Cache-Aware Real-Time Scheduling Simulator: Implementation and Return of Experience

Hai Nam Tran, Frank Singhoff, Stéphane Rubini, Jalil Boukhobza
Univ. Bretagne Occidentale, UMR 6285, Lab-STICC, F-29200 Brest, France
{hai-nam.tran,singhoff,rubini,boukhobza}@univ-brest.fr

ABSTRACT

Evaluating cache related preemption delay (CRPD) in preemptive scheduling context for Real-Time Embedded System (RTES) stays an open issue despite of its practical importance. Indeed, various parameters should be taken into account such as cache utilization of tasks, memory layout of tasks, processor utilization and priority assignment algorithm. In state-of-the-art work, dependencies amongst those parameters are not investigated with precision because of the lack of scheduling analysis tool taking them into account. In this article, we propose a tool to investigate and evaluate scheduling analysis of RTES with cache memory and various scheduling parameters. Both modeling guidelines and implementation is detailed. Implementation is made in Cheddar, an open-source scheduling analyzer, which is freely available; Experiments are conducted in order to illustrate performance and applicability of our tool. Furthermore, we discuss about implementation issues, problems raised and lessons learned from those experiments.

1. INTRODUCTION

Cache memory is a crucial hardware component used to reduce the performance gap between processors and main memories. In the context of real-time embedded system (RTES), the popularization of processors with large size and multi level caches motivates the proposition of verification methods handling this hardware component [14], [5], [2].

Scheduling simulation is a classical verification method performed at design step of RTES. It provides means to investigate that all timing constraints of a RTES are satisfied. To perform scheduling simulation, several assumptions are usually made in order to simplify system modeling. One of them is that preemption costs are negligible. *Preemption cost* is the additional time to process interrupts, manipulate task queues and actually perform context switches.

Integrating cache memory in RTES generally enhances the whole performance in term of execution time, but unfortunately it can lead to an increase in preemption cost and execution time variability [15]. When a task is preempted, memory blocks belonging to the task could be removed from the cache. Once the task resumes, previously removed memory blocks have to be reloaded into the cache. Thus, a new preemption cost named *Cache Related Preemption Delay* (CRPD) is introduced. By definition, CRPD is the additional time to refill the cache with memory blocks evicted

by the preemption. In [15], the authors showed that the cost of context switching raises from the range of $4.2\mu s$ to $8.7\,\mu s$ to the range of $38.6\mu s$ to $203.2\mu s$ when the size of the data sets of programs is larger than the cache. Thus, taking CRPD into account may be important when performing scheduling analysis of a RTES.

Problem statement: Scheduling analysis of RTES with cache memory in fixed priority preemptive scheduling context is complex because there are many parameters affecting the outcome and dependencies amongst them. For example, CRPD analysis, which is the process of evaluating the impact of CRPD on a RTES, cannot be only based on single-task analysis as it depends on correlations amongst tasks. For a set of tasks, CRPD analysis has to take into account parameters such as cache utilizations of tasks, memory layout of tasks, processor utilization but also priority assignment algorithms and WCET tasks.

To address all parameters, many tools are involved, in two domains: WCET/Cache analysis and scheduling analysis. Each of those domains is only dedicated to a sub-part of those parameters. Unfortunately, there are no tools addressing the whole problem in state-of-the-art work.

Contributions : We propose a tool to investigate and evaluate scheduling analysis of RTES with cache memory and various scheduling parameters. The work is implemented in Cheddar, an open-source scheduling analyzer, which is freely available to researchers and practitioners. We propose an approach to use the analysis models and results of WCET/cache analysis tools into a scheduling analysis tool. The programming model we used is compliant with the existing one in [14], [5], [2] and [17]. Experiments are performed to illustrate performance and applicability of our tool.

The rest of the article is organized as follows. Section 2 presents background of our work. In Section 3, we give an overview of our approach. In Section 4, we present development process and detailed information about the implementation of our work. In Section 5, an evaluation of our proposed scheduling simulator is given. Section 6 discusses related works and Section 7 concludes the article.

2. BACKGROUND

In this section, we introduce the system model and we explain how preemption cost and CRPD are computed. We assume an uniprocessor RTES with direct-mapped instruction cache. As far as we know, instruction cache is popular in practical implementation of RTES. The assumption about direct-mapped is used to simplify the data flow analysis, it could be easily relaxed. There are n independent tasks,
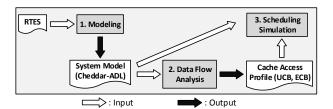
Figure 1: Approach

$\tau_1, \tau_2, ..., \tau_n$ scheduled by a preemptive scheduler. CPRD is bounded by $g \cdot BRT$, where g is an upper bound on the number of cache block reloaded due to preemption, and BRT is an upper-bound on the time necessary to reload a memory block in the cache (block reload time). To analyze the effect of preemption on a preempted task, Lee et al.[14] introduced the concept of useful cache block (UCB):

Definition 1. A memory block m is called a useful cache block (UCB) at program point P, if m may be cached at P and m may be reused at program point P' after P that may be reached from P without eviction of m on this path.

The number of UCB at program point P gives an upper bound on the number of additional reloads due to a preemption at P. The maximum possible preemption cost for a task is determined by the program point with the highest number of UCB. In [22], the authors exploit the fact that for the i-th preemption, only the i-th highest number of UCB has to be considered. However, as shown in [1] and [3], a significant reduction typically only occurs at a high number of preemptions. Thus, we only consider the program point with highest number of UCB.

The impact of preempting task is given by the number of cache blocks that the task may evict during its execution. Busquet et al.[5] introduce the concept of evicting cache block (ECB):

Definition 2. A memory block of the preempting task is called an evicting cache block (ECB), if it is accessed during the execution of the preempting task.

The notation UCB_i and ECB_i are used to present the set of UCBs and ECBs of a task τ_i. Assume that the sets of UCB and ECB of each task are preliminary computed, UCB'_i is the set of UCBs currently in the cache of the preempted task. $\gamma_{i,j}$ is the preemption cost (i.e. the CRPD) when a task τ_j directly preempt task τ_i. In case of a preemption between two tasks τ_i and τ_j, $\gamma_{i,j}$ is computed by:

$$\gamma_{i,j} = \mathrm{BRT} \cdot \mid UCB'_i \cap ECB_j \mid \qquad (1)$$

However, in case of nested preemptions, a task τ_j can preempt more than one task. Thus, computation of CRPD must take all preempted tasks into account.

3. OUR APPROACH

Our approach consists of three steps as shown in Fig. 1. First, we modeled a RTES with components required to apply analysis methods stated in Section 2. Those components include hardware and software parts. Hardware components are processor, core and cache memory. Software components are tasks and control flow graphs of tasks.

Second, from this model, we applied data flow analysis presented in [14] in order to compute the set of UCB and

ECB of each task, which is called cache access profile in the sequel. A detailed description of these first and second steps could be found in [24], which is a preliminary work toward cache integration.

Third, system model with computed cache access profiles are loaded into the scheduling simulator. Scheduling simulation is done and provides various outcomes such as feasibility of the system, worst case response time of tasks, number of preemption, CRPD per task and total CRPD.

In our work, the cache utilization of a program at step 1 is modeled at both low level, which is the control flow graph of program produced by a WCET analysis tool, and at high level, which is a pre-computed set of UCB and ECB. We expect the user to re-use results produced by a WCET analysis tool within Cheddar. We can extract all information related to CRPD from the scheduling simulator as written above to fully investigate the impact of CRPD. Detailed implementation of each step is presented in the next section.

4. IMPLEMENTATION

In this section, we present the implementation of our approach. We introduce our framework, discuss about the development process and point out several implementation issues. The source code of the presented work is available under GNU GPL licence at `http://beru.univ-brest.fr/svn/CHEDDAR/trunk/src/`.

The work took place in the context of the Cheddar project [21]. Cheddar is an open source real-time scheduling analysis tool. Classical feasibility tests, scheduling algorithms and scheduling simulator for real-time systems are implemented in Cheddar. System architecture is defined with Cheddar Architecture Description Language (Cheddar-ADL). Cheddar class files are automatically generated by the tool Platypus [20] through a model-driven process. The Cheddar meta-model defines hardware components such as: *processor*, *core* and *shared_resource*; and software components such as: *task* and *task_group* [11].

The development process consisted of three steps. First, Cheddar-ADL is extended to model RTES with cache memory and cache access profile. Second, we implemented UCB computation by data flow analysis. Third, a cache-aware scheduling simulator is implemented by extending the scheduling simulator of Cheddar. In [24], we presented how Cheddar-ADL is extended to take into account cache memory and how data flow analysis in [14] is implemented. In this article, we provide a summary of our models and focus more on the implementation of the cache-aware scheduling simulator and its application.

4.1 Cache Model and Cache Access Profile

To support cache aware scheduling analysis, the Cheddar meta-model has been extended with the entities below:

```
1   ENTITY Generic_Cache
2       cache_size : Natural;
3       line_size : Natural;
4       associativity : Natural;
5       miss_time : Natural;
6       cache_category : Cache_Type;
7   ENTITY Generic_Task
8       task_type : Natural;
9       capacity : Natural;
10      deadline : Natural;
11      priority : Priority_Range;
12      cache_access_profile_name : String;
```

```
13        cfg_name : String;
14    ENTITY Cache_Access_Profile
15        UCBs : Cache_Blocks_Table;
16        ECBs : Cache_Blocks_Table;
17    ENTITY CFG
18        nodes : CFG_Nodes_Table;
19        edges : CFG_Edges_Table;
20    ENTITY Basic_Block SUBTYPE OF (CFG_Node);
21        instruction_offset : INTEGER;
22        instruction_capacity : INTEGER;
```

Cache memory is modeled by classical attributes such as cache size, line size, associativity and block reloading time or miss time. From this model of cache and control flow graph of a task, we can compute the memory layout, the memory-to-cache mapping and the set of UCB and ECB of a task.

A task is modeled by an entity with classical attributes such as *capacity*, *deadline*, *priority* and *offsets*. The *task_type* specifies the type of a task, such as *periodic*, *aperiodic*, *sporadic* or *poisson* [11].

Two attributes, *UCBs* and *ECBs* are added to handle caches. They are the set of UCBs and ECBs of a task, respectively. The model is designed according to two assumptions: first, a task is assumed to have a fixed set of UCB and ECB. UCB is computed by the program point with highest number of UCB ; second, any partial execution of a task uses all its ECBs and UCBs ; thus, CRPD achieved from scheduling simulation is always upper-bounded.

In order to compute the set of UCB and ECB of a task, we need its control flow graph (which is modeled by the CFG EXPRESS entity). Each node in the CFG stores information about the capacity and the location in main memory of its assembly instruction. At the moment, our analysis method takes into account direct mapped instruction cache because instruction access pattern is simpler to be computed from the CFG of a program.

4.2   Computing Cache Access Profile

UCB analysis is implemented following the work in [14]. The input of the algorithm is an annotated CFG of a program. From the input CFG, our tool analyzes and computes the set of UCB and ECB.

4.3   The Cache-Aware Scheduling Simulator

The scheduling simulation in Cheddar works as follows. First, a system architecture model, including hardware-/software components, is loaded. Then, the scheduling is computed by three successive steps: computing priority, inserting ready task into queues and electing task. The elected task will receive the processor for the next unit of time.

The scheduling simulator records different events raised during the simulation, such as task releases, task completions and shared resources lockings or unlockings. The result of the scheduling analysis is the set of events produced at simulation time.

We extended the scheduling simulator of Cheddar as follows. First, we extended the set of events Cheddar can produce. For example, an event PREEMPTION, which is raised when a preemption occurs, is added. Second, event RUNNING_TASK, which is raised when a task executes, is extended with the assumption about CRPD that any partial execution of a task uses all its ECBs and UCBs.

The pseudo code of the event handler is written below. The notation τ_i.cUCB represents the set of UCBs of task τ_i in the cache. It is computed from a system model at scheduling simulation time. The function Remove() at line 8 is used to remove an element from a set.

```
1    event SCHED_START
2        for each task τi loop
3            τi.cUCB ← τi.UCB
4        end loop
5    event PREEMPTION
6        τj ← preempting_task
7        for each task τi preempted loop
8            τi.cUCB ← Remove(τi.cUCB, τj.ECB)
9        end loop
10   event RUNNING_TASK
11       τi ← executing_task
12       CRPD ← (τi.UCB − τi.cUCB) * Miss_Time
13       τi.cUCB ← τi.UCB
```

At the start of the scheduling simulation, a SCHED_START event is raised. The WCET of a task is assumed to include the block reloading time when a task is executed non-preemptively. So, on event SCHED_START, the set of UCBs of a task is assumed to be filled.

When a preemption occurs, a PREEMPTION event is raised and the simulator computes the evicted UCBs of preempted tasks using the ECBs of preempting task. The scheduler keeps track of the number of UCBs for each task.

When a task executes, a RUNNING_TASK event is raised. The scheduler firstly checks if all the UCBs of this task are loaded into the cache. If so, the task continues its execution. If not, the task reloads the evicted UCBs. The CRPD is added to the remaining capacity of the task itself. In our implementation, CRPD is not added to the capacity of preempted tasks at the preemption point but at the instant at which those tasks resume execution.

4.4   Implementation Issues

Several issues were raised when designing and implementing the simulator. Most of them were raised because we need to mix timing specifications of different orders of magnitude. Others are related to tools interoperability.

In practice, cache block reload time is significantly smaller than period or capacity of a task. In Cheddar, we do not prescribe 1 unit of time is equivalent to 1 ms or 1 μs, which are the unit of task period and block reload time. Mixing timing specifications of different orders of magnitude makes complex the computation of the *feasibility interval*. We recall that a feasibility interval is an interval for which testing of task feasibility is needed [10]. The scheduling simulation interval needed to verify the schedulability of a task set could be significantly large if a μs is chosen as a time unit. A solution in practice is to design a system with harmonic task set in order to reduce the feasibility interval; however, it is clearly not always possible. In addition, instead of using 1 μs, we use the cache block reload time as a base value for 1 unit of time as in our experiment. Furthermore, a long scheduling simulation interval also raises issues regarding performance and scalability. Even with harmonic task set, the tool must be able to perform scheduling simulation in a large interval to overcome the different between cache block reload time and task period, which may be CPU and memory expensive. As Cheddar stores scheduling simulation results into XML files, it can also be I/O intensive. To reduce memory and I/O overhead, we selected a subset of events the simulator has to handle and store.

A second issue we are facing is about tool interoperability.

The input data of the CRPD analysis in our tool is designed to be compatible with data provided by a WCET analysis tool. We also support import data in XML format. At the moment, we do not enforce tool interoperability and we expect to investigate WCET tools in order to overcome this issue.

5. EXPERIMENT AND DISCUSSION

In this section, we show that our tool can handle parameters compliant with the existing works in [14], [5], [2] in the first experiment. In addition, we discuss about the dependency between CRPD and scheduling parameters. Furthermore, we point out that our tool can run CRPD optimization techniques by taking an example of memory layout optimization by simulated annealing following the work of [17] in the second experiment. We also provide performance and scalability tests of the tool in the third experiment.

Task period and cache utilization generation of our experiments is based on the existing work in [1]. Task sets are generated with the following configuration. Task periods are uniformly generated from 5 ms to 500 ms, as found in most automotive and aerospace hard real-time applications [1]. Generated task sets are harmonic in order to have a low feasibility interval and scheduling simulation period. Task deadlines are implicit, i.e. $\forall i : D_i = T_i$. Processor utilization values (PU) are generated using the UUniFast algorithm [4]. Task execution times are set based on the processor utilizations and the generated periods: $\forall i : C_i = U_i \cdot T_i$, where U_i is the processor utilization of task i. Task offsets are uniformly distributed from 1 to 30 ms.

Cache memory and cache access profile of tasks are generated as follows. The cache is direct mapped. The number of cache blocks is equal to 256. The block-reload time is 8 μs. The cache usage of each task is determined by the number of ECBs. They are generated using UUniFast algorithm for a total cache utilization (CU) of 5. UUniFast may produce values larger than 1 which means a task fills the whole cache. ECBs of each tasks are consecutively arranged from a cache set. For each task, the UCBs are generated according to a uniform distribution ranging from 0 to the number of ECB multiplied by a reuse factor (RF). If the set of ECBs generated exceeds the number of cache sets, the set of ECBs is limited to the number of cache sets. For the generation of the UCBs, the original set of ECBs is used.

5.1 CRPD with Priority Assignment and Processor Utilization (PU)

In this experiment, we present CRPD analysis with different priority assignments or scheduling algorithms. In addition, we discuss about the impact of changing priority assignment/scheduling algorithm and increasing PU to CRPD.

PU is varied from 0.5 to 0.95 in steps of 0.05. RF is fixed at 0.3. For each value of PU, we performed scheduling simulations with 100 task set and computed the average number of preemptions and average total CRPD in a scheduling interval of 1000 ms. Experiments are conducted with two priority assignment algorithms: Rate Monotonic (RM) and one named PA*, which assigns the highest priority level to the task with the largest set of UCB. In addition, we also take into account Earliest Deadline First (EDF) scheduler.

The result of this experiment is Fig. 2. As the graph illustrates, the number of preemptions and the preemption cost increases steadily from the processor utilization of 50%
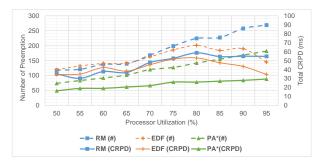

Figure 2: Varying PU, RF=0.3

to 80%. After this point, there is a downward trend in the preemption cost and in the number of preemptions of EDF while there is an upward trend in those data of RM and PA*. Observed from the scheduler, when PU is larger than 80%, many task sets are not schedulable.

In conclusion, first, when PU increases, the total number of preemption and CRPD also increase. However, the change is not linear. Second, a priority assignment algorithm with less number of preemption tends to give lower total CRPD. EDF and PA* generate less preemption and CRPD than RM. In fact, to enforce the fixed priority order, the number of preemptions that typically occur under RM is higher than under EDF [6]. From this experiment, we see that CRPD depends on the chosen priority assignment or scheduler.

In addition, this experiment shows that both scheduling analysis and CRPD analysis should be performed jointly. PA*, a priority assignment taking CRPD into account has a significant lower total CRPD. The decrease in total CRPD of PA* with RM and EDF is roughly 30ms on a scheduling interval of 1000ms. However, comparing to RM and DM, feasibility constraints of tasks are not satisfied with PA*, only total CRPD is reduced. In [25], we proposed a priority assignment heuristic to take into account both feasibility constraints and CRPD.

5.2 CRPD with Memory Layout Optimization by Simulated Annealing

The objective of this experiment is to show that users can use CRPD optimization approaches with our tool. We apply memory layout optimization by simulated annealing (SA) based on the work of [17] with our generated task sets. In our experiment, the objective of SA is to lower the total CRPD after a scheduling simulation over a scheduling interval of 1000 ms.

For each iteration of SA, we perform a swap in memory layout between two random tasks. Changes are made to the layout of tasks in memory, and then mapped to their cache layout for evaluation. The total CRPD is computed by scheduling simulation. The optimum layout is the layout which has the lowest total CRPD. Initial temperature of SA is 1.0, and after every iteration, it is reduced by multiplying it by a cooling rate of 0.5 until it reaches the target temperature of 0.2. Number of iteration for each temperature is 10.

The result of this experiment is Fig. 3. From the graph, we can see the impact of memory layout optimization to total CRPD. We can reduce roughly 20-30 ms of total CRPD. To sum up, this experiment shows that our tool allows users to perform a specific optimization of CRPD.
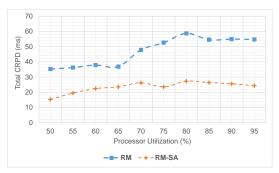
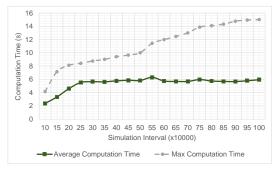**Figure 3: Total CRPD with and without memory layout optimization**

Figure 4: Computation time of the simulator

5.3 Performance/Scalability Analysis

The objective of this experiment is to test the performance and the scalability of our tool when scheduling simulation interval increases. In general, there are four factors affecting the performance of a scheduling simulator: (1) the number of tasks, (2) the scheduling simulation interval, (3) the cache size and (4) the number of events. The first three factors depend on a chosen RTES. The number of events depends on characteristics of the RTES; for example, a higher processor utilization means a higher number of preemption events. In this experiment, we choose to test a system model of 10 tasks and 256 cache blocks. Processor utilization is set to 70 %. Scheduling simulation is ranging from 100,000 to 1,000,000 units of time where 1 unit = 8 μs.

Fig. 4 displays results of our experiment on a PC with Intel Core i5-3360 CPU, running Ubuntu 12.04. For each simulation interval, 100 task sets are generated. We perform scheduling simulation and compute the max and average computation time.

As we can see, while maximum computation time increases slightly when simulation interval increases, average computation time only fluctuates around 6 seconds. This shows that the tool is scalable when simulation interval is high.

6. RELATED WORKS

In this section, we present several real-time scheduling analysis and WCET analysis tools.

MAST[12] is a modeling and analysis suite for real-time applications. The hardware component abstraction of MAST model is generic and it includes processing resources and shared resources. The shared resource component is not supposed to model a cache memory unit. However, MAST considers the overhead parameters of the components that may be used to model CRPD.

STORM[26] and YARTISS [7] are scheduling simulation tools mainly designed for evaluating and comparing real-time scheduling algorithm for multiprocessor architectures.

SymTA/S[13] and RealTime-at-Work [1] are model-based scheduling analysis tools targeting automotive industry. The hardware components supported in those tools are specific to their domains (ECU, CAN and AFDX Networks).

To the best of our knowledge, the support for cache memory does not exists in the tools above.

SimSo[8] is a scheduling simulation tool. It supports cache sharing on multi-processor systems. It takes into account impact of the caches through statistical models and also the direct overheads such as context switches and scheduling decisions. The memory behavior of a program is modeled based on Stack Distance Profiles (SDP) - the distribution of the stack distances for all the memory accesses of a task, where a stack distance is by definition the number of unique cache lines accessed between two consecutive accesses to a same line [18]. The difference between SDP and our model is that SDP is achieved by on-line monitored counters such as valgrind [19] while UCB and ECB are achieved by an off-line WCET analysis tools as below. At the moment, there is no archived comparison between the two. UCB analysis with scheduling simulation could be more pessimistic but safer because it takes into account the program point with largest number of UCB.

Several WCET analysis tools allow designers to perform cache analysis. SymTA/P[23], HEPTANE[9], Chronos[16] and aiT [2] are examples of them. UCB computation of a program is supported by aiT. The analysis of those tools are based on program path analysis of the control flow graph of the program. It is compliant with the requirement of our proposed tool.

In conclusion, WCET tools focus on the evaluation of program's control flow graph to compute the WCET and also a few tools can compute cache access profile. The analysis result could be used as an input for a scheduling analysis tool. In the domain of real-time scheduling analysis, the support for cache and CRPD when performing scheduling analysis is not very well specified. As far as we know, only SimSo clearly supports scheduling simulation with cache analysis based on SDP. Then, we proposed a tool available to the community which can either compute cache access profile of a task from its control flow graph or re-use information obtained from a WCET/cache analysis tool to perform scheduling analysis. Our model is compliant with existing work in [14], [5], [2] and [17]. In addition, because Cheddar provides a large set of scheduling analysis methods, we can fully investigate the dependency between CRPD and other scheduling parameters in order to either adjust or optimize a RTES design to meet its timing constraints.

7. CONCLUSIONS

In the article, we presented an approach to implement a cache-aware scheduling simulator. The work was proceeded in the context of the Cheddar real-time scheduling analyzer, which is open-source, freely available to researchers and practitioners that want to investigate scheduling analysis of RTES with cache memory. Our solution consists of three parts: modeling cache memory and cache access profile, implementing several cache analysis methods and per-

---

[1] RealTime-at-Work, http://www.realtimeatwork.com/

[2] AbsInt Inc., http://www.absint.com/

forming scheduling simulation. We extended Cheddar to be able to deal with cache memory and illustrated the dependency between cache and other scheduling parameters.

There are open problems we are aiming to address in the future. The first one concerns the refinement of other CRPD analysis methods. Several improvements have been proposed in [2] to reduce the upper-bound of the CRPD. In addition, we plan to compare our approach with the approach based on Stack Distance Profile in [8]. Second, we are going to study the problem of feasibility interval of RTES with cache memory.

8. REFERENCES

[1] S. Altmeyer, R. I. Davis, and C. Maiza. Improved cache related pre-emption delay aware response time analysis for fixed priority pre-emptive systems. *Real-Time Systems*, 48(5):499–526, 2012.

[2] S. Altmeyer and C. Maiza Burguière. Cache-related preemption delay via useful cache blocks: Survey and redefinition. *Journal of Systems Architecture*, 57(7):707–719, 2011.

[3] M. Bertogna, O. Xhani, M. Marinoni, F. Esposito, and G. Buttazzo. Optimal selection of preemption points to minimize preemption overhead. In *Real-Time Systems (ECRTS), 2011 23rd Euromicro Conference on*, pages 217–227. IEEE, 2011.

[4] E. Bini and G. C. Buttazzo. Measuring the performance of schedulability tests. *Real-Time Systems*, 30(1-2):129–154, 2005.

[5] J. V. Busquets-Mataix, J. J. Serrano, R. Ors, P. Gil, and A. Wellings. Adding instruction cache effect to schedulability analysis of preemptive real-time systems. In *Proceedings of the 2^{nd} IEEE Real-Time Technology and Applications Symposium (RTAS)*, pages 204–212, 1996.

[6] G. C. Buttazzo. Rate monotonic vs. edf: judgment day. *Real-Time Systems*, 29(1):5–26, 2005.

[7] Y. Chandarli, F. Fauberteau, D. Masson, S. Midonnet, M. Qamhieh, et al. Yartiss: A tool to visualize, test, compare and evaluate real-time scheduling algorithms. In *Proceedings of the 3rd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems*, pages 21–26, 2012.

[8] M. Chéramy, A.-M. Déplanche, P.-E. Hladik, et al. Simulation of real-time multiprocessor scheduling with overheads. In *International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH)*, 2013.

[9] A. Colin and I. Puaut. Worst-case timing analysis of the rtems real-time operating system. *Rapport Nâ̊Ů ̨e PI1277, IRISA, France*, 1999.

[10] L. Cucu and J. Goossens. Feasibility intervals for fixed-priority real-time scheduling on uniform multiprocessors. In *IEEE Conference on Emerging Technologies and Factory Automation (ETFA)*, pages 397–404, 2006.

[11] C. Fotsing, F. Singhoff, A. Plantec, V. Gaudel, S. Rubini, S. Li, H. N. Tran, L. Lemarchand, P. Dissaux, and J. Legrand. Cheddar architecture description language. 2014.

[12] M. González Harbour, J. Gutiérrez García, J. Palencia Gutiérrez, and J. Drake Moyano. Mast: Modeling and analysis suite for real time applications. In *Real-Time Systems, 13th Euromicro Conference on, 2001.*, pages 125–134. IEEE, 2001.

[13] R. Henia, A. Hamann, M. Jersak, R. Racu, K. Richter, and R. Ernst. System level performance analysis–the symta/s approach. *IEE Proceedings-Computers and Digital Techniques*, 152(2):148–166, 2005.

[14] C.-G. Lee, H. Hahn, Y.-M. Seo, S. L. Min, R. Ha, S. Hong, C. Y. Park, M. Lee, and C. S. Kim. Analysis of cache-related preemption delay in fixed-priority preemptive scheduling. *IEEE Transactions on Computers*, 47(6):700–713, 1998.

[15] C. Li, C. Ding, and K. Shen. Quantifying the cost of context switch. In *Proceedings of the 2007 workshop on Experimental computer science*. ACM, 2007.

[16] X. Li, Y. Liang, T. Mitra, and A. Roychoudhury. Chronos: A timing analyzer for embedded software. *Science of Computer Programming*, 69(1):56–67, 2007.

[17] W. Lunniss, S. Altmeyer, and R. I. Davis. Optimising task layout to increase schedulability via reduced cache related pre-emption delays. In *Proceedings of the 20th International Conference on Real-Time and Network Systems*, pages 161–170. ACM, 2012.

[18] R. L. Mattson, J. Gecsei, D. R. Slutz, and I. L. Traiger. Evaluation techniques for storage hierarchies. *IBM Systems journal*, 9(2):78–117, 1970.

[19] N. Nethercote and J. Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *ACM Sigplan notices*, volume 42, pages 89–100, 2007.

[20] A. Plantec and F. Singhoff. Refactoring of an ada 95 library with a meta case tool. In *ACM SIGAda Ada Letters*, volume 26, pages 61–70. ACM, 2006.

[21] F. Singhoff, J. Legrand, L. Nana, and L. Marcé. Cheddar: a flexible real time scheduling framework. *ACM SIGAda Ada Letters*, 24(4):1–8, 2004.

[22] J. Staschulat, S. Schliecker, and R. Ernst. Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In *Proceedings of the 17th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 41–48, 2005.

[23] J. Staschulat, S. Schliecker, and R. Ernst. Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In *Euromicro Conference on Real-Time Systems (ECRTS)*, Palma de Mallorca, Spain, 2005.

[24] H. N. Tran, F. Singhoff, S. Rubini, and J. Boukhobza. Instruction cache in hard real-time systems: modeling and integration in scheduling analysis tools with AADL. In *Proceedings of the 12^{th} IEEE International Conference on Embedded and Ubiquitous Computing*, Milan, Italy, 2014.

[25] H. N. Tran, F. Singhoff, S. Rubini, and J. Boukhobza. Addressing cache related preemption delay in fixed priority assignment. In *Proceedings of the 20th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, Luxembourg, 2015.

[26] R. Urunuela, A. Deplanche, and Y. Trinquet. Storm, a simulation tool for real-time multiprocessor scheduling evaluation. In *IEEE Conference on Emerging Technologies and Factory Automation (ETFA)*, pages 1–8, 2010.