

PAStime: Progress-aware Scheduling for Time-critical Computing

Soham Sinha
Department of Computer Science
Boston University
Boston, USA
soham1@bu.edu

Richard West
Department of Computer Science
Boston University
Boston, USA
richwest@bu.edu

Abstract—Over-estimation of worst-case execution times (WCETs) of real-time tasks leads to poor resource utilization. In a mixed-criticality system (MCS), the over-provisioning of CPU time to accommodate the WCETs of highly critical tasks can lead to degraded service for less critical tasks. In this paper, we present PAStime, a novel approach to monitor and adapt the runtime progress of highly time-critical applications, to allow for improved service to lower criticality tasks. In PAStime, CPU time is allocated to time-critical tasks according to the delays they experience as they progress through their control flow graphs. This ensures that as much time as possible is made available to improve the Quality-of-Service of less critical tasks, while high-criticality tasks are compensated after their delays.

In this paper, we integrate PAStime with Adaptive Mixed-criticality (AMC) scheduling. The LO-mode budget of a high-criticality task is adjusted according to the delay observed at execution checkpoints. Using LITMUS^{RT} to implement both AMC and AMC-PAStime, we observe that AMC-PAStime significantly improves the utilization of low-criticality tasks while guaranteeing service to high-criticality tasks.

I. INTRODUCTION

In real-time systems, computing resources are typically allocated according to each task’s worst-case execution time (WCET), to ensure timing constraints are met. However, worst-case conditions for an application are rather rare, resulting in poor resource utilization. Previous research work [49] shows that the worst-cases lie at the tiny tail-end of the probability distribution curve of the execution times for many programs. Instead, average-case execution times are significantly more likely, taking a fraction of the worst-case times.

Mixed-criticality systems (MCSs) provide a way to avoid over-estimation of resource needs, by considering the schedulability of tasks according to different estimates of their execution times at different criticality or assurance levels [48]. Higher criticality tasks are afforded more execution time at the cost of less time for lower criticality tasks, when it is impossible to meet all task timing constraints. There have been multiple proposals [9], [10], [14], [18] since Vestal’s work on MCSs [48]. Most prior work focuses on meeting task deadlines and ignores Quality-of-Service (QoS) metrics [47] such as frames-per-second or average utilization. Although timely completion is most important for high-criticality applications, QoS is a significant metric for lower criticality tasks [13], [37], [44]. This has motivated our work on PAStime (Progress-

Aware Scheduling for *time*-critical computing), to maximize the QoS for low-criticality tasks.

In PAStime, we first identify a checkpoint in a high-criticality application at an intermediate stage in its source code. The application is then profiled offline to measure the time to reach the marked checkpoint. Using this timing data, the application evaluates its progress at the checkpoint during runtime. Based on the delay at the checkpoint, PAStime predicts the expected execution time of a high-criticality application. We consequently adjust the runtime of the application, given that the change does not hamper schedulability of co-running tasks. If at runtime, a highly critical program is deemed to be making insufficient progress, it is given greater CPU time.

We combine PAStime with Adaptive Mixed-criticality (AMC) scheduling [10], to improve the performance of low-criticality tasks. In AMC, the system is started in *LO-mode*, where all tasks are scheduled with their LO-mode budgets. When a high-criticality task runs for more than its LO-mode budget, the system is switched to *HI-mode*. In HI-mode, all low-criticality tasks are stopped, and the high-criticality tasks are given their increased HI-mode budgets. However, switching to HI-mode should be avoided as much as possible [8], because it affects the performance of low-criticality tasks, which are not executed in HI-mode.

Several works extend the mixed-criticality task model to improve the performance of low-criticality tasks, such as providing an offline extra budget allowance to the high-criticality tasks [41], and estimating multiple budgets [36], [39] and periods [42], [44] for low-criticality tasks. However, these approaches do not utilize runtime information.

We extend AMC with PAStime to avoid mode switches, by dynamically adjusting the LO-mode budget for a high-criticality task, based on progress to execution checkpoints. Then, we predict the expected execution time of a high-criticality task based on the observed delay until a checkpoint. We carry out an efficient online schedulability test to determine whether we can increase the LO-mode budget of the delayed high-criticality task to our predicted execution time. In case the taskset is still schedulable with the increased budget, we extend the LO-mode budget of the high-criticality task to the predicted execution time. When a high-criticality task finishes

within its extended execution budget, we keep the system in LO-mode and avoid a mode switch that would otherwise happen in AMC scheduling. Thus, PASTime improves the QoS of low-criticality tasks by keeping the system in LO-mode for a longer time.

Factors such as I/O events and hardware micro-architectural delays lead to actual execution times exceeding those predicted by PASTime. Any high criticality task running at the end of its predicted execution time causes a timer interrupt to switch the system into HI-mode, as is the case with AMC scheduling. Thus, a high-criticality task never misses its deadline.

Contributions: The central idea of PASTime is to help the OS make scheduling decisions based on a program’s runtime progress. We implement both AMC and a PASTime extension to AMC scheduling in LITMUS^{RT} [12], [16]. We test our implementation with real-world applications: an object classification application from the Darknet neural network framework [40], and an MPEG video decoder [2].

We show that PASTime increases the average utilization of low-criticality tasks by 1.5 to 9 times for 2 to 20 tasks. We also demonstrate that our implementation of AMC-PASTime has minimal and bounded additional overhead in LITMUS^{RT}.

We provide a C library for PASTime to instrument checkpoints in high-criticality programs. In addition, we modify the LLVM compiler [30] to automatically identify potential locations of checkpoints during profiling for simple time-critical applications.

The next section describes our approach to PASTime. Section III details the theoretical background behind AMC and its extension with PASTime. Section IV describes the design and implementation of PASTime, which is then evaluated in Section V. Finally, we describe related work, followed by conclusions and future work in Sections VI and VII, respectively.

II. APPROACH

Compiler infrastructures such as LLVM are capable of producing a program’s control flow graph (CFG). A CFG represents the interconnection between multiple *Basic Blocks* (BBs), where a block is a sequence of straight-line code without any internal branches. However, CFGs are not typically utilized by an OS to manage time for different computing resources, in spite of being a rich source of analytical information about a program. Consequently, current OSs are oblivious to a program’s computing requirements (e.g., CPU utilization) at different points in its execution. A developer of an application can, instead, help the OS make decisions related to resource management, by providing runtime information about a program at certain points in its source code.

PASTime dynamically decides a program’s execution budget based on its runtime progress and theoretical analysis of the allowable delay at specific checkpoint locations. At runtime, PASTime measures the time to reach a checkpoint from the start of a task, and then compares that time to a pre-profiled reference value. The execution budget for the task is then adjusted according to actual progress.

Fig. 1: CFG and average estimates of a program

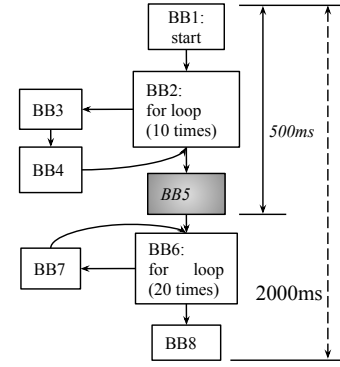


Figure 1 shows the CFG for a program with two loops, starting at BB2 and BB6. In this example, PASTime inserts a checkpoint between the two loops at the end of BB5. BB5 is a potentially good location for a checkpoint because there is one loop before and after this BB, providing an opportunity to increase the execution budget to compensate for the delay until BB5.

Suppose that we derive the LO-mode budget of the whole program to be 2000 ms by profiling, and the time to reach the checkpoint at BB5 is 500 ms. The program is then executed in the presence of other tasks. The execution budget at the checkpoint (in BB5) is adjusted, to account for the program’s actual runtime progress. For example, at the checkpoint in BB5, suppose the program experiences a delay of 100 ms. This means that the checkpoint in BB5 is reached in 600 ms, instead of the expected 500 ms. Therefore, the program is delayed by $(\frac{100}{500} \times 100\%) = 20\%$. Then, we predict the total execution time of 2000 ms to be $(2000 + 20\% \times 2000) = 2400$ ms. This is based on the runtime information until the checkpoint in BB5 which has shown a delay of 100 ms. Hence, we need to increase the program’s execution budget by 400 ms. PASTime uses this available information until a checkpoint to extend the LO-mode budget of a high-criticality application.

A. Benefits in Adaptive Mixed-criticality Scheduling

We see progress-aware scheduling as being beneficial in mixed-criticality systems. Adaptive Mixed-criticality scheduling [10] is the state-of-the-art fixed-priority scheduling policy for Mixed-criticality tasksets. In AMC, a system is first initialized to run in *LO-mode*. In LO-mode, all the tasks are executed with their LO-mode execution budgets. Whenever a high-criticality task overruns its LO-mode budget, the system is switched to *HI-mode*. In HI-mode, all low-criticality tasks are discarded (or given reduced execution time [8], [36]), and the high-criticality tasks are given increased HI-mode budgets. The system’s switch to HI-mode therefore impacts the QoS for low-criticality tasks.

By combing PASTime with AMC (to yield AMC-PASTime), we extend the LO-mode budget of a high-criticality task to its predicted execution time at a checkpoint. Going back to our previous example in Figure 1, we try to extend the LO-mode budget of the task by 400 ms. We carry out an online

schedulability test to determine if we can extend the task's LO-mode budget by 400 ms. If the whole taskset is still schedulable after an extension of the LO-mode budget of the delayed high-criticality task, the increment in the task's LO-mode budget is approved. We let the task run until the extended time and keep the system in LO-mode. In case the high-criticality task finishes within its extended time, we avoid an unnecessary switch to HI-mode. Thus, the low-criticality tasks run for an extended period of time and do not suffer degraded CPU utilization and QoS, as occurs with AMC.

In case the task does not finish within the predicted time, then the system is changed to HI-mode, and low-criticality tasks are finally dropped. This behavior is identical to AMC, and every high-criticality task still finishes within its own deadline. Therefore, we improve the QoS of the low-criticality tasks when the high-criticality tasks finish within their extended LO-mode budgets, and otherwise, we fall back to AMC. When there is no delay at a checkpoint, we do not change a task's LO-mode budget.

III. THEORETICAL BACKGROUND

In this section, we provide a response time analysis for AMC-PAStime by extending the analysis for AMC scheduling. We also describe details about the online schedulability test based on the response time values.

A. Schedulability Analysis for AMC and PAStime

1) *Task Model*: We use the same AMC task model as described by Baruah et al [10]. Without loss of generality, we restrict ourselves to two criticality levels - LO and HI. Each task, τ_i , has five parameters: $C_i(LO)$ - LO-mode runtime budget, $C_i(HI)$ - HI-mode runtime budget, D_i - deadline, T_i - period, and L_i - criticality level of a task, which is either high (*HC*) or low (*LC*). We assume each task's deadline, D_i , is equal to its period, T_i . A HC task has two budgets: $C_i(LO)$ for LO-mode assurance and $C_i(HI)$ ($> C_i(LO)$) for HI-mode assurance. A LC task has only one budget of $C_i(LO)$ for LO-mode assurance.

2) *Scheduling Policy*: Both AMC and AMC-PAStime use the same task priority ordering, based on Audsley's priority assignment algorithm [5]. If a task's response time for a given priority order is less than its period, then the task is deemed schedulable. The details of the priority assignment strategy are discussed in previous research work [6], [10]. We do not change the priority ordering of the tasks at runtime.

AMC-PAStime initializes a system in LO-mode with all tasks assigned their LO-mode budgets. We extend a high-criticality task's LO-mode budget at a checkpoint if the task is lagging behind its expected progress, as long as the extension does not hamper the schedulability of the delayed task and all the lower or equal priority tasks. An increase to the LO-mode budget of a task that violates its own or other task schedulability is not allowed. If a high-criticality task has not finished its execution even after its extended LO-mode budget, then the system is switched to HI-mode.

3) *Response Time Analysis*: The AMC response time recurrence equations for (1) all tasks in LO-mode, (2) HC tasks in HI-mode, and (3) HC tasks during mode-switches are shown in Equation 1, 2 and 3, respectively. $hp(i)$ is the set of tasks with priorities higher than or equal to that of τ_i . Likewise, $hpHC(i)$ and $hpLC(i)$ are the set of high- and low-criticality tasks, respectively, with priorities higher than or equal to the priority of τ_i .

AMC provides two analyses for mode switches: AMC-response-time-bound (AMC-rtb) and AMC-maximum. We use AMC-rtb for our analysis, as AMC-maximum is computationally more expensive. However, AMC-rtb does not allow a taskset which is not schedulable by AMC-maximum. Therefore, AMC-rtb is *sufficient* for schedulability.

$$R_i^{LO} = C_i(LO) + \sum_{\tau_j \in hp(i)} \left\lceil \frac{R_j^{LO}}{T_j} \right\rceil \times C_j(LO) \quad (1)$$

$$R_i^{HI} = C_i(HI) + \sum_{\tau_j \in hpHC(i)} \left\lceil \frac{R_j^{HI}}{T_j} \right\rceil \times C_j(HI) \quad (2)$$

$$R_i^* = C_i(HI) + \sum_{\tau_j \in hpHC(i)} \left\lceil \frac{R_j^*}{T_j} \right\rceil \times C_j(HI) + \sum_{\tau_j \in hpLC(i)} \left\lceil \frac{R_j^{LO}}{T_j} \right\rceil \times C_j(LO) \quad (3)$$

With AMC-PAStime, if a high-criticality task, τ_i , is delayed by $X\%$ at a checkpoint, compared to its pre-profiled progress, then τ_i 's budget is tentatively increased from $C_i(LO)$ to $C'_i(LO)$. $C'_i(LO) = f(C_i(LO), X)$. Here, $f(C_i(LO), X)$ is a function to predict the delayed total execution time, given that the observed delay until the checkpoint is X . For example, by a linear extrapolation of the observed delay of X until a checkpoint up to the full $C_i(LO)$, $f(C_i(LO), X) = (C_i(LO) + \frac{C_i(LO) \times X}{100})$. f could also depend on the other hardware microarchitectural factors like caches. We show multiple execution time prediction techniques in Section V.

An online schedulability test then calculates the the extended LO-mode response time, R_i^{LO-ext} , for τ_i , using Equation 4. Similarly, R_i^{*-ext} is calculated using Equation 5. Equations 4 and 5 are extensions of Equations 1 and 3. AMC-PAStime checks at runtime whether both R_i^{LO-ext} and R_i^{*-ext} are less than or equal to τ_i 's period to determine its schedulability.

The new response times are then calculated for all tasks in LO-mode with priorities less than or equal to τ_i , using Equation 4. Similarly, new response times are calculated for all HC tasks with lower or equal priority to τ_i during a mode switch, using Equation 5. If all newly calculated response times are less than or equal to the respective task periods, the system is schedulable. In this case, AMC-PAStime approves the LO-mode budget extension to τ_i . If the system is not schedulable, then τ_i 's budget remains $C_i(LO)$.

$$R_i^{LO-ext} = C'_i(LO) + \sum_{\tau_j \in hp(i)} \left\lceil \frac{R_j^{LO-ext}}{T_j} \right\rceil \times C'_j(LO) \quad (4)$$

$$R_i^{*-ext} = C_i(HI) + \sum_{\tau_j \in hpHC(i)} \left\lceil \frac{R_j^{*-ext}}{T_j} \right\rceil \times C_j(HI) + \sum_{\tau_j \in hpLC(i)} \left\lceil \frac{R_j^{LO-ext}}{T_j} \right\rceil \times C_j(LO) \quad (5)$$

AMC-PAStime only extends the LO-mode budget of a delayed HC task for its current job. When a new job for the same HC task is dispatched, it starts with its original LO-mode budget. If another request for an extension in LO-mode for the same task appears, AMC-PAStime tests the schedulability with the maximum among the requested extended budgets. The system keeps track of the maximum extended budget for a task and uses that value for online schedulability testing. We explain the AMC-PAStime scheduling scheme with an example taskset in Table I.

TABLE I: A mixed-criticality taskset example

Task	Type	C(LO)	C(HI)	T	Pr	R^{LO}	R^*
τ_1	HC	3	6	10	1	3	6
τ_2	LC	2	-	9	2	5	-
τ_3	HC	5	10	50	3	15	38

Suppose, task τ_1 is delayed by 66% at a checkpoint in the task's source code. PAStime will then try to extend the budget by $(3 \times \frac{66}{100}) \approx 2$ time units. Therefore, the potential extended budget $C_1'(LO)$ for τ_1 would be $(3 + 2) = 5$ time units. PAStime will calculate the response times, R_i^{LO-ext} and R_i^{*-ext} , for τ_1 and the lower priority tasks τ_2 and τ_3 . R_1^{LO-ext} would just be 5, and R_1^{*-ext} would remain the same as $R_1^* = 6$, as τ_1 is the highest priority task. The new R_2^{LO-ext} would be 7 (by Equation 4) which is smaller than its period of 9. Therefore, τ_2 would still be schedulable if we extend τ_1 's LO-mode budget from 3 to 5.

For τ_3 , the new R_3^{LO-ext} and R_3^{*-ext} would be, respectively, 26 (by Equation 4) and 40 (by Equation 5) which are also smaller than τ_3 's period of 50. Therefore, the extended budget of 5 for τ_1 would be approved by AMC-PAStime. In conventional AMC scheduling, the system would be switched to HI-mode if τ_1 did not finish within 3 time units. However, AMC-PAStime will extend τ_1 's LO-mode budget to 5 because of the observed delay at its checkpoint, so the system is kept in LO-mode. Consequently, LC task τ_2 is allowed to run by AMC-PAStime, if τ_1 finishes before 5 time units. If τ_1 does not finish even after 5 time units, the system would be switched to HI-mode.

In this example, 5 jobs of τ_1 are dispatched for every single job of τ_3 , as τ_3 's period of 50 is 5 times the period of τ_1 . Suppose τ_1 asks for the 66% increment in its LO-mode budget for the first job, as we have explained above. In its second job, τ_1 asks for an increment of 33% (1 time unit) in its LO-mode budget. In this case, we again need to calculate the response times for all tasks. If we calculate the online response times for τ_2 and τ_3 assuming $(3 + 1) = 4$ time units for $C_1'(LO)$, then we would not account for the first job of τ_1 , which potentially executes for 5 time units. We would need to keep track of the extended LO-mode budgets for all jobs of τ_1 , to accurately calculate online response times for τ_2 and τ_3 .

To avoid the cost of recording all extended LO-mode budgets for a task, AMC-PAStime simply stores the maximum extended budget for a task. When calculating online response times to check whether to approve an extension to the LO-mode budget, the system uses the maximum extended budget of every high-criticality task. This value is stored in the `max_extended_budget` variable for each HC task. Therefore, when τ_1 asks for 4 time units as its extended LO-mode budget in its second job, the system calculates R_1^{LO-ext} , R_1^{*-ext} , R_2^{LO-ext} , R_3^{LO-ext} , R_3^{*-ext} with $C_1'(LO) = 5$.

AMC-PAStime uses maximum extended budgets to calculate safe upper bounds for online response times. The `max_extended_budget` is reset when a task has not requested a LO-mode budget extension for any of its jobs dispatched within the maximum period of all tasks.

B. Online Schedulability Test

AMC-PAStime performs an online schedulability test whenever a high-criticality task asks for an extension to its LO-mode budget. The test calculates the response times (R_i^{LO-ext} and R_i^{*-ext}) of the delayed high-criticality task and all lower priority tasks. Then, it checks whether the response times are less than or equal to the task periods. If any task's response time is greater than its period, the online schedulability test returns false, and the extension in LO-mode for the high-criticality task is denied. If the schedulability test is successful the high-criticality task is permitted to run for its extended budget in LO-mode.

1) *Initial Value for Online Response Times Calculation:* Online response time calculations may take significant time, depending on the number of iterations of Equations 4 and 5. However, the response time values from Equations 1, 2, 3 are already calculated offline to determine the schedulability of a taskset using Audsley's priority assignment algorithm [5] with AMC scheduling. Since AMC-PAStime uses the same priority ordering as AMC scheduling, the offline response times for schedulability remain same.

AMC-PAStime initializes the online R_i^{LO-ext} in Equation 4 with $R_i^{LO} + e$, where R_i^{LO} is calculated offline by Equation 1 and $e (> 0)$ is the extra budget of the delayed task. e is set to $\frac{C_i(LO) \times X}{100}$ if τ_i is delayed by $X\%$ at its checkpoint.

Since the budget of a delayed task is extended by e , R_i^{LO-ext} must be greater than or equal to $R_i^{LO} + e$. Hence, this is a good initial value to start calculating R_i^{LO-ext} online. In Section V, we establish an upper bound on the total number of iterations needed to check the schedulability of random tasksets, if the highest priority task's budget is extended by different amounts.

2) *Details of the Online Test:* Algorithm 1 is the pseudocode for the online schedulability test. The offline response time values (R_i^{LO} and R_i^*) are stored in the timing properties for each task τ_i . When a high-criticality task τ_k is delayed, it asks for an extension of its LO-mode budget by e time units.

Line 6 in Algorithm 1 determines whether the newly requested extension, e , is more than a previously saved maximum extended budget for τ_k . In lines 10–13, the $C_i'(LO)$ is set to the maximum extended budget for all the lower priority

tasks and τ_k . As we have explained above, it is practically infeasible to store the execution times of every job of all the tasks to calculate online response times. Therefore, we store the maximum extended budget of every task and use this to calculate response times online. For the currently delayed task, τ_k , $C'_k(LO)$ is set to the maximum value of a previously saved `max_extended_budget` and the currently requested $C_k(LO) + e$. Considering the example in Table I, line 12 would translate to $C'_i(LO) = \max(5, 4)$ for the second LO-mode extension request.

Then, Equation 4 is solved in line 15 by initializing R_i^{LO-ext} to the R_i^{LO} plus extra budget e' from line 6. If a lower priority task is a high-criticality task, then R_i^{*-ext} is calculated in line 19, with the newly derived value of R_i^{LO-ext} .

The initial value in the recurrence relations is set *offline* to significantly decrease the number of *online* iterations, and hence overhead, needed to determine taskset schedulability. In addition, we solve both recurrence relations with an incremental response time algorithm [17] to further minimize the overhead.

Algorithm 1 Online Schedulability Test

```

1: Input: tasks - set of all tasks in priority order
2:    $\tau_k$  - delayed task
3:    $e$  - extra budget for  $\tau_k$ 
4: Output: true or false
5: function ISBUDGETCHANGEAPPROVED(tasks,  $\tau_k$ ,  $e$ )
6:    $e' = \max(\tau_k.\text{max\_extended\_budget}, C_k(LO) + e) -$ 
7:      $C_k(LO)$ 
8:   for each task  $\tau_i$  in  $\{\tau_k \cup \text{lower priority tasks than } \tau_k$ 
9:     in } \tau_k\} do
10:     $C'_i(LO) = \tau_i.\text{max\_extended\_budget}$ 
11:    if  $\tau_i$  is  $\tau_k$  then
12:       $C'_i(LO) = \max(C'_i(LO), C_i(LO) + e)$ 
13:    end if
14:    Initialize  $R_i^{LO-ext}$  for Equation 4 with  $R_i^{LO} + e'$ 
15:    Solve Equation 4
16:    if  $R_i^{LO-ext} \leq T_i$  then
17:      if  $\tau_i$  is high-criticality then
18:        Initialize  $R_i^{*-ext}$  for Equation 4 with  $R_i^{*-ext}$ 
19:        Solve Equation 5 with the new  $R_i^{LO-ext}$ 
20:        if  $R_i^{*-ext} > T_i$  then Return false
21:      end if
22:    end if
23:    else Return false
24:  end if
25: end for
26:  $\tau_k.\text{max\_extended\_budget} =$ 
27:    $\max(C_k(LO) + e, \tau_k.\text{max\_extended\_budget})$ 
28: Return true
29: end function

```

IV. DESIGN AND IMPLEMENTATION OF PASTIME

In this section, we first describe the overall design of AMC-PAStime in LITMUS^{RT} [12], [16]. This is followed

by a description of how checkpoints are instrumented in an application's source code. We then show the algorithm to determine and insert checkpoints, which is integrated into the LLVM compiler infrastructure. A high-criticality task requires profiling to determine the placement of checkpoints, before it is ready for execution with other tasks. We describe how the profiling and execution of a high-criticality task is performed, along with the scheduling mechanism in PASTime. The source code for PASTime in LITMUS^{RT} will be made publicly available.

A. Design Overview

AMC-PAStime has two phases: a Profiling phase for high-criticality tasks, and an Execution phase. In the Profiling phase, one or more checkpoints are placed at key stages in a program's source code. The *average* time to reach each checkpoint is then measured. After profiling all high-criticality tasks, the system switches into the Execution phase. The time taken to reach each checkpoint in every high-criticality task is observed by the system at runtime. Each observed time is compared against the profiled time to reach the same checkpoint. Any high-criticality task lagging behind its profiled time to a checkpoint is tentatively given increased LO-mode budget, according to the approach described in Section III.

Fig. 2: Implementation of AMC-PAStime in LITMUS^{RT}

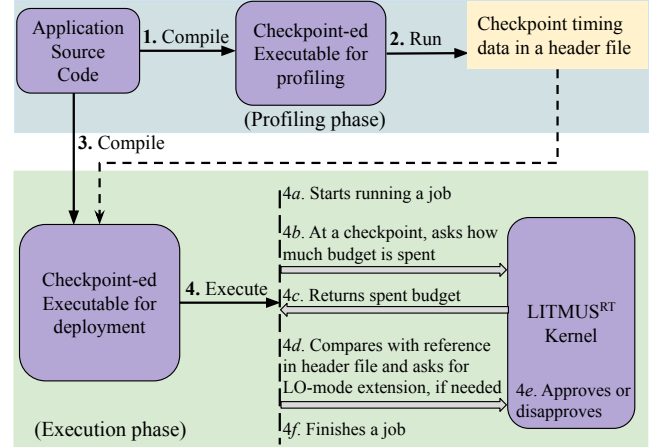


Figure 2 shows an overview of the design of AMC-PAStime in LITMUS^{RT}. Step 1 is the compilation of a high-criticality application's source code in the Profiling phase, which uses our compilation procedure [38]. The compiled executable has checkpoints embedded into its code for profiling. Step 2 executes the program to generate timing metadata for each checkpoint in a *timeinfo.h* file. Step 2 is performed multiple times with different program inputs to generate an average time to reach each checkpoint.

Step 3 compiles the source code along with the checkpoint timing metadata header file (*timeinfo.h*) for the Execution phase, producing a binary image that is used for deployment under working conditions. Finally, Step 4 runs the code in the Execution phase along with all other tasks. At some point

after the system is started, Step 4a starts running a job for a high-criticality task. When a checkpoint is reached, Step 4b asks the LITMUS^{RT} kernel how much budget it has consumed. Step 4c returns the spent budget from the LITMUS^{RT} kernel to the application.

After receiving the spent budget, t_{spent} , the application compares it to the reference timing, t_{ref} , for the checkpoint from the *timeinfo.h* header file. It calculates the extra budget that it needs in LO-mode, using Equation 6. If $e > 0$ in the equation, Step 4d asks LITMUS^{RT} for extra budget. The AMC-PAS_{time} scheduling policy in the LITMUS^{RT} kernel runs an online schedulability test using Algorithm 1. If Algorithm 1 returns true, the LO-mode budget of the current task is extended by e . If the algorithm returns false, the task budget is not altered. Finally, Step 4f finishes the current job of the running task.

$$e = \frac{C_i(LO) \times (t_{\text{spent}} - t_{\text{ref}})}{t_{\text{ref}}} \quad (6)$$

B. Checkpoint Instrumentation and Detection

A checkpoint is a key stage in an application’s code, used to evaluate the progress of a currently running task. Well placed checkpoints balance the number of instructions that are executed prior to the checkpoint, with those that remain to the next checkpoint or the end of the program. Ideally, there should be a meaningful number of instructions leading up to a checkpoint to determine progress. Likewise, there should be sufficient instructions after a checkpoint to increase the likelihood that a task is adequately compensated for execution delays using an extended budget.

A developer of a high-criticality application finds a potential checkpoint location in the program’s source code for its Execution phase, after trying out multiple locations in the Profiling phase. PAS_{time} includes a development library to instrument checkpoints for the two different phases. Additional modifications to the LLVM compiler [30] are used to detect and instrument checkpoints in the Profiling phase.

1) *Checkpoint Library*: We have developed a C library to instrument checkpoints in the two PAS_{time} phases. The main purpose of the library is to generate the necessary checkpoint timing information during the Profiling phase and then request a task’s extended LO-mode budget from the LITMUS^{RT} kernel during the Execution phase.

In the Profiling phase, a `noteTime` function call from our library is inserted into the application code at a desired checkpoint. `noteTime` takes a unique ID for each checkpoint. The function logs the time to reach that checkpoint in the source code since the start of a job. The average of multiple such timing entries is saved in *timeinfo.h* after the Profiling phase.

During the Execution phase, a preprocessor macro called `EXTEND_BUDGET` is inserted at a checkpoint. This macro obtains the spent budget from the LITMUS^{RT} kernel via a `get_current_budget` system call. Then, it compares the spent budget with the reference budget from the *timeinfo.h* header file, and calculates the extra budget using

Equation 6. If extra budget is needed, the macro makes a `set_rt_task_param` system call.

2) *Manually Inserted Checkpoint*: A developer may attempt various strategies to identify a key stage [15], [23], [46] of an application to instrument a checkpoint. The developer uses either the `noteTime` function for the Profiling phase, or the `EXTEND_BUDGET` macro for the Execution phase to instrument a checkpoint. Checkpoints should generally be avoided inside tight loops. Visiting a checkpoint every loop iteration incurs a small overhead that is accumulated across multiple iterations.

3) *Automatic Instrumentation of a Checkpoint*: We have written a compiler pass in LLVM to automatically instrument checkpoints for the Profiling phase of PAS_{time}. The instrumented code is run in the Profiling phase, and multiple checkpoint timing information is generated. Finally, the developer chooses one such checkpoint for the Execution phase.

The compiler pass automatically inserts checkpoints in the basic block preceding each loop in a function, except the first loop. The first loop is excluded so that enough instructions are executed before a checkpoint to determine meaningful delays.

For nested loops, we consider only the outer loop. Automatic instrumentation works with only simple program structures and ignores intersecting loops. LLVM’s `LoopInfo` analysis identifies only natural loops [3]. We utilize the `LoopInfo` class in our checkpoint instrumentation implementation.

Algorithm 2 Determine and Insert Checkpoints

```

1: isLoopBefore: identifies if a loop is in the paths
2:                   from the starting BB to another BB
3: visited: set of already visited BBs in DFS
4: function INSERTCHECKPOINT(function)
5:   startingBB = function.getEntryBlock()
6:   DODFS(startingBB)
7: end function
8: function DODFS(currentBB)
9:   if currentBB is in visited then return
10:  end if
11:  LoopID = getLoopFor(currentBB)
12:  if LoopID != null then
13:    if isLoopBefore[currentBB]  $\wedge$ 
14:      isLoopHeader(currentBB) then
15:      insert checkpoint before currentBB
16:    end if
17:  end if
18:  insert currentBB in visited
19:  for each s in successors of currentBB do
20:    DODFS(s)
21:  end for
22: end function

```

a) *Checkpoint Location Algorithm*: Our algorithm to identify checkpoint locations for the Profiling phase is given in Algorithm 2. We start a Depth First Search (DFS) from the starting Basic Block (BB) of a function, by calling `DODFS`

in line 6. We check whether a BB is part of a loop using `LoopInfo`'s `getLoopFor` member function. This function returns a unique `LoopID` for every new loop. An inner loop within a nested loop has the same ID as its outer loop. If a BB is not part of a loop, then it returns `null`.

If a BB is part of a loop, we first check in line 12 whether there is any loop before the current loop using a dictionary `isLoopBefore`. `isLoopBefore` is pre-populated for every BB in the CFG to indicate whether there is at least one loop seen in the paths from the starting BB to the current BB. To pre-populate `isLoopBefore`, we just check whether there is a path to a BB from the loop BBs.

Then, in Algorithm 2, we verify that the current BB is a header of a loop using `LoopInfo`'s `isLoopHeader` member function. A header is the entry-point of a natural loop. We insert a checkpoint in the predecessor BB of a header.

Finally, if there is at least one loop before the current header BB, we add a checkpoint before the current BB in line number 13. As natural loops have only one header, we avoid inserting a checkpoint for any inner loop in nested loops. We continue the DFS by marking the current BB as visited.

b) Compiler Pass: We have implemented the above algorithm in a compiler pass within LLVM [30], to automatically detect and instrument appropriate function calls as checkpoints in a C language program. This uses our previously described Checkpoint library for the Profiling phase. A developer uses our modified LLVM compiler with their high-criticality application written in C. Our compiler pass operates at the LLVM Intermediate Representation (IR) level. It takes a piece of IR logic as input, figures out the points of interest according to the above algorithm for checkpoints in a particular function, and generates the instrumented IR. These IRs are compiled into executable machine code. As the compiler pass operates at the IR level, it is easily extensible to other high-level languages and back-ends supported by LLVM.

C. Profiling and Execution Phases

The Profiling phase of PASTime determines viable checkpoints for use in the Execution phase, and also the LO- and HI-mode budgets for a high-criticality application. A checkpointed program is run multiple times in the Profiling phase using a set of test cases. The program is allowed to have multiple test checkpoints, which are either generated automatically using our modified LLVM compiler, or manually by a developer. Each checkpoint has a unique ID given to the checkpoint function call `noteTime`. We have developed a Python package to help run the Profiling phase and collect the average times to reach different checkpoints.

Step 4 in Figure 2 is the start of the Execution phase. One checkpoint from those generated in the Profiling phase for a high-criticality application is instrumented with an `EXTEND_BUDGET` macro. Although in general it is possible to use multiple checkpoints within the same application in the Execution phase, our experience shows that one is sufficient to improve LO-mode service. Multiple checkpoints add overhead to the task execution. Moreover, later checkpoints account for

execution delays that make earlier checkpoints redundant, as long as they are reached before the LO-mode budget expires.

The key issue in deciding on a single checkpoint for the Execution phase is to ensure it is not placed too late in the instruction stream. If it is placed too far into the program code a mode switch may occur before the task's LO-mode budget is extended due to delays. We show in Section V-D the effects of using checkpoints at different locations in a program's code.

D. Implementation Details in LITMUS^{RT}

As a first step to applying PASTime for use in adaptive mixed-criticality scheduling, we extended LITMUS^{RT} with AMC support. This required modifications to the existing partitioned fixed-priority scheduling policy, to include the following new variables in the task properties: `c_lo`, `c_hi`, `r_lo`, `r_star`, `c_extended`, `max_extended_budget`.

Tasks are divided into low- and high-criticality classes. By default, the HI-mode budget for low-criticality tasks, `c_hi`, is set to zero. If desired, we also support a reduced, non-zero HI-mode execution budget for low-criticality tasks, as discussed in the work on Imprecise Mixed-criticality scheduling [36]. Task priorities are determined offline, along with all response times for a system operating in LO-mode (`r_lo`) and during a mode switch (`r_star`). These parameters are then initialized in the kernel when the system starts executing a set of tasks.

a) Policy for mode switches: A system-wide mode is initialized in the LITMUS^{RT} kernel. The system is started in LO-mode, with all tasks assigned their `c_lo` budgets. Whenever a high-criticality task is out of its LO-mode budget, an enforcement timer handler (based on Linux's high-resolution timer [34]) is fired, and the system is switched to HI-mode.

For AMC-PASTime, `c_extended` is the extended LO-mode budget for a delayed job of a high-criticality task. An enforcement timer handler is therefore triggered only when a high-criticality task is still unfinished after the depletion of its `c_extended` time.

As Baruah et al. suggest [10], an AMC system switches back to LO-mode when none of the high-criticality tasks have been running for more than their LO-mode budgets. In the case of AMC-PASTime, the system will switch to a lower criticality level when none of the high-criticality tasks have been running for more than their extended LO-mode budgets. A list is used to keep track of which high-criticality tasks complete within their (extended) LO-mode budgets, to determine when to revert to a lower system criticality level.

b) LO-mode Budget Extension: A task's LO-mode budget is extended by AMC-PASTime by making a `set_rt_task_param` system call inside the `EXTEND_BUDGET` macro. We have implemented the `task_change_params` callback of a LITMUS^{RT} `sched_plugin` interface to support runtime adjustment of task parameters. In the `task_change_params` callback, the system runs an online schedulability test according to Algorithm 1. If the test returns `true`, the budget extension is approved, and the enforcement timer for the current task is adjusted accordingly. The task's `c_extended` variable

is updated, along with the maximum value of its extended budget in `max_extended_budget`.

V. EVALUATION

We test our implementation of AMC and AMC-PAStime in LITMUS^{RT} with real-world applications on an Intel NUC Kit [25]. The machine has an Intel Core i7-5557U 3.1GHz processor with 8GB RAM, running Linux kernel 4.9. We use two real-world applications in our evaluations: a high-criticality object classification application from the Darknet neural network framework [40], and a low-criticality MPEG video decoder [2]. These applications are chosen for their relevance to the sorts of applications that might be used in infotainment and autonomous driving systems.

For object classification, we use the COCO dataset [33] images for both profiling and execution. For the video decoder application, we use the *Big Buck Bunny* video [1] as the input. We have turned off memory locking with `mlock` by the user-space `liblitmus` library, as multiple object classification tasks collectively require more RAM than the physical 8GB machine limit.

A. Task Parameters

Table II shows the LO- and HI-mode budgets for the two applications. Each object classification task consists of a series of jobs that classify objects in a single image. Each video decoder task decodes 30 frames in a single job.

TABLE II: Applications and their budgets

Application	C(LO)	C(HI)
object classification	345 ms	627 ms
video decoder	250 ms	-

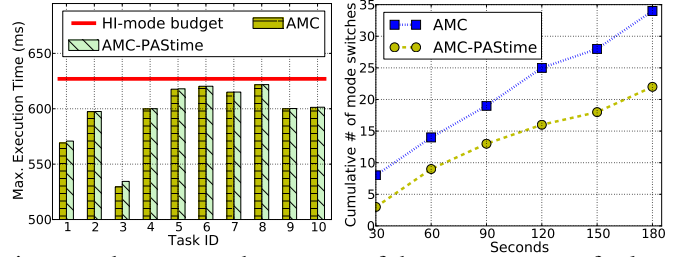
The LO-mode budget, $C(LO)$, is the average time that a task takes to complete its job. In Section V-C, we also present experiments where we increase our LO-mode budget estimate. The HI-mode budget, $C(HI)$, accounts for the worst-case running time of the high-criticality task for any of its jobs. The LO-mode utilization of each individual task is generated by the UUnifast algorithm [11]. We then calculate a task's period by dividing its LO-mode budget by its utilization.

Figure 3 shows the maximum execution times of 10 high-criticality object classification tasks in the presence of 10 other low-criticality video decoder tasks. We see that none of the tasks exceed the HI-mode budget of 627 ms, which is also the case in all subsequent experiments. Consequently, none of the high-criticality tasks miss their deadlines in any of our tests.

B. QoS Improvements for Low-criticality Tasks

We compare the QoS for low-criticality tasks using AMC and AMC-PAStime in different cases. In every case, each taskset has an equal number of high-criticality object classification tasks and low-criticality video decoder tasks. We experiment with ten schedulable tasksets in all cases except the base case described below. We run each of the tests ten

Fig. 3: Object classification Fig. 4: Switches to HI-mode




times, and we report the average of the measurements for low-criticality tasks. As stated earlier, all high-criticality tasks meet their timing requirements in each case.

1) *Base Case - Two Tasks*: Our base case is to run one high-criticality object classification task and one low-criticality video decoder task for 180 seconds. Here, we set the periods of both tasks to 1000 ms rather than using the UUnifast algorithm, yielding a total LO-mode utilization of ~60%.

Figure 5a shows the cumulative number of frames decoded by the video decoder task. The LO-mode Upper Bound (UB) line shows the cumulative number of decoded frames if the system is kept in LO-mode for the full 180 seconds. This line represents a theoretical upper bound for a decoded frame playback rate of 30 frames per second over the entire experimental run.

We see in Figure 5a that AMC-PAStime has a 9–21% increase in the cumulative number of decoded frames compared to AMC scheduling. The performance of the low-criticality task is related to the number of HI-mode switches in the two scheduling policies. Figure 4 shows that AMC-PAStime decreases the number of HI-mode switches by 35% over the entire execution run, compared to AMC scheduling.

Fig. 5: Two tasks - Video Decoder performance



(a) Cumulative number of frames

(b) Average utilization

Although the number of decoded frames is an illustrative metric for a video decoder's QoS, the average utilization of an application is a more generic metric. Average utilization represents the CPU share a task receives over a period of time. Figure 5b shows that AMC-PAStime achieves 10% more utilization on average for the video decoder task, compared to AMC scheduling. The LO-mode UB line is the maximum utilization of the video decoder task, which is 25% (i.e., $C(LO)/Period = 250 \text{ ms}/1000 \text{ ms}$).

2) *Scalability*: To test system scalability, we increase the number of tasks in a taskset up to 20 tasks. As explained above, we generate the periods of the tasks by distributing the

total LO-mode utilization of $\sim 60\%$ to all the tasks using the UUnifast algorithm. This setup is inspired by the theoretical parameters in previous mixed-criticality research work [10]. The LO-mode utilization bound for low-criticality tasks remains between 25–35%.

Figure 6 shows the average utilization of the low-criticality tasks, when the total tasks vary from 2 to 20. Each task in this case consists of 20 jobs. We see that the average utilization drops for AMC scheduling as the number of tasks increases.

AMC-PAStime achieves significantly greater average utilization for the low-criticality tasks, by deferring system switches to HI-mode until much later than with AMC scheduling. This is because the LO-mode budgets for the high-criticality tasks are extended due to runtime delays.

Table IIIa shows that AMC-PAStime decreases the number of mode switches by 28–55%. AMC-PAStime’s resistance to switching into HI-mode allows low-criticality tasks to make progress. This in turn improves their QoS. In these experiments, AMC-PAStime improves the utilization of the low-criticality tasks by a factor of 3, 5 and 9, respectively, for 8, 14 and 20 tasks.

TABLE III: Number of mode switches

# tasks	AMC	AMC-PAStime	Utilization (%)	AMC	AMC-PAStime
2	4	2	40	11	4
8	9	4	50	11	4
14	7	5	60	9	4
20	5	3	70	10	5
			80	11	9

(a) Varying number of tasks

(b) Varying LO-mode utilization

Fig. 6: Varying # of tasks





Fig. 7: Varying LO-utilization



3) *Varying the Initial Total LO-mode Utilization:* In this test, we vary the initial total LO-mode utilization for 8 tasks from 40% to 80% by adjusting the periods of all tasks. The initial utilization does not account for increases caused by LO-mode budget extensions to high-criticality tasks.

Figure 7 demonstrates that AMC-PAStime improves average utilization of the low-criticality tasks by more than 3 times, up to 70% total LO-mode utilization. After that, AMC-PAStime and AMC scheduling converge to the same average utilization for low-criticality tasks. This is because there is insufficient surplus CPU time in LO-mode for AMC-PAStime to accommodate the extended budget of a high-criticality task.

Therefore, the LO-mode extension requests are disapproved by AMC-PAStime.

Table IIIb demonstrates that AMC-PAStime improves the low-criticality tasks’ QoS by reducing the number of mode switches by $\sim 50\%$ until 70% LO-mode utilization.

We note that the schedulability of random tasksets decrease with higher LO-mode utilization in AMC scheduling. Therefore, many real-world tasksets may not be schedulable because of their HI-mode utilizations. Thus, AMC-PAStime’s improved performance is significant for practical use-cases.

Fig. 8: Overestimated $C(LO)$





Fig. 9: Checkpoint location



C. Estimation of LO-mode Budget

In our evaluations until now, we estimate a LO-mode budget based on the average execution time of the high-criticality object classification application. In the next set of experiments, we estimate the LO-mode budget of a high-criticality task to be a certain percentage above the average profiled execution time. An increased LO-mode budget for high-criticality tasks benefits AMC scheduling. This is because high-criticality tasks are now given more time to complete in LO-mode, and therefore low-criticality tasks will still be able to execute as well. As a result, the utilization of low-criticality tasks is able to increase.

Suppose that $C(LO)$ is an average execution time estimate for the LO-mode budget of a high-criticality task. Let $(C(LO) + o)$ be an overestimate of the LO-mode budget. As before, AMC-PAStime detects an $X\%$ lag at a checkpoint and predicts the total execution time to be $C(LO) + e$, where e is derived from Equation 6. If the actual LO-mode budget is $(C(LO) + o)$ then AMC-PAStime requests for an extra budget of $(e - o)$, assuming $(e - o) > 0$.

Figure 8 shows that AMC-PAStime still improves utilization for the low-criticality tasks by more than a factor of 3 up to an overestimation of 40%. Overestimation helps in reducing the number of mode switches for AMC scheduling after 40%, as high-criticality tasks have larger budgets in LO-mode.

There is no improvement by AMC scheduling after 60% LO-mode budget overestimation. AMC-PAStime also shows no benefits with increased overestimation, because the LO-mode budget extensions are disapproved by the online schedulability test. Therefore, the system is switched to HI-mode by an overrun of a high-criticality task.

D. Checkpoint Location

We use our modified LLVM compiler in the Profiling phase to determine a viable checkpoint for the high-criticality

Fig. 10: Prediction accuracy




Fig. 11: Offline iterations




Fig. 12: Online iterations





Fig. 13: Budget extension



object classification task. We instrument checkpoints in the `forward_network` function of the Darknet neural network module. We consider four checkpoint locations in the Profiling phase, which are both automatically and manually instrumented. In the Execution phase, we measure performance for each of these checkpoint locations. Figure 9 shows the variation in the number of mode switches against the location of a checkpoint. The x-axis is the approximated division point of a checkpoint location with respect to LO-mode budget. For example, $0.1 \times$ means that the checkpoint is at $(0.1 \times C(LO))$.

We see that the number of mode switches decreases if the location of a checkpoint is more towards the middle of the code. However, a checkpoint near the start and the end of the source code have nearly the same number of mode switches, as with AMC scheduling. A checkpoint near the beginning of a program is not able to capture sufficient delay to increase the LO-mode budget enough to prevent a mode switch. Likewise, a checkpoint near the end of a program is often too late. A HI-mode switch may occur before the high-criticality task even reaches its checkpoint. Hence, a checkpoint at $0.8 \times$ in the source code of a program reduces the number of mode switches by just 1.

E. Effectiveness of a Checkpoint

We now investigate whether a checkpoint is effective in detecting a delayed high-criticality object classification task. We compare the predicted and actual times taken by the high-criticality tasks when LO-mode is extended by AMC-PAStime. This experiment is performed with 8 tasks for 40 jobs each. The initial total LO-mode utilization is 60%, before applying budget extensions.

We show ten of the extended jobs in Figure 10. We see that the predicted execution times are close to the actual budget spent. The predicted times are higher than the spent budgets by just 0.88% on average for this experiment, when the predicted times are more than the actually spent budgets. In Figure 10, Job ID 6, 8 and 10 show lower predicted times than the actual spent budgets. For these jobs, the system is switched to HI-mode because the extended LO-mode is not enough for a task to complete its job. In these cases, the predicted times are smaller than the spent budgets by 0.49%.

This experiment shows that the checkpoint is effectively being used to predict the execution time of a high-criticality task in most cases. The budget extensions are a reasonable estimate of the actual task requirements.

F. Overheads

The main overhead of AMC-PAStime over AMC is the online schedulability test. We first derive an upper bound on this overhead by offline analysis and compare with the experimental measurements.

1) *Offline Upper Bound*: Our offline upper bound is the total number of iterations in solving the response time recurrence relations during the schedulability test in AMC-PAStime. We generate 500 random tasksets of 20 tasks for different initial LO-mode utilizations. Initial LO-mode utilization ranges from 40% to 90%. The utilization of each individual task is generated using the UUnifast algorithm, and each period is taken from 10 to 1000 simulated time units, as done in previous work [10], [24]. As our experimental taskset has a criticality factor ($CF = \frac{C(HI)}{C(LO)}$) of ~ 1.8 , we test with a CF of 1.8 as well.

Among the schedulable tasks with AMC scheduling, we increase the demand in LO-mode budget of the highest priority task. Then, we calculate the total number of iterations needed to determine whether an extension of the budget can be approved by an offline version of Algorithm 1. Here, one iteration is a single update to the response time in any one of the recurrence equations (in Equations 4 and 5) used to test for schedulability.

In Figure 11, we show the maximum number of iterations taken to decide the schedulability of a taskset against a demand of 10 to 80% extra budget in LO-mode by the highest priority task. Each point in the figure represents the maximum iterations across the 500 tasksets to either approve or disapprove of schedulability.

The figure shows the case with 60% initial total LO-mode utilization (before applying budget extensions). The extra demand shown on the x-axis cannot go beyond 80% as the CF is 1.8. As expected, disapproval takes fewer iterations than approval of a taskset.

We have carried out offline analyses with other initial total LO-mode utilizations, as stated above. We have observed the number of iterations to be as high as 120. Therefore, we set 120 as the highest number of allowed iterations for the online schedulability test. When the number of iterations exceed 120 at runtime, we disapprove a LO-mode budget extension. This strategy maintains a safe and known upper bound on the online overhead of AMC-PAStime.

2) *Microbenchmarks*: Each iteration of a response time calculation has a worst-case time of $1\mu s$, for our implementation of Algorithm 1 in LITMUS^{RT}. Therefore, we bound the worst-case delay for the online schedulability test in LITMUS^{RT} at $120\mu s$ for our test cases.

Additionally, the worst-case user-space execution time for our EXTEND_BUDGET macro is $10\mu s$. Hence, the maximum total overhead of an EXTEND_BUDGET macro call, accounting for the schedulability test, is $130\mu s$. The $130\mu s$ overhead is factored into the LO-mode extension inside the LITMUS^{RT} kernel.

Figure 12 shows the number of iterations for the online schedulability test for a taskset of 20 tasks with 60% initial total LO-mode utilization. We see that the offline bound of 120 iterations is much higher than the actually observed number of iterations. Hence, we never need to abandon the schedulability test because of excessive overheads. In addition, disapproval takes less time than approval online, which corroborates our offline observation in Figure 11.

Figure 13 shows the maximum and average times for a LO-mode budget extension decision by our online version of Algorithm 1. It demonstrates that the extension approval decision, along with the schedulability test, takes more time with increasing number of tasks. However, the maximum times are still significantly lower than the offline upper bound of $130\mu s$ for 20 tasks. In general, budget extension and schedulability test overheads can be bounded according to the number of tasks in the system.

VI. RELATED WORK

The problem of determining tight worst-case execution time (WCET) bounds for tasks [49] is compounded by timing variations caused by caches, buses and other hardware features. Recently, mixed-criticality systems (MCSs) [8]–[10], [14], [18], [48] have gained popularity as they allow tasks to have multiple estimates of execution time at different criticality, or assurance, levels. Baruah et al. proposed Adaptive Mixed Criticality (AMC) as a fixed-priority scheduling policy for mixed-criticality systems [10]. AMC dominates other fixed-priority scheduling schemes, such as Static Mixed-criticality with Audsley’s priority assignment and Period Transformation for random tasksets [10], [24].

AMC scheduling affects the QoS of low-criticality tasks by dropping them in HI-mode. Further research work explored adjustments to the task model, including stretching the periods [22], [42]–[45], and using reduced budget in HI-mode [7], [8], [36], [39], to provide improved service to the low-criticality tasks. In this paper, we extend the original AMC task model with a runtime scheduling policy based on the execution progress of a high-criticality task. Other task models are complementary to PASTime’s scheduling approach.

Santy et al. proposed the idea of a task allowance [41]. The allowance is statically calculated, and a possible online implementation of the allowance is provided. In our current work, we do not compute any allowance offline. Rather, we

dynamically decide whether a task is given extra budget in LO-mode based on the observed runtime delay at a checkpoint. PASTime is then able to decide which task’s budget to extend in LO-mode, given the slack in computational resources [18]. In addition, we provide a working implementation of PASTime, coupled with AMC, in LITMUS^{RT}.

The work by Kritikakou et al. uses run-time monitoring and control in mixed-criticality systems, to increase task parallelism [27]–[29]. The authors run high- and low-criticality tasks together and monitor high-criticality tasks at multiple *observations points* embedded into their control flow graphs. If interference from the low-criticality tasks is too prohibitive for the high-criticality tasks, low-criticality tasks are stopped to ensure that the high-criticality tasks meet their deadlines. The authors use static execution time analysis to decide whether to run the low-criticality tasks after an observation point in a high-criticality task. In our work, we dynamically adjust LO-mode budgets of the high-criticality tasks when we detect delays at intermediate checkpoints. We decide about the LO-mode budgets based on the observed progress, instead of using static offline remaining time as used by the other work. In addition, we have implemented an LLVM compiler pass to instrument checkpoints in real-world applications, which have been tested with a PASTime implementation in LITMUS^{RT}. Notwithstanding, Kritikakou et al. provide an important WCET analysis using CFGs for high-criticality tasks.

Previous ideas of progress-based scheduling were proposed to improve GPU performance [4], [26], fairness among multiple threads [20], [21] and to account for instruction cycles [19]. Most of these works run a task in an isolated environment and compare its progress to an online parallel execution with other tasks. Somewhat similar to the motivation behind Jeong et al’s work [26], we also meet the deadlines of high-criticality tasks. However, PASTime uses CFGs to monitor progress rather than application-specific features such as frame processing rates in multimedia applications as done in the other works.

VII. CONCLUSIONS AND FUTURE WORK

This paper presents PASTime, a scheduling strategy based on the execution progress of a task. Progress is measured by observing the time taken for a program to reach a designated checkpoint in its control flow graph (CFG). We integrate PASTime in mixed-criticality systems by extending AMC scheduling. PASTime extends the LO-mode budget of a high-criticality task based on its observed progress, given that the extension does not violate the schedulability of any tasks. We carry out a bounded online schedulability test based on offline response time values. We also provide the response time equations for online calculations. Our extension to AMC scheduling, called AMC-PASTime, is shown to improve the QoS of low-criticality tasks.

Moreover, we provide an algorithm to detect viable locations for program checkpoints. We modify the LLVM compiler to automatically instrument checkpoints for use in task profiling and runtime execution.

We have implemented both AMC and AMC-PAStime in LITMUS^{RT} and compared their performance. While both meet deadlines for all high-criticality tasks, AMC-PAStime improves the average utilization of low-criticality tasks by 1.5–9 times for 2–20 total tasks. For 8 tasks, PAStime increases CPU time to low-criticality tasks until the system reaches 70% LO-mode utilization, after which it converges with AMC. AMC-PAStime is shown to yield improved performance for low-criticality tasks while reducing the number of mode switches. It also has a bounded overhead for its online schedulability test.

In future work, we will explore other uses of progress-aware scheduling in timing-critical systems. We plan to extend the Linux kernel SCHED_DEADLINE policy [31], [32], [35] to support progress-aware scheduling. The aim is to improve the QoS of non-time-critical tasks in Linux while timing critical SCHED_DEADLINE tasks meet their deadlines.

We believe that PAStime is applicable to timing-sensitive cloud computing applications, where it is possible to adjust power (e.g., via Dynamic Voltage Frequency Scaling) based on progress. Application of PAStime to domains outside real-time computing will also be considered in future work.

REFERENCES

- [1] “Big Buck Bunny,” <https://www.bigbuckbunny.org>, 2018.
- [2] “FFmpeg Multimedia Framework,” <https://www.ffmpeg.org/>, 2019.
- [3] “LLVM LoopInfo Class,” https://lvm.org/doxygen/LoopInfo_8h_source.html, Last Accessed: May 2019.
- [4] J. Anantpur and R. Govindarajan, “PRO: Progress-aware GPU Warp Scheduling Algorithm,” in *Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2015, pp. 979–988.
- [5] N. C. Audsley, “On Priority Assignment in Fixed Priority Scheduling,” *Information Processing Letters*, vol. 79, no. 1, pp. 39–44, 2001.
- [6] S. Baruah and A. Burns, “Fixed-priority Scheduling of Dual-criticality Systems,” in *Proceedings of the 21st International conference on Real-Time Networks and Systems*. ACM, 2013, pp. 173–181.
- [7] S. Baruah, A. Burns, and Z. Guo, “Scheduling Mixed-criticality Systems to Guarantee Some Service under All Non-erroneous Behaviors,” in *2016 28th Euromicro Conference on Real-Time Systems (ECRTS)*. IEEE, 2016, pp. 131–138.
- [8] S. Baruah, H. Li, and L. Stougie, “Towards the Design of Certifiable Mixed-criticality Systems,” in *Proceedings of the 16th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2010, pp. 13–22.
- [9] S. Baruah and S. Vestal, “Schedulability Analysis of Sporadic Tasks with Multiple Criticality Specifications,” in *Proceedings of the Euromicro Conference on Real-Time Systems (ECRTS)*, 2008, pp. 147–155.
- [10] S. K. Baruah, A. Burns, and R. I. Davis, “Response-time Analysis for Mixed Criticality Systems,” in *Proceedings of the 32nd IEEE Real-Time Systems Symposium (RTSS)*, 2011, pp. 34–43.
- [11] E. Bini and G. C. Buttazzo, “Measuring the Performance of Schedulability Tests,” *Real-Time Systems*, vol. 30, no. 1-2, pp. 129–154, 2005.
- [12] B. Brandenburg and J. H. Anderson, “Scheduling and Locking in Multiprocessor Real-time Operating Systems,” Ph.D. dissertation, PhD thesis, The University of North Carolina at Chapel Hill, 2011.
- [13] A. Burns and S. Baruah, “Towards a more practical model for mixed criticality systems,” in *Workshop on Mixed-Criticality Systems (colocated with RTSS)*, 2013.
- [14] A. Burns and R. I. Davis, “A Survey of Research into Mixed Criticality Systems,” *ACM Comput. Surv.*, vol. 50, no. 6, pp. 82:1–82:37, Nov. 2017. [Online]. Available: <http://doi.acm.org/10.1145/3131347>
- [15] K. Burr and W. Young, “Combinatorial Test Techniques: Table-based Automation, Test Generation and Code Coverage,” in *Proceedings of the International Conference on Software Testing Analysis & Review*, 1998.
- [16] J. M. Calandrino, H. Leontyev, A. Block, U. C. Devi, and J. H. Anderson, “LITMUS^{RT}: A Testbed for Empirically Comparing Real-time Multiprocessor Schedulers,” in *2006 27th IEEE International Real-Time Systems Symposium (RTSS’06)*. IEEE, 2006, pp. 111–126.
- [17] R. I. Davis, A. Zabus, and A. Burns, “Efficient Exact Schedulability Tests for Fixed Priority Real-time Systems,” *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1261–1276, 2008.
- [18] D. De Niz, K. Lakshmanan, and R. Rajkumar, “On the Scheduling of Mixed-criticality Real-time Task Sets,” in *Proceedings of the 30th IEEE Real-Time Systems Symposium (RTSS)*, 2009, pp. 291–300.
- [19] S. Eyermer and L. Eeckhout, “Per-thread Cycle Accounting in SMT Processors,” *ACM Sigplan Notices*, vol. 44, no. 3, pp. 133–144, 2009.
- [20] J. Feliu, J. Sahuquillo, S. Petit, and J. Duato, “Addressing Fairness in SMT Multicores with a Progress-aware Scheduler,” in *Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2015, pp. 187–196.
- [21] Feliu, Josue and Sahuquillo, Julio and Petit, Salvador and Duato, Jose, “Perf&Fair: A Progress-aware Scheduler to Enhance Performance and Fairness in SMT Multicores,” *IEEE Transactions on Computers*, vol. 66, no. 5, pp. 905–911, 2017.
- [22] C. Gill, J. Orr, and S. Harris, “Supporting Graceful Degradation through Elasticity in Mixed-Criticality Federated Scheduling,” in *Proc. 6th Workshop on Mixed Criticality Systems (WMC)*, RTSS, 2018, pp. 19–24.
- [23] Google, “trucov,” <https://code.google.com/archive/p/trucov/>, 2018.
- [24] H.-M. Huang, C. Gill, and C. Lu, “Implementation and Evaluation of Mixed-criticality Scheduling Approaches for Sporadic Tasks,” *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 13, no. 4s, p. 126, 2014.
- [25] Intel, “Intel® NUC Kit NUC6i7KYK (Skull Canyon): Features and Configurations,” <https://www.intel.com/content/www/us/en/products/docs/boards-kits/nuc/nuc-kit-nuc6i7kyk-features-configurations.html>, 2019.
- [26] M. K. Jeong, M. Erez, C. Sudanthi, and N. Paver, “A QoS-aware Memory Controller for Dynamically Balancing GPU and CPU Bandwidth Use in an MPSoC,” in *Proceedings of the 49th annual Design Automation Conference*. ACM, 2012, pp. 850–855.
- [27] A. Kritikakou, O. Baldellon, C. Pagetti, C. Rochange, and M. Roy, “Run-time Control to Increase Task Parallelism in Mixed-critical Systems,” in *Proceedings of the 26th Euromicro Conference on Real-Time Systems (ECRTS)*, 2014.
- [28] A. Kritikakou, T. Marty, and M. Roy, “DYNASCORE: DYNAMIC Software COntroller to increase REsource Utilization in Mixed-critical Systems,” *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 23, no. 2, 2017.
- [29] A. Kritikakou, C. Rochange, M. Faugère, C. Pagetti, M. Roy, S. Girbal, and D. G. Pérez, “Distributed Run-time WCET Controller for Concurrent Critical Tasks in Mixed-critical Systems,” in *Proceedings of the 22nd International Conference on Real-Time Networks and Systems*. ACM, 2014.
- [30] C. Lattner and V. Adve, “LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation,” in *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*. IEEE Computer Society, 2004, p. 75.
- [31] J. Lelli, G. Lipari, D. Faggioli, and T. Cucinotta, “An Efficient and Scalable Implementation of Global EDF in Linux,” in *Proceedings of the 7th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPert)*, 2011, pp. 6–15.
- [32] J. Lelli, C. Scordino, L. Abeni, and D. Faggioli, “Deadline Scheduling in the Linux kernel,” *Software: Practice and Experience*, vol. 46, no. 6, pp. 821–839, 2016.
- [33] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft COCO: Common Objects in Context,” in *Proceedings of the European Conference on Computer Vision*. Springer, 2014, pp. 740–755.
- [34] Linux, “hrtimers - Subsystem for High-resolution Kernel Timers,” <https://www.kernel.org/doc/Documentation/timers/hrtimers.txt>, Last Accessed: May 2019.
- [35] Linux, “SCHED_DEADLINE Scheduling Policy,” <https://www.kernel.org/doc/Documentation/scheduler/sched-deadline.txt>, Last Accessed: May 2019.
- [36] D. Liu, N. Guan, J. Spasic, G. Chen, S. Liu, T. Stefanov, and W. Yi, “Scheduling analysis of imprecise mixed-criticality real-time tasks,” *IEEE Transactions on Computers*, vol. 67, no. 7, pp. 975–991, 2018.

- [37] D. Liu, J. Spasic, N. Guan, G. Chen, S. Liu, T. Stefanov, and W. Yi, "Edf-vd scheduling of mixed-criticality systems with degraded quality guarantees," in *2016 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2016, pp. 35–46.
- [38] NA, "Will be made available after publication," 2019.
- [39] S. Ramanathan, A. Easwaran, and H. Cho, "Multi-rate Fluid Scheduling of Mixed-criticality Systems on Multiprocessors," *Real-Time Systems*, vol. 54, no. 2, pp. 247–277, 2018.
- [40] J. Redmon, "Darknet: Open Source Neural Networks in C," <http://pjreddie.com/darknet/>, 2013–2016.
- [41] F. Santy, L. George, P. Thierry, and J. Goossens, "Relaxing mixed-criticality scheduling strictness for task sets scheduled with fp," in *2012 24th Euromicro Conference on Real-Time Systems*. IEEE, 2012, pp. 155–165.
- [42] H. Su, P. Deng, D. Zhu, and Q. Zhu, "Fixed-priority Dual-rate Mixed-criticality Systems: Schedulability Analysis and Performance Optimization," in *2016 IEEE 22nd International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 2016, pp. 59–68.
- [43] H. Su, N. Guan, and D. Zhu, "Service guarantee exploration for mixed-criticality systems," in *2014 IEEE 20th International Conference on Embedded and Real-Time Computing Systems and Applications*. IEEE, 2014, pp. 1–10.
- [44] H. Su and D. Zhu, "An Elastic Mixed-criticality Task Model and its Scheduling Algorithm," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2013, pp. 147–152.
- [45] H. Su, D. Zhu, and S. Brandt, "An Elastic Mixed-Criticality Task Model and Early-Release EDF Scheduling Algorithms," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 22, no. 2, p. 28, 2017.
- [46] M. M. Tikir and J. K. Hollingsworth, "Efficient Instrumentation for Code Coverage Testing," in *ACM SIGSOFT Software Engineering Notes*, vol. 27, no. 4. ACM, 2002, pp. 86–96.
- [47] M. Vanga, A. Bastoni, H. Theiling, and B. B. Brandenburg, "Supporting Low-Latency, Low-Criticality Tasks in A Certified Mixed-Criticality OS," in *Proceedings of the 25th International Conference on Real-Time Networks and Systems*. ACM, 2017, pp. 227–236.
- [48] S. Vestal, "Preemptive Scheduling of Multi-criticality Systems with Varying Degrees of Execution Time Assurance," in *Proceedings of the 28th IEEE Real-Time Systems Symposium (RTSS)*, 2007, pp. 239–243.
- [49] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra *et al.*, "The Worst-case Execution-time Problem - Overview of Methods and Survey of Tools," *ACM Transactions on Embedded Computing Systems*, vol. 7, no. 3, p. 36, 2008.