

CAPA: A Framework for Contention-Aware and Progress-Aware Multi-Core Real-Time Systems

Weifan Chen*, Ivan Izhbirdeev*, and Renato Mancuso

Abstract—Real-time systems face significant challenges in managing shared resources on multi-core platforms while maintaining temporal predictability. This paper introduces the CAPA (Contention-Aware Progress-Aware) framework for real-time systems to export and leverage runtime progress information to enable informed resource management decisions. CAPA represents a practical model for tracking task progress, extending existing Timely Progress Assessment (TPA) techniques to support complex control flows and concurrent execution on commercial multi-core hardware. Building on this model, it is possible to design progress-aware multi-core schedulers capable of dynamically regulating task execution to meet timing constraints. We implement CAPA on a commercial off-the-shelf (COTS) platform and evaluate its performance. Results demonstrate CAPA’s ability to provide controlled performance degradation, meet timeliness constraints, and improve schedulability. Thus, CAPA represents a significant step forward toward practical, contention-aware scheduling in multi-core real-time systems.

Index Terms—Adaptive Control, Real-Time, Scheduling, Resource Contention, Multicore

I. INTRODUCTION

THE quality of service in many systems degrades nonlinearly as demand increases, particularly as the system’s performance bottleneck approaches saturation. This concept has profound implications in computing systems. When shared hardware resources are accessed by multiple contenders (e.g., processes running on different cores), the execution time of each contender can become highly variable, potentially compromising the temporal correctness of the system, a critical concern in real-time applications.

Contention and Unpredictability. Unpredictable performance degradation due to contention is a well-recognized issue in multicore systems [1], [2]. Without proper shared resource management, deploying hard real-time workloads on complex high-performance embedded systems becomes unsafe. On the other hand, over-provisioning can result in system underutilization. This creates a tension between ensuring predictability and maximizing resource utilization.

Several shared resource management paradigms have been studied. In broad categories, (1) strict partitioning/regulation approaches attempt to mitigate contention by enforcing static per-core/per-task reservations of performance-critical hardware to reduce sharing [3]–[6]; (2) profile-driven workload consolidation approaches aim to use offline-acquired data about the expected behavior of applications to tune resource partitioning [7]–[10]; (3) profile-driven dynamic management combines offline analysis and online tracking of performance

metrics to adjust partitioning parameters dynamically [11]–[14]. (4) Design time analysis or principles that quantify or manage interference effects [15], [16].

Progress-aware Management. Recently, a new line of work has emerged: progress-driven workload management. The idea is to enable a system to live-track the fraction of work that a task accomplishes in relation to the total needed to complete. Then, contention-induced slowdowns can be promptly assessed, and management decisions can be taken accordingly and well ahead of task deadlines. Early proposals of progress-driven workload management relied on software instrumentation of the tasks [17]–[19], on firmware instrumentation [20], or on custom hardware support [21]. More recently, however, the work in [22] demonstrated that live progress tracking, a.k.a. *Timely Progress Assessment* (TPA), as coined by Chen et al. [22], can be achieved on commercially available hardware, with low overhead, and without any changes to application binaries and OS/hypervisor-level software. While foundational, this work has three important shortcomings. (1) Only tasks whose control flow can be determined prior to execution can be monitored accurately. (2) The original progress monitoring approach could not live-track tasks on multiple cores. (3) The presented management strategy only reacts to tasks suffering unsatisfactory progress, however, does not exploit the situation where some tasks progress faster than expectation.

Contribution. In this paper, we acknowledge that there is immense potential in TPA-based workload management and build upon the state-of-the-art in three main directions. (1) We broaden the progress tracking capability to incorporate tasks with complex control flows at runtime and in multicore settings. (2) Importantly, this work pioneers a fractional, linear definition of progress with respect to the execution path leading to the worst-case execution time (WCET). This practically estimates the progress even when the execution does not follow the worst-case path. (3) We reason about integrating live progress information into the design of multicore schedulers. In particular, we propose a class of *timely progress-aware multicore schedulers* capable of regulating the pace at which concurrent applications progress to meet temporal constraints. In summary, our contributions are the following:

- 1) Propose a Contention-Aware-Progress-Aware (CAPA) task model, namely CAPATASK, that can be used to track the progress in percentage of complex control flow.
- 2) Design, implement, and test on commercial hardware a system-level CAPA support, namely CAPASYS. CAPASYS showcases for the first time that TPA of concurrent tasks on a COTS multicore platform is possible by leveraging existing hardware support;
- 3) Propose and evaluate a first class of schedulers that can

*These authors contributed equally to this work.




Fig. 1: Two jobs illustrated in progress- Fig. 2: The left shows τ_1 misses its deadline due to contention. The right shows the same jobs but τ_0 pauses in between to achieve the timeliness of both.

leverage CAPASYS to enact progress-driven multicore scheduling with the purpose-built CAPASIM simulator.

The rest of the paper is structured as follows. Section II provides a motivating example, while Section III reviews the related work. Section IV presents the CAPATASK. Next, Section V describes key design elements of CAPASYS, with implementation details and demonstrating CAPASYS in action on a COTS platform. Section VI describes CAPASIM and presents evaluations of CAPASYS employed atop traditional multicore schedulers. The paper concludes in Section VII.

II. MOTIVATING EXAMPLE

Can runtime progress information unlock better system resource management? To better illustrate the intuition behind, we introduce a *progress-aware diagram*. Two examples are shown in Fig. 1. In the diagram, x-axis represents the time ticks. In each vertical slice at a time tick, the solid color represents completed progress. The top of the diagram illustrates a job completes around half progress at $t = 2.5$, while the bottom shows a job completion at $t = 4$.

Progress Management. Consider two jobs on separate cores. They may contend for shared resources, causing execution slowdowns—visible as flatter progress slopes due to reduced progress rates. In Fig. 2-left, both jobs finish at $t = 4.5$, but τ_1 misses its deadline without intervention. In contrast, Fig. 2-right shows that with runtime progress awareness, the system can temporarily pause τ_0 , allowing τ_1 to accelerate and meet its deadline. τ_0 then resumes and completes on time. These scenarios are demonstrated in both our implementation (Section V-D) and simulation (Section VI-D).

III. RELATED WORKS

Two aspects play a considerable role in the example above: (i) inter-core shared resource contention can affect progress rate; (ii) the progress can be extracted at runtime. Both are foundational to this work and have been extensively studied by the real-time community.

Shared Resource Contention Aware Task Models. In most of these models, a job is typically treated as an “indivisible atom,” i.e., its impact on the performance of the co-runners is uniform regardless of which phase the job resides. For example, Aceituno et al. [23] factor out the worst-case interference time if the application were to encounter resource contention. Bini [24] assumes that the current mode of the system, not those of tasks, determines each task’s progress rate. Kim et al. [25] directly prevent tasks with considerable interference running concurrently. Unfortunately, considering tasks as “indivisible” units of work fails to reflect that jobs at different phases impact shared resources differently.

Workload Profiling. Numerous works have illustrated that shared resource usage can vary throughout the execution of a task, and that profiling is crucial for informed management. For example, approaches to inspect the cache activity as concurrent tasks execute were proposed [7], [8]. Roozkhosh et al. [9] inspect the activity in main memory at the cache-line granularity. Sohal et al. [10] demonstrate using memory consumption envelopes for partition tuning. The ability to inspect at runtime the performance of a target application while minimally disrupting its temporal behavior has been thoroughly investigated in [26].

Progress-aware Task Models. Various expressions for progress exist, e.g. the number of instructions retired [13], [27], deadline partitioning [28], [29]. However, the most relevant approach to this work involves TPA [17]–[22], whose essence is the selection of specific addresses within the task through offline profiling. These addresses are typically the entries/exits of single-entry-single-exit code blocks, and addresses inside loops. Then, a runtime monitoring mechanism tracks the time when the addresses are reached, thereby adaptively assessing progress timeliness. Most of the works [17], [18], [20], [22] aim at protecting the timeliness of a single task in a multicore environment via TPA by throttling the activities of the co-runners on different cores. Only Kritikakou et al. [19], [21] investigate protecting the timeliness of multiple tasks in a multicore environment. Nevertheless, these works assume that the execution order of visits to the selected addresses can be determined prior to runtime, and none define a notion of universal progress. As a result, the progress of different tasks cannot be effectively compared. This work addresses these limitations.

IV. THE CAPA TASK MODEL

Contention-Aware-Progress-Aware task model (CAPATASK for short) is presented in this section. The goal of CAPATASK is to express the progress of a task as a percentage. Despite its simple form, the definition of progress must remain robust even for tasks with complex control flow, e.g., in the presence of branches that cannot be statically predicted before runtime, or in the presence of loops where the number of iterations is only known at runtime. The key intuition to achieve this is:

Intuition I: Normalizing progress over all possible execution paths with respect to the worst-case path identified during offline task profiling or static analysis.

Task profiling, in this context, refers to measuring execution times at specific control flow “observation points” during its

execution as the task is repeatedly executed in a contention-free environment during the analysis phase. This procedure is used in all existing TPA-related methods [17]–[22]. We use the terms *contention-free*, *without contention*, *solo*, and *isolation* interchangeably to describe the condition where a task is running alone in the system. While WCET traditionally refers to the end-to-end execution time of a task, we also use WCET to denote the execution time of a portion of a task. When doing so, the specific context will make this distinction clear. While we adopt a measurement-based approach, equivalent definitions apply for static analysis-based methods.

A. Task Model

The proposed CAPATASK integrates support for trackable progress within traditional real-time taskset semantics. Each CAPATASK task τ_i is a sequence of jobs with sporadic or periodic release where the generic job $\tau_{i,j}$ can be described as the tuple $\tau_{i,j} = \{T_i, D_i, \mathcal{G}_i\}$. Here, T_i represents a minimum inter-arrival time in the case of sporadic tasks, or an exact release period for periodic tasks. Next, D_i is the relative deadline, i.e., a constraint on the job’s response time once it is released. Finally, \mathcal{G}_i is a Timed Milestone Graph (TMG) proposed in [22]. A recap of TMG is presented in the next section. In CAPATASK, TMGs replace typical WCET information. CAPATASK is not intended to replace WCET analysis; instead, it leverages WCET estimates—whether from static or measurement-based methods—for progress estimation. Moreover, it mitigates WCET pessimism by detecting surplus progress (or lack thereof) for better resource management.

B. Recap: Timed Milestone Graph

CAPATASK is inspired by the Timed Milestone Graph (TMG) and is a transformation over an existing TMG. In this work, we follow the same method as in [20] to construct the initial TMG and then transform it into CAPATASK. This section presents the existing definition and construction method of TMG. Conceptually, a TMG retains two types of timing information acquired through profiling/analysis: (1) a *tail time* $\mathcal{T}_{\text{tail}}$ indicates the WCET for execution to reach a specific instruction; (2) a *move time* $\mathcal{T}_{\text{move}}$ indicates the maximum time it takes for the execution to transfer from one specific instruction to another.

Definition I: Timed Milestone Graph

A *Timed Milestone Graph* (TMG) is a weakly connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} and \mathcal{A} are sets of vertices and arcs respectively. The graph has the following properties:

- There exists only one vertex $v_{\text{entry}} \in \mathcal{V}$ that has no predecessors—in-degree is zero;
- There exists only one vertex $v_{\text{exit}} \in \mathcal{V}$ that has no successors—out-degree is zero.

Additionally, two maps are defined:

- $\mathcal{T}_{\text{tail}} : \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ assigns a **tail time** to each vertex;
- $\mathcal{T}_{\text{move}} : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ assigns a **move time** to each arc.

Typically, a vertex $u \in \mathcal{V}$, a.k.a., a *milestone* in this context, represents an instruction location, such as the address of the

first assembly instruction in τ . At runtime, we assert that a milestone u is reached as the program counter (PC) points to u . We also name it a vertex/milestone-hit.

Experimental Construction. Like all other TPA methods, $\mathcal{T}_{\text{tail}}$ and $\mathcal{T}_{\text{move}}$ are derived by collecting the behavior of the task across multiple runs/samples in solo execution. A milestone might be reached multiple times in the same run if it is part of a loop. Let’s define the random variable $T(u; i, j)$ indicating the time elapsed since the job release till u is reached for the j^{th} time in the i^{th} run. Then $\mathcal{T}_{\text{tail}}(u)$ is assigned to be the maximum in all i, j for $T(u; i, j)$:

$$\mathcal{T}_{\text{tail}}(u) = \max_{i,j} \{T(u; i, j)\}. \quad (1)$$

Thus, the tail time $\mathcal{T}_{\text{tail}}(u)$ indicates the worst (experimentally observed) execution time to reach u for the very last time. By construction, $\mathcal{T}_{\text{tail}}(v_{\text{entry}}) = 0$.

The move time $\mathcal{T}_{\text{move}}(u, v)$ is a *conservative* estimation of the maximum time for the execution to transfer from u to v .¹ The largest time elapse seen during the profiling is taken as its move time. Formally, let $T(u, v; i, j)$ be the time elapsed for the execution to transfer from u to v for the j^{th} time in the i^{th} run. Then

$$\mathcal{T}_{\text{move}}(u, v) = \max_{i,j} \{T(u, v; i, j)\}. \quad (2)$$

Milestone Selection. The milestone selection method we use in this paper is akin to what is described in [20]. We hereby present the key steps. (1) Given the task’s binary, construct its control flow graph. Several tools are available, such as `angr` [30] and `GNU cflow`. (2) In the main function, tentatively place initial milestones on all the return addresses of top-level function calls and all the successors of any strongly connected components (SCC). (3) Measure $\mathcal{T}_{\text{tail}}$ and $\mathcal{T}_{\text{move}}$ for these initial milestones. If the temporal distance between any two subsequent milestones is too small, then randomly remove one of them. If too large, place more milestones in between. Two cases exist in the latter scenario: (i) If the successor milestone is the return address of a function call, then unfold the function and repeat from step (2), thus placing milestones within the unfolded function instead of `main`; (ii) If the milestone corresponds to the successor of an SCC, this typically implies that one of the loops in the SCC takes a large number of iterations. In this case, randomly place one milestone inside the SCC, and repeat from step (3). The procedure above is applied iteratively until a satisfactory TMG is obtained. This approach is empirically practical because a TMG with fine granularity can usually be obtained by placing one or two milestones within key SCCs that the outlined methodology identifies by progressively annealing the distance between milestones. The iterative nature of the procedure also allows a user to control the granularity of the TMG on demand. In Table II, we quantitatively summarize the TMGs of the benchmarks used in this work.² In the next section, we present our method to transform a TMG into a CAPATASK,

¹Precisely speaking, it is the execution to transfer from u to v without reaching any other vertices during the transfer.

²Appendix I offers a detailed discussion on the offline cost of TMG construction.




Fig. 3: A TMG and its transformed graph. Each SCC (in dashed boxes) is replaced with arcs (in red).

which not only provides a normalized progress value at each vertex, but also conservatively evaluates the progress when control flow diverges at runtime.

C. Normalized Progress over TMG

To enhance readability, we provide intuitive explanations at key stages of the transformation and include a glossary in Table I to facilitate reading.

Symbol	Description
\mathcal{G}	A timed milestone graph (TMG)
\mathcal{G}^{DAG}	An acyclic TMG, a.k.a., <i>birdseye</i> TMG
w	A generic walk or subwalk on \mathcal{G}
\tilde{w}	Walk leading to solo WCET; a.k.a., the <i>nominal</i> walk
$C(w)$	The WCET in a solo run of executing walk w
C_i	The WCET in a solo run from start to i^{th} vertex-hit
$\Delta C_i = C_i - C_{i-1}$	The solo WCET between $(i-1)^{\text{th}}$ and i^{th} vertex-hit
$\chi_i \in [0, 1)$	Progress between $(i-1)^{\text{th}}$ and i^{th} vertex hits
$P(v_i) \in [0, 1]$	Progress at i^{th} vertex hit
$P(t) \in [0, 1]$	Progress at time t
$\tilde{P}(t) \in [0, 1]$	Nominal progress over time t
\tilde{k}	Nominal progress rate (reciprocal of solo WCET)

TABLE I: Summary of notation.

A TMG captures all the potential control flow transfers. However, it cannot predict whether a branch is taken. This accompanies the presence of loop components, making expressing progress as a single percentage quantity hard.

Intuition II: Loops complicate progress tracking, but if we focus only on entry and exit of a loop—ignoring iterations—the problem becomes simpler. Progress within loops can then be addressed separately.

A strongly connected component (SCC) is a subgraph in which every vertex is reachable from every other vertex. A loop is just a simple form of SCC. If no SCCs are present, building the desired progress would be simpler, because the TMG will be acyclic. Notice that a vertex is either in an SCC or not. We can “encapsulate” each SCC and replace it with a set of arcs that reflect the net result of the control transfer. The transformation is illustrated in Fig. 3. With that, the problem becomes: (1) how to define the progress over an acyclic TMG, and (2) in an SCC. We now focus on solving (1). Said SCC encapsulation is intuitive. We now define it formally. Let S_i denote the set of all vertices in an SCC. A vertex u is a predecessor of the SCC if and only if $u \notin S_i$ and there exists an arc (u, v) such

that $v \in S_i$. A vertex v is a successor of the SCC if and only if $v \notin S_i$ and there exists an arc (u, v) such that $u \in S_i$. Let $\mathcal{P}(S_i)$ be the set of predecessors of S_i , similarly, $\mathcal{S}(S_i)$ for the set of successors. The resulting transformed graph, called the **birdseye TMG** \mathcal{G}^{DAG} , is formally defined as follows. The acyclic nature of \mathcal{G}^{DAG} facilitate progress definition.

Definition II: Birdseye TMG

Given a TMG

$$\mathcal{G} = \{\mathcal{V}_{\text{scc}} \cup \mathcal{V}_{\text{scc}}^c, \mathcal{A}_{\text{scc}} \cup \mathcal{A}_{\text{scc}}^c\} \quad (3)$$

where \mathcal{V}_{scc} is the set of vertices of some SCCs; $\mathcal{V}_{\text{scc}}^c$ is the complementary set; \mathcal{A}_{scc} is the set of arcs, each of which has at least one endpoint within some SCCs; $\mathcal{A}_{\text{scc}}^c$ is the complementary set. A birdseye TMG

$$\mathcal{G}^{\text{DAG}} = \{\mathcal{V}_{\text{scc}}^c, \mathcal{A}_{\text{scc}}^c \cup \mathcal{A}'\} \quad (4)$$

is an acyclic weakly connected graph, where \mathcal{A}' consists of additional arcs added by the following procedures. For an SCC S_i in \mathcal{G} , the arcs added are defined as:

$$\mathcal{A}_i = \{(u, v) \mid u \in \mathcal{P}(S_i) \wedge v \in \mathcal{S}(S_i)\} \quad (5)$$

With $\mathcal{T}_{\text{move}}(u, v) = \mathcal{T}_{\text{tail}}(v) - \mathcal{T}_{\text{tail}}(u)$. All $\mathcal{T}_{\text{tail}}$ remain the same. Then \mathcal{A}' is the union of all \mathcal{A}_i

$$\mathcal{A}' = \bigcup \mathcal{A}_i \quad (6)$$

Walk over TMG. A TMG or Birdseye TMG do not contain the information regarding actual execution for a concrete execution. For any concrete execution, it can be represented as a sequence of vertices on a TMG or Birdseye TMG. We name the said sequence as walk. A walk is a path corresponding to the control transfers in a concrete run, defined as

Definition III: Walk over TMG

Given a TMG $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, a walk over \mathcal{G} is a sequence of vertices $w = (v_1, \dots, v_N)$ such that for all $i \in \{1, \dots, N-1\}$, the arc $(v_i, v_{i+1}) \in \mathcal{A}$.

The subscript denotes the order in which PC visits a vertex. For example, v_i is the i^{th} vertex reached by PC. A vertex can be reached more than once due to the presence of loops. The walk is well-defined on both \mathcal{G} and \mathcal{G}^{DAG} . Now define a function $f_w : \mathcal{V} \times \mathcal{V} \rightarrow W$ that maps two vertices to the set of all possible walks from the former to the latter. On a TMG, $f_w(u, v)$ can be an infinite set due to the existence of SCC(s), but on \mathcal{G}^{DAG} the set is always finite. The finiteness allows us to find the walk that results in contention-free WCET.

Intuition III: The acyclic property of a birdseye TMG allows us to enumerate all possible walks³, inspect the timing information, and find the walk that results the contention-free WCET.




Fig. 4: Left: red arcs indicate \tilde{w} ; evaluate the progress along with the alternative walk indicated by the green arrow; Middle: from the sub-graph (solid line), find the new nominal walk w' , represented by the solid red lines; Right: take the relevant sub-walk w'' , represented by the solid red lines.

Given a \mathcal{G}^{DAG} , and denote $W^{\text{DAG}} = f_w(v_{\text{entry}}, v_{\text{exit}})$ as the set of walks from entry to exit. We want to find the walk $\tilde{w} \in W^{\text{DAG}}$ that corresponds to the path which leads to the contention-free WCET. We name \tilde{w} as the **nominal walk**. In the scope of this work, “nominal” means the worst-case behavior in a contention-free environment. Now we present the **nominal walk normalization procedure** which (1) identifies \tilde{w} ; (2) calculates the ratio of total progress made between two vertices in \tilde{w} ; (3) defines the progress at each of the vertices in \tilde{w} . To identify \tilde{w} , we need to evaluate the nominal execution time for a generic walk as follows. For a walk $w \in W^{\text{DAG}}$, denote the nominal execution time when PC reaches v_i as C_i whose value is defined inductively

$$C_0 = 0 \quad (7)$$

$$C_{i+1} = \min\{\mathcal{T}_{\text{tail}}(v_{i+1}), C_i + \mathcal{T}_{\text{move}}(v_i, v_{i+1})\} \quad (8)$$

Eq. 8 takes the minimum of the two, because $\mathcal{T}_{\text{move}}$ is the longest time observed to transfer between v_i and v_{i+1} across all profiling data, and, unlike $\mathcal{T}_{\text{tail}}$, it does not take into account potential temporal dependencies. Thus, merely summing up all $\mathcal{T}_{\text{move}}$ might result in a pessimistic estimation. The use of $\mathcal{T}_{\text{tail}}$ here not only mitigates the pessimism, but also ensures safety, because $\mathcal{T}_{\text{tail}}$ captures the contention-free WCET until execution reaches v_{i+1} .

For all possible walks from entry to exit, the one with the largest nominal time is chosen as the nominal walk, formally

$$\tilde{w} = \underset{w \in W^{\text{DAG}}}{\text{argmax}}\{C(w)\} \quad (9)$$

Progress over Nominal Walk. Now define the progress when the execution follows \tilde{w} . Define $\chi_i \in [0, 1]$ as the *fractional progress* contributed by the execution from v_{i-1} to v_i :

$$\chi_i = \frac{\Delta C_i}{C(\tilde{w})}, \quad (10)$$

where $\Delta C_i = C_i - C_{i-1}$. Thus, if at runtime the task’s execution follows exactly \tilde{w} and has reached v_j , then its

current progress $P(v_j)$ is the summation of each fractional progress contribution:

$$P(v_j) = \sum_{i=0}^j \chi_i \quad (11)$$

This definition of progress has the following properties. First (1) it is normalized to $[0, 1]$ with

$$P(v_{\text{entry}}) = 0 \quad (12)$$

$$P(v_{\text{exit}}) = 1 \quad (13)$$

Second, (2) if the execution does follow \tilde{w} , the progress remains in a linear relationship w.r.t. the task’s worst-case execution time in isolation. This concludes the nominal walk normalization procedure.

Progress over Alternative Walks. As part of the progress normalization procedure, each arc along \tilde{w} is assigned a fractional progress χ_i , and at each vertex v_i , $P(v_i)$ is also well defined. However, progress for arcs and vertices along alternative walks is not yet defined. To define this quantity, we make the following observation.

Intuition IV: An alternative walk always has at least one vertex down the path shared by \tilde{w} . When reaching a shared vertex, the current progress must be the same as if the execution had been following \tilde{w} thus far.

We therefore present an iterative algorithm to set values for unassigned arcs and vertices. The idea is to progressively focus on arcs (u, v) where the predecessor u is already associated a notion of progress, while the successor’s progress $P(v)$ is still undefined. For the very first call of the algorithm, all arcs (u, v) such that $u \in \tilde{w} \wedge v \notin \tilde{w}$ are eligible—see left side of Fig. 4. Consider all paths starting from u where v is visited next. All of them will eventually merge into some v_{merger} that already has progress assigned. For example, v_{merger} in Fig. 4 is u' . During the first iteration of the algorithm, v_{merger} is guaranteed to be a vertex along \tilde{w} , but the subsequent iteration of the algorithm can merge into a vertex that has been assigned a progress value even if not necessarily along \tilde{w} . Among all such paths, we pick the path whose v_{merger} has the least progress. If multiple paths with the same v_{merger} exist, we reuse the aforementioned nominal walk normalization procedure to pick the one with the largest nominal time—i.e., the worst-case sub-path starting with arc (u, v) . We can form a new birdseye TMG $\mathcal{G}_{(u,v)}^{\text{DAG}}$ which includes u as its entry vertex, the arc (u, v) , and all vertices and arcs reachable from v through some path to v_{merger} . Next we apply nominal walk normalization procedure on $\mathcal{G}_{(u,v)}^{\text{DAG}}$ to identify nominal walk w' of $\mathcal{G}_{(u,v)}^{\text{DAG}}$. Before applying the procedure, we need to subtract $\mathcal{T}_{\text{tail}}(u)$ from the $\mathcal{T}_{\text{tail}}$ of each vertex in $\mathcal{G}_{(u,v)}^{\text{DAG}}$, because $\mathcal{T}_{\text{tail}}$ is the time elapsed since the execution of the TMG.

Let u' be the first vertex where convergence, i.e., the v_{merger} , occurs (center of Fig. 4). We want the sub-walk w'' from u to u' (right side of Fig. 4) to contribute the same amount of progress as if the original nominal walk \tilde{w} from u to u' , i.e., the progress contribution equals to $P(u') - P(u)$, where $P(u')$

³This enumeration does not suffer from exponential complexity in practice, because the number of vertices and their out-degree is generally small for a TMG. A TMG with less than 20 vertices can already offer satisfactory monitoring granularity [20], [22].

is the progress at u' following \tilde{w} . Now, we treat w'' as a TMG, and apply the nominal walk normalization procedure again to calculate the fractional progress $\chi_i^{w''}$ for each edge in w'' (Eq. 10). The superscript w'' indicates this fractional progress is based on the TMG w'' . The final fractional progress from v_{i-1} to v_i along the alternative path is

$$\chi_i = \chi_i^{w''} (P(u') - P(u)) \quad (14)$$

With χ_i assigned, the progress at each vertex can also be assigned according to Eq. 11. This algorithm is invoked iteratively until the progress at all the vertices has been assigned. At runtime, when the PC hits a vertex v on \mathcal{G}^{DAG} , the application's progress is $P(v)$. The progress has the following properties: For any paths leading to v , the sum of all χ_i along that path equals to $P(v)$, i.e., the progress at each vertex is unique regardless of path taken, and the progress always monotonically increases due to $\chi_i > 0$.⁴

Progress over an SCC. Now that the progress over \mathcal{G}^{DAG} is well-defined. We focus our attention on the progress while executing inside an SCC. Given the arbitrary control flow, the execution can enter and exit the SCC at any predecessors and successors, and iterates inside SCC for a nondeterministic number of times. To define the progress under this circumstance, we realize

Intuition V: While the execution is inside an SCC, treat it as if it will eventual leave the SCC at the successor which has the least progress.

Despite the many possible control flows within an SCC, this complexity does not affect the reasoning at the level of birdseye TMG, because any SCCs are *squashed* into arcs. Thus, as per intuition above, when the execution enters the SCC from a predecessor, we can conservatively assume that it will leave the SCC via the successor with the least progress as calculated in Eq. 11. At runtime, if it exits via other successors, it is always possible to inflate the progress to match the corresponding successor. This guarantees that the worst-case solo progress is never overestimated.

Assume the execution enters it via predecessor u , and the successor corresponding to the least amount of progress is v . Denote the least possible progress contributed by the SCC as $\chi_{\text{sc}} = P(v) - P(u)$. Then the question becomes how to distribute χ_{sc} while the execution is inside the SCC. Similar to Eq. 10 and Eq. 14, where the denominator is the solo worst-case time elapsed for \tilde{w} or some sub-walk, we also need a similar quantity for the SCC, \mathcal{T}_{sc} . We can conservatively approximate \mathcal{T}_{sc} as

$$\mathcal{T}_{\text{sc}} = \mathcal{T}_{\text{tail}}(v) - \mathcal{T}_{\text{tail}}(u). \quad (15)$$

Inside an SCC, the progress contributed by the execution from the $(i-1)^{\text{th}}$ to i^{th} vertex-hit is inductively calculated as

$$\chi_i = \frac{\Delta A_i}{\mathcal{T}_{\text{sc}}} \chi_{\text{sc}} \quad (16)$$

where

$$A_0 \leftarrow 0 \quad (17)$$

$$A_{i+1} \leftarrow \min\{\mathcal{T}_{\text{sc}}, A_i + \mathcal{T}_{\text{move}}(v_i, v_{i+1})\} \quad (18)$$

The intuition is that Eq. 18 ensures ΔA_i is at most some $\mathcal{T}_{\text{move}}$, i.e., the longest time PC takes to move from one vertex to the other. The ratio between ΔA_i and \mathcal{T}_{sc} is then used to assign a fraction of χ_{sc} to the progress. The upper bound of A_i is \mathcal{T}_{sc} , which makes the progress estimation conservative.

If at runtime a loop exceeds the maximum number of iterations ever seen during the profiling phase, Eq. 18 will no longer increase, indicating that no further progress is being made. This is the desired behavior because, for each milestone hit, the progress and the job's CPU time are compared to determine whether the execution satisfies the temporal requirement (see Section V-D). In the event of unexpected higher loop iterations, the ever-increasing CPU time compared with stagnant progress will signal a temporal abnormality.

Effective and Nominal Progress The three expressions for χ_i Eq. 10, 14, and 16 define how to evaluate the progress when the execution is on the worst-case path, on some alternative paths, and inside some SCCs. The evaluation is conducted when PC reaches any of the vertices

$$P(v_0) = 0 \quad (19)$$

$$P(v_{i+1}) = P(v_i) + \chi_{i+1} \quad (20)$$

We name the above progress evaluation as **effective progress**, which is a function of the i^{th} reached vertex.⁵ We can also define the effective progress as a function of wallclock time, $P(t)$. Assume the j^{th} vertex-hit is the latest one prior to t . We conservatively define

$$P(t) = P(v_j) \quad (21)$$

In contrast, we name the progress made as if the execution takes the worst-case path without contention as **nominal progress** \tilde{P} . The property of CAPATASK ensures that the \tilde{P} is linear to execution time. Assume the contention-free WCET is \tilde{C} , then the effective nominal progress rate is $\tilde{k} = 1/\tilde{C}$. Then the expression for nominal progress is

$$\tilde{P}(t_{\text{CPU}}) = t_{\text{CPU}}/\tilde{C} = \tilde{k}t_{\text{CPU}} \quad (22)$$

which is a function of processor time t_{CPU} , i.e., the time spent by the task under analysis while actively running on the CPU. In Section V-D and Section VI, we shall see how to leverage $P(v_i)$ and $\tilde{P}(t)$ to compute slack at runtime to achieve progress-aware scheduling.

V. CAPASYS DESIGN: SYSTEM-LEVEL CAPA SUPPORT

This section presents our design of CAPASYS, a runtime infrastructure to instantiate the CAPATASK on one commercial multi-core hardware platform. We first review key platform assumptions (Section V-A), next we describe the main responsibilities of CAPASYS and the corresponding implementation details (Section V-B and V-C). Beyond implementing progress

⁴A proof is provided in Appendix II.

⁵The computation of effective progress at runtime is also summarized as pseudo-code in Appendix III.




Fig. 5: A TMG used by tracer at runtime. From left to right, three snapshots are provided in temporal order: (1) PC is somewhere on the arc (u_1, u_2) , and the TU is waiting for the event indicating u_2 was hit. (2) When PC reaches u_2 , the TU notifies the system and the tracer starts to re-configure the TU. (3) While execution is on the (u_2, u_3) arc, the TU is waiting for u_3 or u_4 to be the next hit.

tracking, CAPASYS also uses the progress information exposed by CAPATASK to add support for progress-driven CPU management primitives, namely **self-halting** and **assurance halt**. Finally, we evaluate the overhead and capabilities of CAPASYS in Section V-D.

A. Assumptions for CAPASYS Deployment

For a practical instantiation of CAPASYS, following assumptions are made. (i) Black-box binaries of the task are available for offline analysis. (ii) The target is a multi-core platform with high-performance cores. The high-performance cores hosting the managed tasks are called *traces*. (iii) Spare (co-)processors exist to instantiate the CAPASYS runtime. The latter, referred to as *tracers*, can have a low-power/low-performance profile compared to the traces. (iv) The traces can conduct self-hosted tracing, i.e., it can generate an execution trace and save it to on-chip memory using a hardware Trace Unit (TU). A trace reports the processor control flow at the level of basic blocks, as described more in-depth in [22]. (v) We assume that the TMGs have been constructed and focus on the instantiation and use of the CAPATASK.

B. Multi-core CAPASYS Responsibilities

The proposed CAPASYS implements multi-core support for TPA. It has three key responsibilities.

(R1) Milestone Detection and TU (Re)configuration. TUs can inform the system when the PC of a given CPU hits a pre-defined virtual address. Nevertheless, the number of addresses that can be concurrently monitored is limited. To overcome this limitation, Fig. 5 depicts the approach used by CAPASYS. As shown in the figure, the module uses the TMG to (re)configure the TUs to only monitor the vertices that can be potentially hit next, i.e., the immediate successors of the current milestone. Thus, CAPASYS must enable efficient milestone detection and TU reconfiguration.

(R2) Multi-core Trace Data Processing. A TU is a per-core asset. When properly configured, each TU emits trace packets upon a vertex hit. Typically, trace packets from different TUs are multiplexed before being routed downstream. CAPASYS must demultiplex the trace stream to recover individual packets




Fig. 6: Job releases at t_0 . The slopes of nominal progress and setpoint are \tilde{k} and $\alpha\tilde{k}$ respectively. The vertical gap between the green line and setpoint is β . The effective progress is the blue line. At t_1 , the slack (vertical gap between the blue line and setpoint) is larger than β . Thus the task self-halts until t_2 . Job completes at t_3 .

and identify the originating core, task, and vertex associated with the hit. Next, the event is timestamped using the cycle counter, and TU reconfiguration is triggered.

(R3) Progress-driven CPU Management. CAPASYS also exports key progress-driven CPU management capabilities, namely (1) *self-halting* and (2) *external blocking*. **Self-halting** refers to the ability of CAPASYS to slow down the progress of a task on a given CPU if its progress is ahead of schedule. This allows for relieving contention effects on shared resources when it is safe to do so. **External blocking** refers to the ability of CAPASYS to slow down progress on cores hosting less important tasks or ahead-of-schedule tasks if the progress of the task on the current core falls below the desired progress setpoint. We will quantify this in Section V-C (R3).

C. Implementation of Key Responsibilities

We implemented CAPASYS on the ZCU102 development board featuring an Xilinx UltraScale+ MPSoC. It consists of two CPU clusters: (1) the APU cluster contains four ARM Cortex-A53 CPUs at 1.3GHz serving as the traces; (2) the RPU cluster contains two ARM Cortex-R5 CPUs at 500MHz, used as the tracers. The target includes ARM CoreSight infrastructure modules that allow self-hosted tracing, in which the *Embedded Trace Macrocells* (ETMs) serve as TUs. Fig. 7 shows an overview. When ARM CoreSight is not available, several alternatives exist to perform multi-core progress tracking. (i) First, equivalent hardware tracing capabilities are offered by Intel Processor Trace, RISC-V E-Trace and N-Trace, MIPS PDtrace, etc. (ii) Second, hardware breakpoints can be used to achieve monitoring [20]. (iii) Finally, software-based approaches have been explored [17], [19].

(R1) Milestone Detection and TU (Re)configuration. In Chen et al. [22], vertex hits are detected by dynamically reprogramming ETM filters, which frequently toggles ETM state and emits extraneous metadata packets. We improved this by utilizing ETM events instead: the PC reaching a user-indicated virtual address can be programmed as an ETM event. When it occurs, only one 1-byte packet is emitted. The ETM filter can be set to a null range (e.g. $[0 \times 0, 0 \times 0]$) to suppress any unnecessary trace data. Once the CPU that originated the event is recovered (see R2 below), the TMG of the corresponding foreground task is queried for the set of possible next milestones. Next, the local ETM is temporarily disabled, reconfigured, and re-enabled.




Fig. 7: Green arrows are trace data stream paths. Gray shapes are trace data interconnects—CoreSight funnels and replicator. The first Cortex-R5 polls the first TMC for trace data. Red arrows are hardware signals for just-in-time flush. Blue arrows represent re-configuration and halt/resume operations.

(R2) Multi-core Trace Data Processing. The ETM only traces the activity on the local core. It does so by producing a stream of trace packets for downstream consumers. When multiple ETMs are active in a multicore platform, trace packets from different ETMs are multiplexed into 16-byte frames according to the Arm CoreSight formatter protocol. The consumer must demux the stream to recover the packets of individual cores. We use the first Cortex-R5 to demux. The recovered packets are written to local memory that is accessible by both Cortex-R5s. The second Cortex-R5 parses the packets which include the ETM event just occurred, indicating the virtual address that PC just reaches (see R1). With this information, tracer can identify the vertex hit in TMG, thus tracking the execution path.

Polling Trace Data. In Chen et al. [22], the trace data go through three trace memory controllers (TMCs) to reach the consumer. We discovered that the trace data can also be acquired by using software to poll any of the TMCs in software FIFO mode. Thus, only one TMC is involved in this implementation. Fig. 7 shows the said two paths.

Just-in-Time Flush. Multiple CoreSight components handle trace routing and must be flushed promptly for real-time delivery. However, frequent flushes can overload the channel with dummy data [22]. To avoid this, CAPASYS leverages ETM event outputs to trigger flushes only on vertex hits, using external pins routed to the TMC (red arrow in Fig. 7).

(R3) Progress-driven CPU Management. The aforementioned self-halting and external blocking rely on CAPASYS’ low-level capability to halt/resume cores, which is enabled by the Cross-Trigger Interface (CTI). The per-core CTI can be used to instruct each core to enter/leave debug state, a technique used in [4], [5], [31] that prevents cache pollution on the traces and that can be kept transparent to the OS. Below are the detailed behaviors of the regulation. CAPASYS allows one to specify a per-task **degradation factor** $\alpha \in (0, 1)$. The **progress setpoint** is then computed as $\alpha \tilde{P}(t)$. Upon a vertex hit, CAPASYS recomputes the current slack as $\Delta(i) = P(v_i) - \alpha \tilde{P}(t)$. A positive slack indicates a task’s progress ahead of the setpoint, while a negative slack indicates that it is falling behind. With this information and the ability to halt/resume CPU activity, CAPASYS implements the following two regulation mechanisms.

Self-halting Regulation. When a task accumulates enough positive slack, the core hosting it will be temporarily halted so that tasks on other cores can progress faster due to reduced contention. CAPASYS accepts a parametric **safety margin** $\beta \in (0, 1)$ capturing the required slack before self-halting.

Upon a vertex hit, if $\Delta > \beta$, CAPASYS will halt the core hosting the task for $\Delta/\tilde{P}(t)$ time units. This brings the effective progress back to the setpoint. Fig. 6 provides a visualization. Additionally, self-halting will not be enacted, if the task is the only running task system-wide.

External Blocking. CAPASYS accepts a configurable *ranking* of CPUs, encoding how stringent the timing constraints are on each CPU in relation to the other ones. This can be used to describe differences in CPU-wide time assurance levels or the presence of non-real-time tasks on some CPUs. This information is then considered by CAPASYS to halt lower-ranking CPUs if tasks on higher-ranking CPUs are not making sufficient progress. Specifically, upon the i^{th} vertex-hit for a task/CPU with higher ranking⁶, if $\Delta(i) < 0$, CAPASYS will halt other lower-ranking cores. Next, when the $(i + 1)^{\text{th}}$ hit occurs, three possible outcomes exist: (a) if $\Delta(i + 1) > 0$, all the cores halted by the considered task are resumed; (b) if $0 > \Delta(i + 1) > \Delta(i)$ (negative slack shrinking) no action is taken, awaiting future vertex hits; (c) if $\Delta(i + 1) \leq \Delta(i)$ (negative slack growing) an additional eligible core to halt is chosen until outcome (a) or (b) occurs⁷.

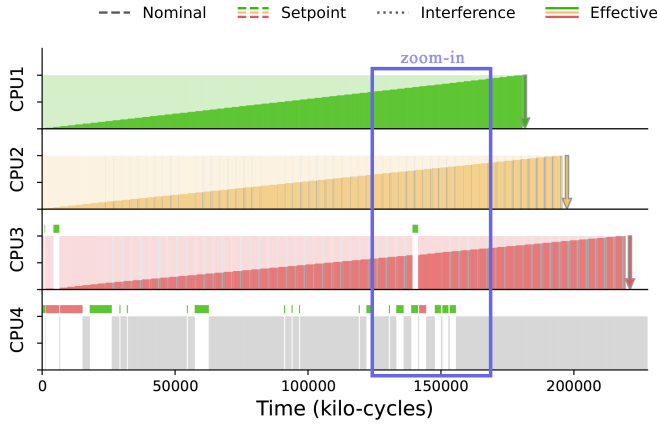
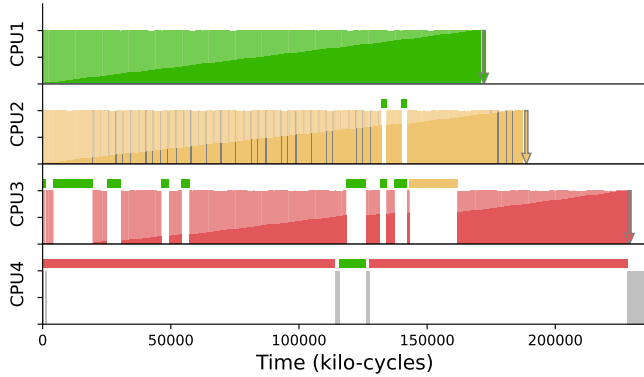
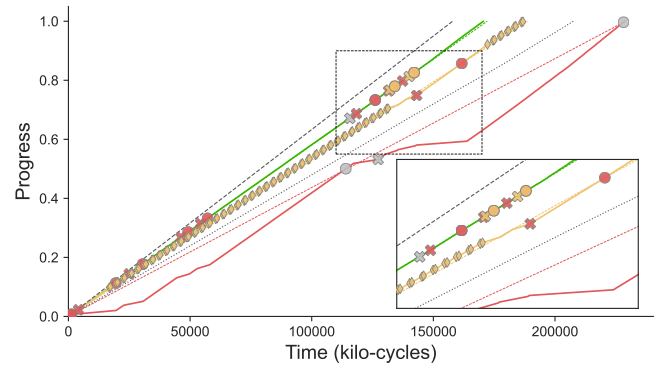
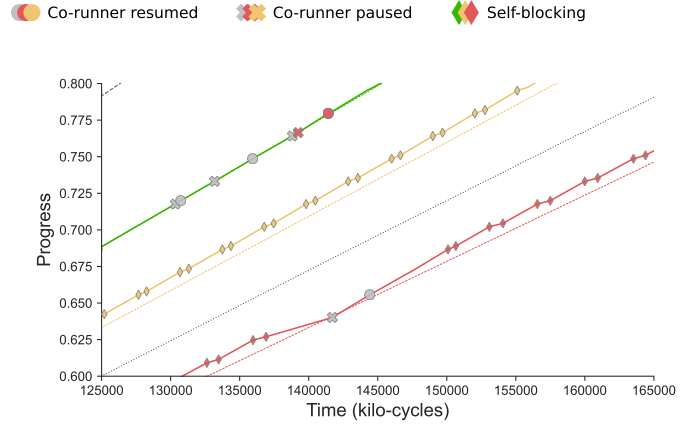
State Change due to Halt/resume. Although halting and resuming are conducted via the debug interface, not altering the core state during the freeze, some shared resources such as the last-level cache (LLC) can still be affected by contenders. When a task is resumed after a period of halting, it might need to warm up the LLC, incurring additional overhead. This effect is negligible in our setup; however, if it becomes significant, the overhead should be subtracted from the slack.

Watchdog Enforcement. Since CAPASYS does not track progress between vertex hits, long gaps may jeopardize the system. To prevent this, CAPASYS can optionally program a per-task watchdog to trigger if the next milestone is not reached within the expected time.

Safety Implication. The safety implication of the ranking configuration is that only tasks running on the highest-ranked core are guaranteed to meet their deadlines, as they can run in isolation when necessary. In contrast, cores with lower rankings receive resources on a best-effort basis. Therefore, the current regulation is more suitable for scheduling soft/firm real-time tasks. Section VI-D further explores how progress information can be leveraged to improve scheduling by combining the spare utilization of some existing scheduling algorithm.

⁶External blocking favors CPUs of higher ranking. Thus, a fairness metric is provided in Appendix V.

⁷A pseudo-code is provided in Appendix IV.

Fig. 8: Three real-time *disparity* tasks with degradation factors (0.87, 0.8, 0.71).Fig. 9: Three real-time *disparity* tasks with degradation factors (0.92, 0.83, 0.69).

Hardware Consideration. Compared with other software-based progress tracking methods [17], [18] in which a core hosts both user tasks and regulation logic, CAPASYS uses additional core(s) to host progress tracking and regulation logic. This is a trade-off between hardware requirements and timing overhead. If the target platform does not have spare (co)processors, the regulation logic of CAPASYS can be hosted on one of the application cores as well. In this case, the time overhead should be accounted for by discounting the corresponding amount of progress at each milestone hit.

D. CAPASYS Evaluation

We evaluate CAPASYS's ability to support progress-driven multi-core management without relying on a scheduler. One task is activated per core in sync, using a fixed ranking (CPU1 > CPU2 > CPU3 > CPU4). We extract the TMG and instantiate the CAPA model using realistic benchmarks from SD-VBS [32] and TACLeBench [33] as real-time workloads on CPUs 1–3. A memory-intensive best-effort task from IsolBench [34] runs on CPU4. CAPASYS is configured with degradation factors $(\alpha_0, \alpha_1, \alpha_2)$ and a fixed safety margin $\beta = 0.01$ for all real-time tasks. Though a smaller value of β will enact more frequent regulation, we find, experimentally, the system performance depends considerably more on the granularity of TMG than the value of β .

Benchmark Selection and Characteristics. *disparity*, *tracking*, *svm*, and *mser* were selected from SD-VBS. From this

suite, *texture_synthesis*, *stitch* and *localization* were excluded because largely CPU-bound, while *multi_ncut* was excluded because it does not run to completion on the selected platform. From the TACLeBench suite, we selected the most memory-intensive benchmark, namely *mpeg2*. All the selected benchmarks, except *mpeg2*, are considerably impacted by interference effects. Additionally, *mser* is particularly representative to illustrate the limitation of CAPASYS when the TMG has low granularity.

Overhead and Latency. The overhead of CAPASYS is computed by comparing the runtime of all the considered applications with and without active progress tracking. We observe that the overhead of CAPASYS on task executions is upper bounded by 0.4%. To measure the latency of milestone-hit detection, we use a synthetic task that self-reports the absolute timestamps of milestones. We then compute how many clock cycles later CAPASYS can detect the same event. In our measurements, the latency is upper-bounded by $16.4\mu\text{s}$ with $4.5\mu\text{s}$ on average.

Multi-core Management. Table II presents the TMG properties and runtime information of the considered benchmarks. We further present the plots for the most representative experiments⁸. In our first experiment, depicted in Fig. 8, three *disparity(s)* benchmarks from SD-VBS are deployed as the real-

⁸Due to space constraints, we omit the plots of benchmarks such as *tracking* and *svm* since they are very similar to what is depicted in Fig. 9 and Fig. 10.




Fig. 10: Three real-time *mpeg2* tasks; all of them have sufficient positive slack, thus enacting self-halting.




Fig. 11: Three real-time *msers* tasks; due to the low-granularity TMG, management is not effective. Additional horizontal dotted lines are added for each milestone hit to visualize the low-granularity issue. It can be seen that many milestone hits are concentrated above $p = 0.5$

benchmarks	vertices	arcs	SCCs	vertex-hits	min/max/avg T_{move} (μs)	nominal WCET (μs)
<i>disparity</i>	11	12	1	137	1270 / 20828 / 11611	1579172
<i>mpeg2</i>	3	3	1	18	6 / 5546 / 5127	87173
<i>msers</i>	10	11	1	10	185 / 207880 / 46447	418025
<i>svm</i>	6	9	1	45	195 / 114485 / 12791	562839
<i>tracking</i>	11	12	1	17	110 / 104201 / 48219	771506

TABLE II: Properties of TMGs for each benchmark. The vertex-hits column is measured at runtime. All other columns are static data obtained during the profiling phase.

time tasks on CPU 1–3. Each task is assigned a color profile in agreement with the CPU ranking: green (CPU1, $\alpha_1 = 0.87$), yellow (CPU2, $\alpha_2 = 0.8$), red (CPU3, $\alpha_3 = 0.71$), and gray (memory-bomb, CPU4). The degradation factor values here are chosen to demonstrate scenarios in which CAPASYS exhibits interesting behaviors. Because the evaluation in this section focuses on inter-core contention effects and considers only one active job per core, schedulability considerations can not be meaningfully discussed yet. In Section VI-D we will explore the advantages of leveraging progress information when multiple tasks per core are active, thus enabling progress-aware, contention-aware multi-core scheduling.

Fig. 8-left describes the scheduling diagram with progress annotated: The gray vertical bars interposed on the green, yellow, and red indicate the duration of the self-blocking. The gap indicates the task is halted due to another task on a higher-ranking CPU. The task on the higher-ranking CPU triggering the halt is shown as color strips on top of the gap.

Fig. 8-right focuses on the per-task progress. The plot is zoomed in to show the activities when execution time is between 125M to 185M cycles. Important progress lines, as in Fig. 6, are shown: nominal, setpoint, and effective progress. An *interference* progress line is also added for reference, which is obtained by running the same setup but where CAPASYS takes no management actions. The first experiment has the interference line ends 210M cycles, which is significantly larger than its contention-free WCET (150M cycles).

A \blacklozenge marker on the effective progress indicates that CAPASYS enacts self-blocking. A color-coded \times marker indicates that external blocking is performed. For example, a gray \times on the green effective progress line indicates that CAPASYS halts the gray task because the green task has a negative slack. A \bullet marker indicates that halted cores are resumed. Thus \times and \bullet typically come in pairs. By setting distinct degradation factors, the first two *disparity* tasks have setpoints above the interference line, the third below. The zoom-in shows various CAPASYS actions to keep the progress of each task to the vicinity of their setpoints.

Behavior under Unsatisfiable Constraints. In this experiment (Fig. 9), we use the same setup as Fig. 8 but with adjusted degradation factors: $\alpha_1 = 0.92$, $\alpha_2 = 0.83$, $\alpha_3 = 0.69$. This places the system slightly outside the feasibility envelope. In this case, the non-real-time task (gray) barely runs. Moreover, the red task finishes with a slight negative slack (deadline miss) and it is halted a number of times to allow higher rank tasks to re-gain positive slacks. CAPASYS effectively maintains schedulability on the top two ranked CPUs, with both the green and yellow tasks completing with positive slack.

Self-blocking for Positive Slack. We evaluate CAPASYS on applications with minimal memory contention using the *mpeg2* benchmark from TACLeBench (Fig. 10, with $\alpha_1 = 0.91$, $\alpha_2 = 0.83$, $\alpha_3 = 0.71$). Here, the *Interference* line closely tracks the *Nominal*, allowing real-time tasks to accumulate positive slack. CAPASYS enacts self-blocking to align progress with

setpoints. While we currently use core halting, other strategies—e.g., executing lower-priority tasks—are also viable. The frequency of monitoring on *mpeg2* is much higher (~ 300 kilo-cycles per invocation), compared to those of *disparity* and *mser* (~ 1333 and ~ 3000 kilo-cycles per invocation). Due to the low-overhead nature of CAPASYS, we do not observe any performance degradation experimentally from high monitoring frequency. However, if the excessive accesses to debug infrastructure is a concern in practice, then some invocations can be skipped to proactively constraint the access. We leave this as future research.

Limitations from Low TMG Granularity. The final experiment highlights how low-granularity TMGs can limit CAPASYS’s ability to make fine-grained decisions. Using the *mser* benchmark from SD-VBS for all tasks ($\alpha_1 = 0.91$, $\alpha_2 = 0.83$, $\alpha_3 = 0.71$), where most milestones occur before progress 0.6, Fig. 11 shows that coarse milestone spacing hinders maintaining progress above the setpoint. Improving TMG granularity is part of our future work.

VI. CAPASIM DESIGN: CAPA-BASED MULTI-CORE SCHEDULING SIMULATOR

Up to and including Section V, all evaluations were conducted on real hardware without assuming a specific scheduler. This section explores the potential of integrating CAPASYS’s management capabilities into a multi-core scheduler—a direction that opens a broad design space for future research. To enable rapid prototyping and hypothesis testing, we also introduce and implement CAPA-based Multi-core Scheduling Simulator (CAPASIM). With the help of CAPASIM, the performance of CAPASYS’s regulation is further evaluated in this section, and a comparative study of shared resource partitioning techniques is also presented.

A. Motivation for CAPASIM

While many scheduling simulators exist [35], [36], these treat jobs as indivisible atoms and do not precisely model shared resource contention effects. Conversely, the proposed CAPASIM models shared resource contention effect. Admittedly, shared resource contention modeling is challenging. The type and topology of the shared resources produce different contention patterns [37]. Thus, our CAPASIM provides the contention model as an interface so that different models can be explored. Nevertheless we propose a baseline statistical contention model. The CAPASIM is open-source to encourage community reuse. The design of CAPASIM is purposefully de-coupled from CAPATASK and CAPASYS, and is a self-contained simulator, so that CAPASIM can be used in a broader scope of research. We will first describe how CAPASIM functions, and then apply CAPATASK and the regulation techniques of CAPASYS in the simulation setup.

B. Simulation Model

Due to space constraints, the complete simulation model can be found in Appendix VI. Here, we only present the key

parameters necessary to follow the rest of the paper. A task in CAPASIM is represented as a sequence of tuples:

$$\{(a_n, b_n) \mid n \in \mathbb{N}^+, a_n \in \mathbb{N}^+, b_n \in [0, 1], n \leq N\} \quad (23)$$

In the n^{th} segment, a_n denotes the number of progress units (PUs) required to complete the segment. A task with N segments requires a total of $\sum_{n=1}^N a_n$ PUs to complete. While b_n specifies the resource consumption factor of the segment, indicating the fraction of PUs in a_n that involve usage of a shared resource. Each simulation cycle, a task will typically advance less than 1 PU if $b > 0$, and the advancement becomes smaller as contention exacerbates. We estimate the number of cycles required to complete a PU that involves shared resource access by leveraging the measurements from Saeed et al. [38], which report memory access latencies under contention-free and maximum contention scenarios. These are denoted as $\mu_{\text{solo}} = 50$ and $\mu_{\text{max}} = 107$ cycles, respectively. CAPASIM utilizes this information to statistically determine the advancement a task makes per simulation cycle. Please refer to Appendix VI for details about how μ_{solo} and μ_{max} are used to compute contention-aware performance degradation effects.

C. Augmenting P-RM and P-EDF with CAPASYS

As a first look at the benefits of integrating the progress-aware management primitives implemented by CAPASYS into multi-core schedulers, we simulate an integration with Partitioned Rate-Monotonic (P-RM) [39] and Partitioned Earliest-Deadline First (P-EDF) [40] scheduling. We refer to the P-RM (resp., P-EDF) scheduler augmented with CAPASYS as CP-RM (resp., CP-EDF). To do so, suitable degradation factors (α_i) have to be set for each task. For the remaining section, any mentioning of utilization corresponds to the contention-free WCET⁹. Spare utilization U' is used as a heuristic to set α_i . Recall that a sufficient condition for a fully preemptive periodic taskset to be schedulable under RM on single-core is $U \leq m(2^{\frac{1}{m}} - 1)$, in which U is the total utilization, and m is the number of task(s). We define per-core spare utilization $U' = m(2^{\frac{1}{m}} - 1) - U$, which can be uniformly distributed to each task. For EDF on a single core, the taskset is schedulable if $U \leq 1$. Thus for each core under P-EDF, the spare utilization is $U' = 1 - U$. The safe execution time of each task can, therefore, be safely increased to

$$C'_i = T_i(U_i + U'/m) \quad (24)$$

With C'_i being the setpoint, it can be calculated $\alpha_i = C_i/C'_i$.

Experimentally, the following two optimizations achieve better performance: (1) Use only self-halting, not external-blocking. Using both improves timeliness on high time-assurance cores, but increases deadline misses on low-assurance cores. While this yields better timeliness on high time-assurance cores than plain P-RM/EDF, the overall schedulability under CP-RM/EDF worsens. In contrast, applying only self-halting leads to better overall performance than

⁹Using contention-free utilization is the choice of the proposed algorithm. Methods that absorb the contention/interference/noisy effect into the utilization exist, such as mixed criticality systems [41].




Fig. 12: Total utilization is increased in increments of 1%. For each utilization value, 100 tasksets are generated and tested against each scheduler. This process is repeated for 100 rounds. The envelope represents the maximum and minimum percentages of schedulable tasksets observed across the 100 rounds. From left to right, the utilization balances y_{ub} are 0.33, 0.43, 0.75, and the consumption factors b_n are sampled from Gaussian distribution with mean 0.3, 0.2, 0.5, and standard deviation 0.025. The μ_{max} is inflated by $6\times$, $3\times$, $3\times$.¹⁰




Fig. 13: For each Gaussian mean and utilization balance pair, all possible total utilization (from 0% to 400% with 1% incremental) are run with 100 independent tasksets for each. Let y in the heat maps on left be the total number of schedulable tasksets for P-RM/EDF, and on right for P-RM/EDF with resource partitioning. y' for CP-RM/EDF. The score in the entry is $(y' - y)/N$ in percentage, where $N = 40000$ is the total number of tasksets tested to produce one entry (100 tasksets for each total utilization 100×400). At $y = 0.25$ on left, all entries are zero, because the optimization from Eq. 25. Six colored boxed entries correspond to the scenarios tested as in Fig. 12.

P-RM/EDF. Moreover, since P-RM/EDF does not distinguish assurance levels, comparing it with self-halting alone—which also lacks such differentiation—is more appropriate. (2) Let \bar{U} be the average per-core utilization. Only cores with

$$U_i < \bar{U}/z \quad (25)$$

enact self-halting, where $z = 1.1$ is a cutoff value. The intuition is that although self-halting causes processor time to be lost, cores with below-average utilization can yield shared resources via self-halting.

D. CP-RM and CP-EDF Simulation Experiments

The simulated system consists of 4 cores, each running 4 tasks, with the following setup: the utilization of the first core is no less than that of the remaining three cores; the three remaining cores have identical utilization; The utilization

balance factor y_{ub} indicates the percentage of total utilization assigned to the first core. Given a total utilization U_{tot} , the first core receives $y_{ub}U_{tot}$, while each of the other three cores receives $(1 - y_{ub})U_{tot}/3$. Each consumption factor b_n is sampled from a Gaussian distribution prior to simulation; The total PUs of a task is sampled from uniform distribution with mean equal to 300. Each task is instrumented with 100 uniformly spaced vertices across its total PUs. To randomize the utilization of each task, UUniFastDiscard [42] is used to generate task utilization given the total utilization received by a core. Implicit deadlines are used. All first jobs release at $t = 0$. For a given taskset, the simulation runs for $3\times$ of the longest period among all tasks. We compare P-RM and P-EDF scheduling against their CAPA variants, i.e., CP-RM and CP-EDF. Fig. 12 shows representative schedulability curves. Though Fig. 12 appears similar to conventional schedulability curves, the nature of the simulation implies that the WCETs under maximum contention is unknown. We will use *feasible* to describe the absence of deadline misses within a finite simulation duration of a taskset. Thus Fig. 12 represents

¹⁰The advancement model of CAPASIM causes the slowdown to be at most twofold. However, experiments show the contention-induced slowdown can be much more severe [20]. Thus we inflate μ_{max} to offset the underestimation.

feasibility curves instead of schedulability curves.

In most cases, CP-RM/EDF outperform their vanilla counterparts. The rightmost diagram shows CP-RM under-performs when the utilization is highly unbalanced. We conduct more extensive experiments by swiping a range of parameters. The results are visualized using heat maps on the left in Fig. 13, in which each cell is an aggregation of a feasibility curve like Fig. 12. In heat maps, the y-axis represents the utilization balance factor y_{ub} . The mean of the b_n factors is varied in the set $\{0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$; μ_{max} is multiplied by $\{1\times, 3\times, 6\times\}$. Green/red-shifted entries indicate more/less tasksets are feasible under CAPA variant.

CP-EDF consistently outperforms P-EDF across most test cases. However, results from (C)P-RM—and some extreme (C)P-EDF scenarios—exhibit counterintuitive behavior: under high contention in unbalanced systems, P-RM can outperform CP-RM. This occurs because brief periods of intense contention can yield more total progress than sustained moderate contention over the same window. While even simple integrations of CAPA into basic multi-core schedulers yield significant gains, a systematic investigation is needed to fully harness progress information in multi-core scheduling.

We compare our progress-based scheduling with shared resource partitioning, including way/set-based cache partitioning [43] and memory bandwidth regulation (e.g., MemGuard, MemPol, MemCoRe [3]–[5]). These methods assign cores specific cache slices or bandwidth quotas, suspending execution upon quota violations. We choose these comparisons for two reasons: (1) our self-halting mechanism resembles how these methods conduct regulation, i.e. via core throttling; and (2) other progress-estimation methods either do not include schedulability/feasibility [18]–[20], [22] or are limited to single-core systems with few milestones [17].

We must stress that our implementation of shared-resource partitioning in the simulation does not faithfully replicate real hardware behavior, because the simulation does not explicitly model hardware resources. Instead, we focus on modeling the effect that access to a slice of shared memory resources has on applications' runtime. This is intended as an illustrative demonstration showing that indiscriminately slowing down the memory phases of applications performs worse than purposefully managing their behavior by leveraging progress information. The heat maps on the right of Fig. 13 show the result. CP-RM/EDF continue to outperform P-RM/EDF with resource partitioning. We model shared resource partitioning as follows. Recall that at each simulation cycle, a running task has a shared resource-consumption factor b that parametrizes the density of memory operations in the current phase. As such, if for g_{window} consecutive simulation cycles, the values of b of a task at each cycle are all larger than $b_{threshold}$, a memory stall is introduced in the form of a 1-simulation-cycle CPU stall. This has the effect of slowing down only the memory-intensive phases of applications, which is consistent with the inherent slowdown introduced by shared resource partitioning techniques. Otherwise, the underlying scheduler remains P-RM/P-EDF. We set $b_{threshold}$ equal to the mean of the Gaussian where b_n is sampled, and g_{window} equal to 65. Although not shown in the heatmap, we also experiment with g_{window} equal

to 16, 33, 120, 240, corresponding to the slowdown of memory phase by 6.25%, 3%, 0.83%, 0.41%. Their performance is very similar to what we observed with $g_{window} = 65$.¹¹

VII. CONCLUSION

This paper presents the CAPA framework—comprising the CAPATASK, CAPASYS, and CAPASIM—to advance progress-aware real-time systems along four dimensions. First, the CAPATASK enables runtime tracking of task progress with complex control flows. Second, CAPASYS extends state-of-the-art TPA implementations to support concurrent multi-core execution with minimal overhead. Third, CAPASIM models task progress and shared resource contention. Fourth, we explore using live progress information to drive system-wide scheduling decisions under CP-RM and CP-EDF, showing promising improvements in feasibility under contention.

While CAPA opens new directions in progress-driven scheduling, further research is needed to fully realize its potential. Future work will explore global scheduling policies and leverage CAPA for enhanced task observability and profile-driven resource management under contention.

REFERENCES

- [1] C. Maiza, H. Rihani, J. Rivas, J. Goossens, S. Altmeyer, and R. Davis, "A Survey of Timing Verification Techniques for Multi-Core Real-Time Systems," *ACM Comput. Surv.*, vol. 52, no. 3, Jun. 2019.
- [2] T. Lugo, S. Lozano, J. Fernández, and J. Carretero, "A survey of techniques for reducing interference in real-time applications on multicore platforms," *IEEE Access*, vol. 10, pp. 21 853–21 882, 2022.
- [3] H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha, "Memguard: Memory bandwidth reservation system for efficient performance isolation in multi-core platforms," in *2013 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2013.
- [4] A. Zuepke, A. Bastoni, W. Chen, M. Caccamo, and R. Mancuso, "Mem-pol: Policing core memory bandwidth from outside of the cores," in *2023 IEEE 29th Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2023, pp. 235–248.
- [5] I. Izhbirdeev, D. Hoornaert, W. Chen, A. Zuepke, Y. Hammad, M. Caccamo, and R. Mancuso, "Coherence-Aided Memory Bandwidth Regulation," in *2024 IEEE Real-Time Systems Symposium (RTSS)*. Los Alamitos, CA, USA: IEEE Computer Society, Dec. 2024.
- [6] D. Ottaviano, F. Ciraolo, R. Mancuso, and M. Cinque, "The Omnivisor: A Real-Time Static Partitioning Hypervisor Extension for Heterogeneous Core Virtualization over MPSoCs," in *36th Euromicro Conference on Real-Time Systems (ECRTS 2024)*, vol. 298, 2024, pp. 7:1–7:27.
- [7] J. Marshall, R. Gifford, G. Bloom, G. Parmer, and R. Simha, "Precise cache profiling for studying radiation effects," *ACM Trans. Embed. Comput. Syst.*, vol. 20, no. 3, Mar. 2021.
- [8] D. Tarapore, S. Roozkhosh, S. Brzozowski, and R. Mancuso, "Observing the invisible: Live cache inspection for high-performance embedded systems," *IEEE Transactions on Computers*, mar 2022.
- [9] S. Roozkhosh, D. Hoornaert, and R. Mancuso, "Caesar: Coherence-aided elective and seamless alternative routing via on-chip fpga," in *2022 IEEE Real-Time Systems Symposium (RTSS)*, 2022, pp. 356–369.
- [10] P. Sohal, R. Tabish, U. Drepper, and R. Mancuso, "E-warp: A system-wide framework for memory bandwidth profiling and management," in *2020 IEEE Real-Time Systems Symposium (RTSS)*, 2020.
- [11] S. Baruah, A. Burns, and R. Davis, "Response-time analysis for mixed criticality systems," in *2011 IEEE 32nd Real-Time Systems Symposium*, 2011.
- [12] S. Baruah, V. Bonifaci, G. D'angelo, H. Li, A. Marchetti-Spaccamela, S. Van Der Ster, and L. Stougie, "Preemptive uniprocessor scheduling of mixed-criticality sporadic task systems," *J. ACM*, May 2015.
- [13] R. Gifford, N. Gandhi, L. T. X. Phan, and A. Haeberlen, "DNA: Dynamic resource allocation for soft real-time multicore systems," in *27th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS '21)*, May 2021.

¹¹The plots can be found in Appendix VII.


- [14] D. Lo, M. Ismail, T. Chen, and G. E. Suh, "Slack-aware opportunistic monitoring for real-time systems," in *2014 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2014.
- [15] T. Degioanni and I. Puaut, "StAMP: Static Analysis of Memory Access Profiles for Real-Time Tasks," in *20th International Workshop on Worst-Case Execution Time Analysis (WCET 2022)*, ser. Open Access Series in Informatics (OASISs), vol. 103, Dagstuhl, Germany, 2022, pp. 1:1–1:13.
- [16] R. Pellizzoni, E. Betti, S. Bak, G. Yao, J. Criswell, M. Caccamo, and R. Kegley, "A predictable execution model for cots-based embedded systems," in *2011 17th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2011, pp. 269–279.
- [17] S. Sinha, R. West, and A. Golchin, "PAStime: Progress-Aware Scheduling for Time-Critical Computing," in *32nd Euromicro Conference on Real-Time Systems (ECRTS 2020)*, vol. 165, 2020, pp. 3:1–3:24.
- [18] A. Kritikakou, C. Pagetti, O. Baldellon, M. Roy, and C. Rochange, "Run-time control to increase task parallelism in mixed-critical systems," in *2014 26th Euromicro Conference on Real-Time Systems*, 2014.
- [19] 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*, ser. RTNS '14, New York, NY, USA, 2014, p. 139–148.
- [20] D. Oliveira, W. Chen, S. Pinto, and R. Mancuso, "Shared Resource Contention in MCUs: A Reality Check and the Quest for Timeliness," in *36th Euromicro Conference on Real-Time Systems (ECRTS 2024)*, vol. 298, 2024, pp. 5:1–5:25.
- [21] A. Kritikakou, T. Marty, and M. Roy, "Dynascore: Dynamic software controller to increase resource utilization in mixed-critical systems," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 23, no. 2, oct 2017.
- [22] W. Chen, I. Izhbirdeev, D. Hoornaert, S. Roozkhosh, P. Carpanedo, S. Sharma, and R. Mancuso, "Low-Overhead Online Assessment of Timely Progress as a System Commodity," in *35th Euromicro Conference on Real-Time Systems (ECRTS 2023)*, 2023.
- [23] J. M. Aceituno, A. Guasque, P. Balbastre, J. Simó, and A. Crespo, "Hardware resources contention-aware scheduling of hard real-time multiprocessor systems," *Journal of Systems Architecture*, 2021.
- [24] E. Bini, "Adaptive fair scheduler: Fairness in presence of disturbances," in *Proceedings of the 24th International Conference on Real-Time Networks and Systems*, ser. RTNS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 129–138.
- [25] H. Kim, D. de Niz, B. Andersson, M. Klein, and J. Lehoczy, "Addressing multi-core timing interference using co-runner locking," in *2021 IEEE Real-Time Systems Symposium (RTSS)*, 2021, pp. 54–67.
- [26] L. Passig Horstmann, J. L. Conradi Hoffmann, and A. A. Fröhlich, "Monitoring the performance of multicore embedded systems without disrupting its timing requirements," *Design Automation for Embedded Systems*, vol. 27, no. 4, pp. 217–239, dec 2023.
- [27] J. Anantpur and R. Govindarajan, "Pro: Progress aware gpu warp scheduling algorithm," in *2015 IEEE International Parallel and Distributed Processing Symposium*, 2015, pp. 979–988.
- [28] S. Moulik and Y. Sharma, "Fresh: Fault-tolerant real-time scheduler for heterogeneous multiprocessor platforms," *Future Generation Computer Systems*, 2024.
- [29] S. Chakraborty, Y. Sharma, and S. Moulik, "Treatfet: Temperature-aware real-time task scheduling for finfet based multicores," *ACM Transactions on Embedded Computing Systems*, 05 2024.
- [30] Y. Shoshitaishvili, R. Wang, C. Salls, N. Stephens, M. Polino, A. Dutcher, J. Grosen, S. Feng, C. Hauser, C. Kruegel, and G. Vigna, "SoK: (State of) The Art of War: Offensive Techniques in Binary Analysis," in *IEEE Symposium on Security and Privacy*, 2016.
- [31] A. Zuepke, A. Bastoni, W. Chen, M. Caccamo, and R. Mancuso, "Mempol: polling-based microsecond-scale per-core memory bandwidth regulation," *Real-Time Systems*, pp. 1–44, 2024.
- [32] S. K. Venkata, I. Ahn, D. Jeon, A. Gupta, C. Louie, S. Garcia, S. Belongie, and M. B. Taylor, "SD-VBS: The san diego vision benchmark suite," in *2009 IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 55–64.
- [33] H. Falk, S. Altmeyer, P. Hellinckx, B. Lisper, W. Puffitsch, C. Rochange, M. Schoeberl, R. B. Sørensen, P. Wägemann, and S. Wegener, "TACLeBench: A benchmark collection to support worst-case execution time research," vol. 55, Dagstuhl, Germany, 2016.
- [34] P. K. Valsan, H. Yun, and F. Farshchi, "Taming non-blocking caches to improve isolation in multicore real-time systems," in *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*.
- [35] M. Chérany, P.-E. Hladik, and A.-M. Déplanche, "Simsos: A simulation tool to evaluate real-time multiprocessor scheduling algorithms," in *Proc. of the 5th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems*, ser. WATERS, 2014.
- [36] F. Singhoff, J. Legrand, L. Nana, and L. Marcé, "Cheddar: a flexible real time scheduling framework," *Ada Lett.*, Nov. 2004.
- [37] S. Zhuravlev, S. Blagodurov, and A. Fedorova, "Addressing shared resource contention in multicore processors via scheduling," 2010.
- [38] A. Saeed, D. Hoornaert, D. Dasari, D. Ziegenbein, D. Mueller-Gritschneider, U. Schlichtmann, A. Gerstlauer, and R. Mancuso, "Memory Latency Distribution-Driven Regulation for Temporal Isolation in MPSoCs," in *35th Euromicro Conference on Real-Time Systems (ECRTS 2023)*, 2023.
- [39] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *J. ACM*, 1973.
- [40] J. Xu and D. Parnas, "Scheduling processors with release times, deadlines, precedence and exclusion relations," *IEEE Trans. Softw. Eng.*, vol. 16, no. 3, p. 360–369, Mar. 1990.
- [41] A. Burns and R. I. Davis, "A survey of research into mixed criticality systems," *ACM Comput. Surv.*, vol. 50, no. 6, Nov. 2017.
- [42] R. I. Davis and A. Burns, "Priority assignment for global fixed priority pre-emptive scheduling in multiprocessor real-time systems," in *2009 30th IEEE Real-Time Systems Symposium*, 2009.
- [43] M. Caccamo, M. Cesati, R. Pellizzoni, E. Betti, R. Dudko, and R. Mancuso, "Real-time cache management framework for multi-core architectures," in *Proceedings of the 2013 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2013.

AUTHOR BIOGRAPHIES

Weifan Chen is a Ph.D. candidate in the Computer Science Department at Boston University. Weifan's work focuses on the performance of real-time systems on modern system-on-a-chip architectures. He has received multiple awards, including one Outstanding Paper, one Best Paper, and one Best Presentation.



Ivan Izhbirdeev is a Ph.D. student in Computer Science Department at Boston University and a member of the Cyber-Physical Systems Lab, supervised by Prof. Renato Mancuso. Ivan's research focuses on memory management and regulation, real-time systems performance, and predictable computing on modern system-on-chip and FPGA-based platforms. He has received multiple awards, including two Outstanding Paper Awards and one Best Paper Award.



Renato Mancuso is an associate professor in the Computer Science Department at Boston University and the director of the BU Cyber-Physical Systems Lab. Renato received his Ph.D. from the Department of Computer Science at the University of Illinois at Urbana-Champaign in 2017. His work focuses on real-time performance management techniques tailored to modern high-performance platforms. He has received several awards, including multiple Best Paper, Outstanding Paper, Best Demo, and Best Presentation awards. Prof. Mancuso also received the

NSF CAREER Award and the Gerald and Deanne Gitner Family Award for Innovation in Teaching with Technology in 2023, and the ACM SIGBED Early Career Researcher Award in 2024. His research has been supported by federal agencies and industrial partners such as the NSF, Bosch, Red Hat, Cisco, and Xilinx. He is a senior member of the IEEE and the ACM SIGBED.

