- Martin S, Minet P, George L (2007) Non pre-emptive fixed priority scheduling with fifo arbitration: uniprocessor and distributed cases. Tech rep, INRIA Rocquencourt
- Meumeu Yomsi P, Sorel Y (2007) Extending rate monotonic analysis with exact cost of preemptions for hard real-time systems. In: Proceedings of the 19th Euromicro conference on real-time systems. IEEE Computer Society, Washington, pp 280–290
- Petters SM, Farber G (2001) Scheduling analysis with respect to hardware related preemption delay. In: Workshop on real-time embedded systems
- Ramaprasad H, Mueller F (2006) Tightening the bounds on feasible preemption points. In: Proceedings RTSS, pp 212–224
- Regehr J (2002) Scheduling tasks with mixed preemption relations for robustness to timing faults. In: Proceedings RTSS, pp 315–325
- Schneider J (2000) Cache and pipeline sensitive fixed priority scheduling for preemptive real-time systems. In: RTSS'2000, pp 195–204

- Sha L, Rajkumar R, Lehoczky JP (1990) Priority inheritance protocols: an approach to real-time synchronization. *IEEE Trans Comput* 39:1175–1185
- Staschulat J, Schliecker S, Ernst R (2005) Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In: *Proceedings ECRTS*, pp 41–48
- Tan Y, Mooney V (2007) Timing analysis for preemptive multi-tasking real-time systems with caches. *ACM Trans Embed Comput Syst* 6(1):1210275
- Tomiyama H, Dutt ND (2000) Program path analysis to bound cache-related preemption delay in preemptive real-time systems. In: *Proceedings CODES*, pp 67–71
- Wang Y, Saksena M (1999) Scheduling fixed-priority tasks with pre-emption threshold. In: *Proceedings RTCSA*, pp 328–338



Sebastian Altmeyer received his M.Sc. in computer science at the Saarland University, Germany in 2006. He is currently a Ph.D. student and research assistant in the team of Prof. R. Wilhelm, department of computer science, Saarland University, Germany. His research interests include embedded systems, timing and scheduling analysis and abstract interpretation.



Robert I. Davis is a Senior Research Fellow in the Real-Time Systems Research Group at the University of York, and a Non-executive Director of Rapita Systems Ltd. He received his D.Phil. in Computer Science from the University of York in 1995. Since then he has founded three startup companies, all of which have succeeded in transferring real-time systems research into commercial products. Robert's research interests include scheduling algorithms and schedulability analysis for real-time systems and networks.



Claire Maiza is an assistant professor of Computer Science at Grenoble INP (France) where she has been employed since 2010. She received her Ph.D. degree in computer science from the university of Toulouse, France in 2008. From September 2008 to August 2010, she was a postdoc in the team of Prof. R. Wilhelm, department of computer science, Saarland University, Germany. Her research activities are with the VERIMAG laboratory in the “synchrone” group. Her research interests include timing analysis, abstract interpretation, pre-emptive systems analysis, architecture of real-time systems and model of hardware and extra-functional properties (energy consumption).