A Hybrid Performance Analysis Technique for Distributed Real-Time Embedded Systems

JUNCHUL CHOI, Seoul National University
HYUNOK OH, Hanyang University
SOONHOI HA, Seoul National University

It remains a challenging problem to tightly estimate the worst case response time of an application in a distributed embedded system, especially when there are dependencies between tasks. We discovered that the state-of-the art techniques considering task dependencies either fail to obtain a conservative bound or produce a loose upper bound. We propose a novel conservative performance analysis, called hybrid performance analysis, combining the response time analysis technique and the scheduling time bound analysis technique to compute a tighter bound fast. Through extensive experiments with randomly generated graphs, superior performance of our proposed approach compared with previous methods is confirmed.

1. INTRODUCTION

For the design of embedded systems that support real-time applications, it is required to guarantee the satisfaction of real-time constraints. After applications are mapped to a candidate architecture, we check the feasibility of the architecture by estimating the performance. Fast estimation enables us to explore the wider design space of architecture selection and application mapping. More accurate estimation will reduce the system cost. The performance analysis problem addressed in this paper is to estimate the worst case response time (WCRT) of an application that is executed on a distributed embedded system. A good example can be found in an intelligent safety application in a car where there is a tight requirement on the worst case response time from the sensor input to the actuator output.

Despite a long history of research over two decades, it still remains a challenging problem to tightly estimate the WCRT of an application in a distributed embedded system based on a fixed priority scheduling policy. Since the response time of an application is affected by interference between applications as well as execution time

variation of tasks, all possible execution scenarios should be considered to obtain the exact WCRT. There are some approaches proposed, such as a model checking approach [Brekling et al. 2008] and an ILP-based approach [Kim et al. 2012], to find the accurate WCRT. However, they require exponential time complexity. The exact WCRT analysis problem is known to be NP-complete [Yen and Wolf 1998].

Analytical techniques have been extensively researched to obtain a tight upper bound of the WCRT with diverse assumptions on target architectures and applications. This paper assumes that an application is given as a task graph that represents data dependency between tasks and the execution time of a task may vary. It is assumed that each task has a fixed priority. In addition, we support arbitrary mixture of preemptive and non-preemptive processing elements in the system. To analyze the WCRT of an application, this paper proposes a performance analysis technique, called hybrid performance analysis (HPA), combining a scheduling time bound analysis and a response time analysis (RTA). The proposed technique is proven to be conservative and experimental results show that it provides a tighter bound of WCRT than the other state-of-the-art techniques.

The rest of this paper is organized as follows. In Section 2 we overview the related work and highlight the contributions of this work. In Section 3, the application model and the system model assumed in this paper are formally described. Section 4 reviews the Y&W method and proves that it is not conservative by showing some counter examples. The proposed technique and some optimization techniques are explained in Sections 5 and 6 respectively. We summarize the overall algorithm in Section 7. Experimental results are discussed in Section 8. Finally, we conclude this paper in Section 9.

2. RELATED WORK

Response time analysis (RTA) was first introduced for a single processor system based on preemptive scheduling of independent tasks that have fixed priorities, fixed execution times, and relative deadline constraints equal to their periods [Lehoczky et al. 1989]. Extensive research efforts [Lehoczky 1990, Audsley et al. 1993] have been performed to release the restricted assumptions. Pioneered by K. Tindell et al. [Tindell and Clark 1994], a group of researchers extended the schedulability analysis technique to distributed systems; for example, dynamic offset of tasks [Palencia and Harbour 1998], communication scheduling [Tindel et all. 1995], partitioned scheduling with shared resources [Schliecker and Ernst 2010], and earliest deadline first (EDF) scheduling [Pellizzoni and Lipari 2007]. There exist some researches that consider precedence constraints between tasks, by assigning the offset and deadline of each task conservatively considering every possible execution ordering between tasks [Palencia and Harbour 1998, Tindel et all. 1995, Pellizzoni and Lipari 2007], which usually incur significant overhead of overestimation and computation complexity. On the other hand, we handle the dependent tasks directly, assuming that a task is released immediately after all predecessors complete.

To the best of our knowledge, the state of the art RTA method for dependent tasks was proposed by Yen and Wolf [Yen and Wolf 1998], denoted as the Y&W method hereafter. It considers data dependencies between tasks and variable execution times of tasks directly in the response time analysis but supports only preemptive systems. Several extensions have been made to the Y&W method, with considering communication [Yen and Wolf 1995] and control dependencies [Pol et al. 2000]. Unfortunately, the conservativeness has not been proven for the Y&W method. This paper discovers that the Y&W method is actually not conservative by showing counter-examples for which the Y&W method produces a shorter response time than the worst case response time in Section 4.

Non-preemptive processing elements are supported in the MAST suite [Harbour 2001] that includes several schedulability analysis techniques. But they support only chain-structured graphs where a task has a single input and/or a single output port.

There is a compositional approach distinguished from the holistic RTA-based approaches. SymTA/S [Henia et al. 2005] which is a well-known compositional analysis, performs the analysis in a modular manner. It analyzes the performance for each processing component and abstracts its result as an event stream at the component boundary. While the compositional approach achieves scalability, it sacrifices estimation accuracy by ignoring the release time constraint coming from data dependencies between tasks running on different processing elements.

Recently, a holistic WCRT analysis approach, called scheduling time bound analysis (STBA), has been proposed [Kim et al. 2013]. It computes the conservative time bound for each task within which the task will be scheduled, considering all possible scheduling patterns. In the STBA approach, however, the task graphs should be expanded up to the least common multiple (LCM) of their periods, which limits the scalability of the technique. While the proposed technique adopts the basic time bound idea of the STBA method to analyze how data dependencies affect the release times of tasks in the same application, it considers inter-application interference analytically, based on the response time analysis.

3. PROBLEM DEFINITION

(a) Task graphs

Task Graph	$\mathcal{T}^p, \mathcal{T}^j$
$\mathcal{T}_0$	100, 0
$\mathcal{T}_1$	100, 0

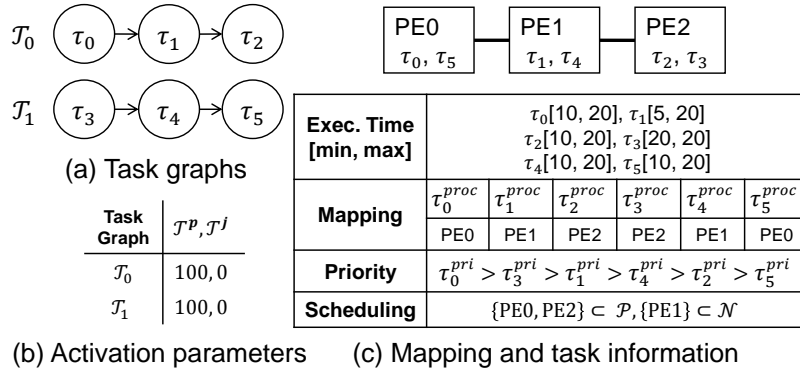(b) Activation parameters

(c) Mapping and task information

Fig. 1.   (a) Example task graphs, (b) task graph information, (c) task mapping and task information

We formally describe the application model and the system model assumed in this paper. An input application, $\mathcal{T}_i$, is represented as an acyclic task graph as illustrated in Fig. 1 (a). In a task graph, $\mathcal{T} = \{\mathcal{V}, \mathcal{E}\}$, $\mathcal{V}$ represents a set of tasks and $\mathcal{E} = \{(\tau_1, \tau_2)|\tau_1, \tau_2 \in \mathcal{V}\}$ a set of edges to represent execution dependencies between tasks. If a task has more than one input edge, it is released after all predecessor tasks are completed. An application $\mathcal{T}$ can be initiated periodically or sporadically, characterized by a tuple $(\mathcal{T}^p, \mathcal{T}^j)$ where $\mathcal{T}^p$ and $\mathcal{T}^j$ represent the period and the maximum jitter, respectively. For sporadic activation, $\mathcal{T}^p$ denotes the minimum initiation interval. Task graph $\mathcal{T}$ is given a relative deadline $\mathcal{T}^d$ to meet once activated. We assume that $\mathcal{T}^d$ is not greater than $\mathcal{T}^p$ in this paper. The task graph that task τ_i belongs to is denoted by $\mathcal{T}_{\tau_i}$.

A system consists of a set of processing elements (PEs) as shown in Fig. 1 (c). A task is a basic mapping unit onto a processing element. We assume that task mapping is given and fixed. The processing element that the task τ is mapped to is de-

noted by τ^{proc}. For each task τ, the varying execution time is represented as a tuple $[\tau^{BCET}, \tau^{WCET}]$ indicating the lower and the upper bound on the mapped PE. Note that a communication network can be modeled as a separate PE. For instance, the PE graph of Fig. 1 (c) represents a system that consists of two processors (PE0 and PE2) connected to a bus (PE1). Tasks mapped to a communication network deliver messages between two computation tasks; for example τ_1 indicates message communication between two computation tasks, τ_0 and τ_2.

We assume that the scheduling policy of a PE can be either a fixed-priority preemptive scheduling or a fixed-priority non-preemptive scheduling. $\mathcal{P}$ denotes a set of PEs that have preemptive scheduling policy, and $\mathcal{N}$ denotes a set of PEs with non-preemptive scheduling policy. A PE belongs to either $\mathcal{P}$ or $\mathcal{N}$. In Fig. 1, PE0 and PE2 use preemptive scheduling while PE1 serves the communication tasks in a non-preemptive fashion; a higher-priority message cannot preempt the current message delivery. We assume that all tasks mapped to each PE have distinct priorities to make the scheduling order deterministic. The priority of the task τ is denoted by τ^{pri}.

The WCRT of task graph $\mathcal{T}$, denoted by $\mathcal{R}_{\mathcal{T}}$, is defined as the time difference between the latest finish time and the earliest release time among tasks in the task graph.

4. REVIEW OF THE Y&W METHOD

Since the Y&W method is known as a state-of-the art schedulability analysis technique that considers dependency between tasks directly, we select it as the reference technique for comparison in this paper. In this section we review the key ideas of the Y&W method and prove that it fails to find a conservative upper bound of the WCRT. Since dependency between tasks constrains the release times of tasks, the Y&W method proposed three techniques: *separation analysis*, *phase adjustment*, and *period shifting*.
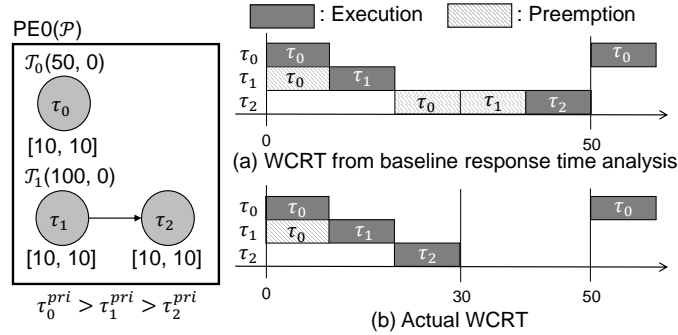
Fig. 2. A simple example of WCRT estimation

The baseline response time analysis for independent tasks defines the response time r_i of a task τ_i as an iterative form:

$$r_i = \tau_i^{WCET} + \sum_{\tau_j \in \{\tau_j | \tau_i^{pri} < \tau_j^{pri}\}} \left\lceil \frac{r_i}{\mathcal{T}_{\tau_j}^p} \right\rceil \cdot \tau_j^{WCET} \qquad (1)$$

where set $\tau_j \in \{\tau_j | \tau_i^{pri} < \tau_j^{pri}\}$ represents the set of higher priority tasks that may preempt task τ_i. Note that the second term subsumes all possible preemption delays by higher priority tasks for conservative estimation. When the baseline RTA is directly applied to tasks with dependency, the analytical WCRT of a task graph may

become much larger than the actual WCRT. Consider an example in Fig. 2. Unless the task dependency is considered, τ_0 can preempt both τ_1 and τ_2, and τ_1 can preempt τ_2. Therefore, WCRT becomes $r_1 + r_2 = 50$ as shown in Fig. 2 (a) while the tight WCRT is 30. If the dependency is considered, τ_0 can preempt either τ_1 or τ_2 but not both, and τ_1 cannot preempt τ_2, as shown in Fig. 2 (b).

In the Y&W method, the release times and the finish times of tasks are computed, considering data dependencies between tasks. If the earliest release time (τ_j^{minR}) of a higher priority task τ_j is larger than the latest finish time (τ_i^{maxF}) of a given task τ_i, τ_j cannot preempt τ_i. And if there is a direct path between a higher priority task and a given task, no preemption will occur. These cases are identified in the *separation analysis* to check the preemption possibility between two tasks in the same task graph; the *separation analysis* finds that τ_1 cannot preempt τ_2 in the example of Fig. 2.

The *phase adjustment* technique in the Y&W method computes the distance, called phase, between the release times of the preempting and the preempted tasks to identify the impossible preemptions along dependent tasks. The request phase $\phi_{i,j}^r$ means the minimum distance from the request(release) time of τ_i to the next release time of τ_j. Note that if the dependency is not considered, the request phase will be 0 since the worst case preemption occurs when the preempting task is released at the same time as the preempted task. The request phase is dependent on the finish phase of its predecessors. The finish phase $\phi_{i,j}^f$ means the minimum distance from the finish time of τ_i to the next release time of τ_j. The *phase adjustment* technique modifies the response time formula of equation (1) as follows:

$$\phi_{i,j}^r = \max\left(0, \min_{\tau_k \in pred(\tau_i)} (\phi_{k,j}^f + \tau_k^{maxF}) - \max_{\tau_k \in pred(\tau_i)} \tau_k^{maxF}\right) \tag{2}$$

$$r_i = \tau_i^{WCET} + \sum_{\tau_j \in \{\tau_j | \tau_i^{pri} < \tau_j^{pri}\} - separated[\tau_i]} \left\lceil \frac{\max(0, r_i - \phi_{i,j}^r)}{\mathcal{T}_{\tau_j}^p} \right\rceil \cdot \tau_j^{WCET} \tag{3}$$

$$\phi_{i,j}^f = \begin{cases} \max(0, \phi_{i,j}^r - r_i), & \text{if } \tau_j \notin \{\tau_j | \tau_i^{pri} < \tau_j^{pri}\} - separated[\tau_i] \\ (\phi_{i,j}^r - r_i) \bmod \mathcal{T}_j^p & \text{otherwise} \end{cases} \tag{4}$$

where $separated[\tau_i]$ is a set of higher priority tasks excluded in the preemption delay computation by the *separation analysis*, and $pred(\tau_i)$ is the immediate predecessor set of task τ_i. In the modified formula, the request phase is subtracted from the response time in the preemption count computation. For conservative computation, the request phase and the finish phase are made non-negative. For detailed explanation of the formula, refer to [Yen and Wolf 1998].

For a preempting task τ_j, we compute the request phase and the finish phase for each task τ_i of a task graph. For a source task τ_i, $\phi_{i,j}^r$ is set to 0, meaning that the preempting task can be released at the same time to make τ_i experience the maximum number of preemptions by τ_j. After computing the response time, finish phase $\phi_{i,j}^f$ is updated. The request phase for a non-source task is updated in turn based on the finish phases of its predecessors. In the example of Fig. 2, the request phase $\phi_{1,0}^r$ is initialized to 0 and the response time of τ_1 is 20. The finish phase $\phi_{1,0}^f$ is updated to $(0 - 20) \bmod 50 = 30$, and the request phase $\phi_{2,0}^r$ is inherited from $\phi_{1,0}^f$. Since τ_1 and τ_2 are separated, the response time of τ_2 becomes $10 = 10 + \lceil \frac{max(0,10-30)}{50} \rceil \cdot 10$. If we sum the response times of τ_1 and τ_2, we obtain the actual WCRT that is 30.

Consider a preempting task with varying release time. The release time may vary due to finish time variation of its predecessor. Then, the preempting task will not be scheduled periodically, but in a bursty fashion, which increases the preemption count. The *period shifting* technique is used in the Y&W method to account for this effect. Let τ_j be a task that may preempt τ_i and of which release time varies between τ_j^{minR} and τ_j^{maxR} that indicate the minimum and maximum release time of τ_j respectively. To compute the maximum number of preemptions by τ_j onto τ_i, we need to consider the release time difference.

The authors of the Y&W method did not clarify how to integrate the *period shifting* technique and the *phase adjustment* technique in a single formula. If both techniques are applied separately in sequence, the conventional RTA equation is modified to the following formula, which is used as the Y&W method in this paper.

$$r_i = \tau_i^{WCET} + \sum_{\tau_j \in \{\tau_j | \tau_i^{pri} < \tau_j^{pri}\} - separated[\tau_i]} \left\lceil \frac{\max(0, r_i - \phi_{i,j}^r + \tau_j^{maxR} - \tau_j^{minR})}{\mathcal{T}_{\tau_j}^p} \right\rceil \cdot \tau_j^{WCET}$$

(5)

4.1. The Y&W method is NOT conservative

Unfortunately, the Y&W method fails to find the maximum preemption count. Consider a simple example of Fig. 3. Since τ_2 has no predecessor and a zero jitter, release time is always 0. Then the response time of τ_1 from the Y&W method becomes $25 = 20 + \lceil \frac{(25+0-0)}{30} \rceil \cdot 5$ as shown in Fig. 3 (a), in which τ_1 can be preempted once by τ_2. But τ_2 can preempt τ_1 twice and the response time of τ_1 can be as large as 30 as shown in Fig. 3 (b) when the start of τ_2 is delayed by the preemption of τ_0. Since the Y&W method only considers the release time variation by predecessors but not the start time variation by preemption, it fails to obtain a conservative WCRT.
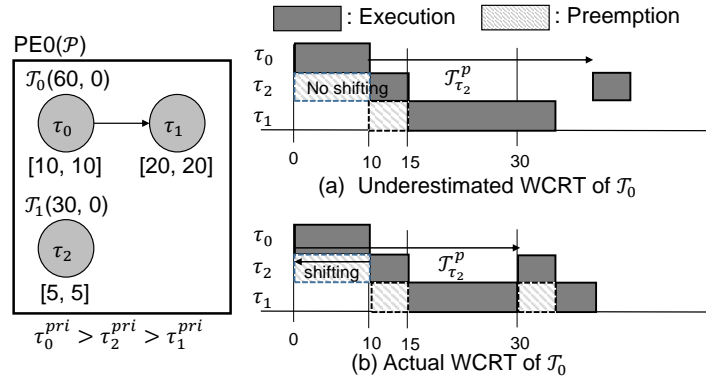
Fig. 3. An underestimated WCRT example in the Y&W method

Another under-estimate example of Fig. 4 shows that period shifting and phase adjustment are inter-dependent while the Y&W method treats them separately. In the Y&W method, $\phi_{1,3}^f$ and $\phi_{2,3}^r$ are 60 assuming that task τ_3 is first released at time 50 simultaneously with τ_1 and the next release will be 100 time units after ignoring the period shifting effect. But the worst case of preemption occurs when τ_3 is first released at time 0. Then the next release of τ_3 can appear after τ_2 executes 10 time units as

shown in Fig. 4. So, the WCRT of $\mathcal{T}_0$ is 110 since τ_3 appears once per $\mathcal{T}_0$ execution in the Y&W method while it is 130 since τ_3 appears twice per $\mathcal{T}_0$ execution in the worst case. Motivated from this example, the proposed technique considers period shifting and phase adjustment holistically, which will be explained in the next section.

Fig. 4.   An example of incorrect phase adjustment

4.2. Overestimation made by the Y&W method
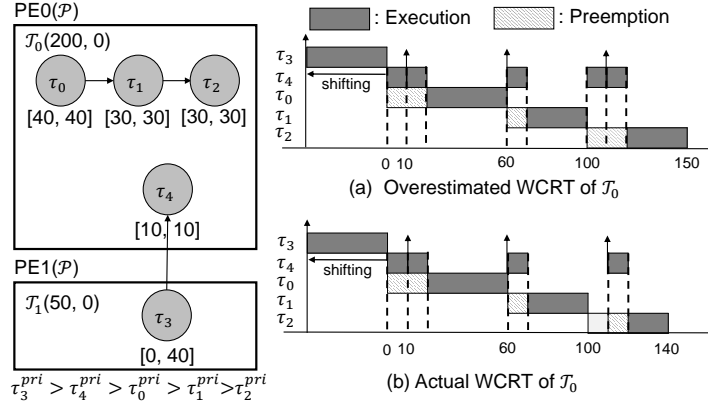
Fig. 5.   Overestimation due to repeated period shifting

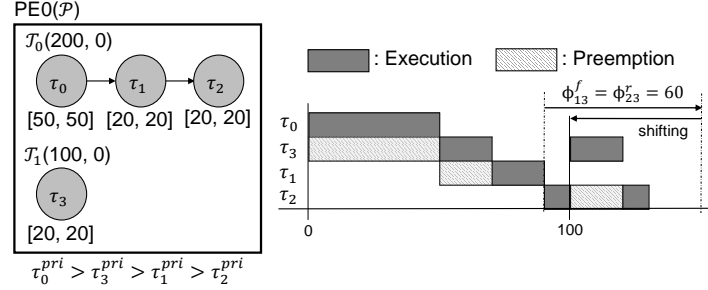In the period shifting technique, the release time difference is added to the response time in the preemption delay computation. It should be applied only once in a sequence of releases of the preempting task. But the Y&W method applies period shifting to all tasks independently, as shown in the example of Fig. 5 where period shifting is applied to all tasks, τ_0 to τ_2, in $\mathcal{T}_0$. As a result τ_4 makes 5 preemptions in total to make the WCRT of the task graph be 150. But the actual WCRT is 140 as shown in Fig. 5 (b) since τ_4 can make 4 preemptions at most. In the Y&W method, period shifting value for τ_4 is 40 because of the execution time variation of its predecessor τ_3. Then τ_4 preempts τ_0 twice to make the WCRT of τ_0 be $60 = 40 + \lceil \frac{(60-0+40)}{50} \rceil \cdot 10$. $\phi_{1,4}^r$ is computed to $40 = (0 - 60) \bmod 50$, and the WCRT of τ_1 is $40 = 30 + \lceil \frac{(40-40+40)}{50} \rceil \cdot 10$. Since $\phi_{2,4}^r$ becomes $0 = 40 - 40$, the WCRT of τ_2 is $50 = 30 + \lceil \frac{(50-0+40)}{50} \rceil \cdot 10$, experiencing two preemptions. By integrating period shifting into phase adjustment we could improve

the tightness of the WCRT by removing the problem of redundant application of period shifting, which will be explained in the next section.
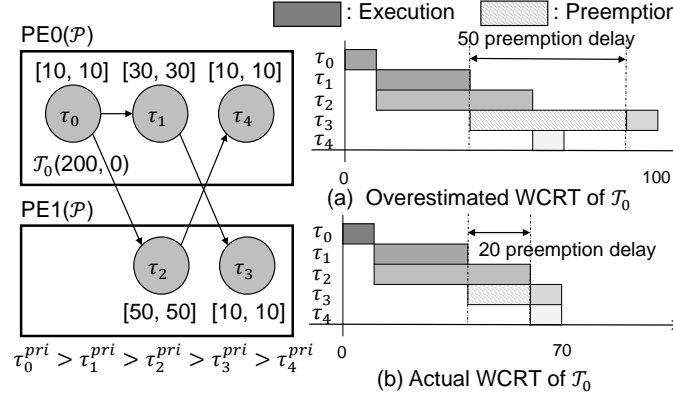

Fig. 6.   An example of overestimation (Y&W method)

Another cause of overestimation comes from the fact that the Y&W method does not consider partial preemptions. In the example of Fig. 6, τ_1 cannot preempt τ_4 since two tasks are separated. On the other hand, τ_2 and τ_3 are not separated because τ_3 can be released during the execution of τ_2. In the Y&W method, however, the WCET of τ_3 is wholly added to the preemption delay of τ_2. Partial preemption may occur between tasks in the same task graph as shown in this example, which is not considered in the baseline response time analysis. In the proposed analysis, however, we perform scheduling of tasks in the same task graph so that we could detect this kind of partial preemption precisely.

5. PROPOSED ANALYSIS TECHNIQUE: HYBRID PERFORMANCE ANALYSIS

The proposed technique called HPA (hybrid performance analysis) extends and combines the STBA approach and the RTA method: the former is to account for interference between tasks in a same task graph and the latter for interference from the other task graphs.

Fig. 7 shows the algorithm flow of the proposed technique. First, we compute three pairs of time bound information for each task: release time bound(τ_i^{minR}, τ_i^{maxR}), start time bound (τ_i^{minS}, τ_i^{maxS}), and finish time bound (τ_i^{minF}, τ_i^{maxF}). Unlike the STBA technique [Kim et al. 2013] that assumes to schedule all task graphs together, the proposed technique schedules each task graph separately at the time bound computation step. The interference from the other task graphs is considered by the holistic phase adjustment technique based on the period shifting amount computed in the previous iteration. A period shifting amount is updated after time bound computation. This process is repeated until every value is converged.
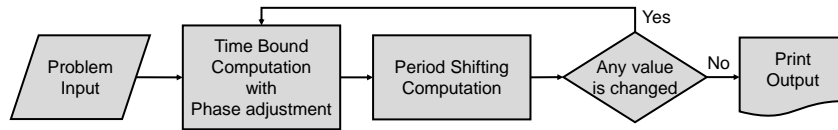

Fig. 7.   The proposed HPA technique overall flow

In this section, we explain the key techniques of the proposed analysis and discuss how to achieve a safe and tighter bound of the WCRT. At first, how to compute time bounds is explained. Secondly, we derive and optimize a conservative bound of a period shifting amount from τ_i to τ_t which bounds the maximum preemption count of τ_i between release time and the maximum start time of τ_t. Finally, *period shifting* is combined with *phase adjustment* technique in the proposed *holistic phase adjustment* technique.

5.1. Time Bound Computation

Let $\tau_t^{release}$, τ_t^{start}, and τ_t^{finish} denote the actual release time, start time, and finish time of task τ_t. In our task graph model, the release time of a task is the maximum finish time of its immediate predecessors , as summarized in the following definition.

DEFINITION 1.

$$\tau_t^{release} = \begin{cases} \mathcal{T}_{\tau_t}^{release}, & \text{if } \tau_t \text{ is a source task} \\ \max_{\tau_p \in pred(\tau_t)} \tau_p^{finish}, & \text{otherwise} \end{cases}$$

where $\mathcal{T}_{\tau_t}^{release}$ denotes the release time of the task graph. Then the minimum (τ_t^{minR}) and maximum (τ_t^{maxR}) release time bound pair is computed as follows:

$$\tau_t^{minR} = \begin{cases} 0, & \text{if } \tau_t \text{ is a source task} \\ \max_{\tau_p \in pred(\tau_t)} \tau_p^{minF}, & \text{otherwise} \end{cases} \tag{6}$$

$$\tau_t^{maxR} = \begin{cases} \mathcal{T}_{\tau_t}^{j}, & \text{if } \tau_t \text{ is a source task} \\ \max_{\tau_p \in pred(\tau_t)} \tau_p^{maxF}, & \text{otherwise} \end{cases} \tag{7}$$

The earliest and the latest release times of a non-source task are defined as the maximum value among the earliest and the latest finish times of predecessors, respectively since it becomes executable only after all predecessor tasks are finished.

LEMMA 1. τ_t^{minR} and τ_t^{maxR} are conservative, or $\tau_t^{minR} \leq \tau_t^{release} \leq \tau_t^{maxR}$.

PROOF.   If τ_t is a source task, $\tau_t^{minR} = 0 \leq \tau_t^{release} \leq \mathcal{T}_{\tau_t}^{j} = \tau_t^{maxR}$ since $0 \leq \mathcal{T}_{\tau_t}^{release} \leq \mathcal{T}_{\tau_t}^{j}$.

For non-source task τ_t, assume that it holds for all predecessor tasks of τ_t. Since $\tau_t^{minR} = \max_{\tau_p \in pred(\tau_t)} \tau_p^{minF} \leq \max_{\tau_p \in pred(\tau_t)} \tau_p^{finish} = \tau_t^{release} \leq \max_{\tau_p \in pred(\tau_t)} \tau_p^{maxF} = \tau_t^{maxR}$, $\tau_t^{minR} \leq \tau_t^{release} \leq \tau_t^{maxR}$.

By induction, the lemma holds. Q.E.D.   □

For task τ_t to start, it should be already released and the processor must be available: The start time of τ_t is not smaller than the release time and the maximum time among finish times of tasks that have higher priority, start earlier, and finish after task τ_t is released. Formally, the start time is defined as follows:

DEFINITION 2. $\tau_t^{start} = \max(\tau_t^{release}, \max_{\tau_s \in \Gamma_{\tau_t}} \tau_s^{finish})$

where set Γ_{τ_t} is defined as $\Gamma_{\tau_t} = \{\tau_s | \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_t^{release} < \tau_s^{finish}, \tau_s^{start} \leq \tau_t^{start}\}$ for the preemptive scheduling policy, and $\Gamma_{\tau_t} = \{\tau_s | \tau_s^{proc} = \tau_t^{proc}, (\tau_s^{pri} > \tau_t^{pri}, \tau_t^{release} < \tau_s^{finish}, \tau_s^{start} \leq \tau_t^{start}) or (\tau_s^{pri} < \tau_t^{pri}, \tau_s^{start} < \tau_t^{release} < \tau_t^{finish})\}$ for the non-preemptive scheduling policy.

One the other hand, the earliest start time τ_t^{minS} is formulated as follows:

$$\tau_t^{minS} = \max(\tau_t^{minR}, \max_{\tau_s \in \mathcal{A}_{\tau_t}} \tau_s^{minF}) \tag{8}$$

where set $\mathcal{A}_{\tau_t}$ for the preemptive scheduling policy is defined as $\mathcal{A}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_t^{minR} < \tau_s^{minF}, \tau_s^{maxS} \leq \tau_t^{minS}\}$ and for the non-preemptive scheduling policy as $\mathcal{A}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, (\tau_s^{pri} > \tau_t^{pri}, \tau_t^{minR} < \tau_s^{minF}, \tau_s^{maxS} \leq \tau_t^{minS}) or (\tau_s^{pri} < \tau_t^{pri}, \tau_s^{maxS} < \tau_t^{minR} < \tau_s^{minF})\}$. If higher priority task τ_s always starts before τ_t^{minS} and the earliest finish time of τ_s is later than τ_t^{minR}, τ_t should wait for the completion of τ_s. In case a non-preemptive scheduling is used, a lower priority task that starts before τ_t^{minR} is included. Note that we only consider the tasks in the same task graph since the other tasks can appear at any time, so they should be ignored for conservative estimation of the minimum start time.

To estimate the maximum start time τ_t^{maxS} for conservative estimation, we should consider all possible preemptions.

$$\tau_t^{maxS} = \tau_t^{maxR} + Delay_t^l + Delay_t^h \tag{9}$$

where $Delay_t^l$ and $Delay_t^h$ denote the amounts of preemption between the release time and the start time by lower and higher priority tasks respectively. For the preemptive scheduling policy, $Delay_t^l$ is zero. In case a lower priority task is running when τ_t is released, τ_t should wait until the current lower priority task finishes in the non-preemptive scheduling policy, which is accounted as follows:

$$Delay_t^l = \begin{cases} 0, & \text{if } \forall_{\tau_p \in pred(\tau_t)}\left(\tau_p^{proc} = \tau_t^{proc}\right) \\ \max \left(\begin{array}{c} \max\limits_{\tau_s \in \mathcal{B}_{\tau_t}} \min(\tau_s^{WCET}, \tau_s^{maxF} - \tau_t^{maxR}), \\ \max\limits_{\tau_s \in \mathcal{C}_{\tau_t}} \tau_s^{WCET} \end{array} \right), & \text{otherwise} \end{cases} \tag{10}$$

where $\mathcal{B}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} < \tau_t^{pri}, \tau_s^{minS} < \tau_t^{maxR} < \tau_s^{maxF}\}$ and $\mathcal{C}_{\tau_t} = \{\tau_s | \tau_s \notin \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} < \tau_t^{pri}\}$. Set $\mathcal{B}_{\tau_t}$ includes lower priority tasks in the same task graph that may start earlier than τ_t and delay the start time of τ_t. Note that partial blocking is considered in the formula by $\tau_s^{maxF} - \tau_t^{maxR}$, which can be smaller than τ_s^{WCET}. On the other hand, set $\mathcal{C}_{\tau_t}$ includes all lower priority tasks in the other task graphs. Since they can appear at any time, we take the maximum WCET for conservative estimation. In case every predecessor is mapped to the same PE, $Delay_t^l$ is zero.

$Delay_t^h$ is commonly formulated for both scheduling policies as follows:

$$Delay_t^h = \sum_{\tau_s \in \mathcal{D}_{\tau_t}} \min(\tau_s^{WCET}, \tau_s^{maxF} - \tau_t^{maxR}) + \\ \sum_{\tau_s \in \mathcal{E}_{\tau_t}} \left\lceil \frac{\max(0, \tau_t^{maxS} - \tau_t^{maxR} + 1 - \phi_{t,s}^r)}{\mathcal{T}_{\tau_s}^p} \right\rceil \cdot \tau_s^{WCET} \tag{11}$$

where $\mathcal{D}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_s^{minS} \leq \tau_t^{maxS}, \tau_t^{maxR} < \tau_s^{maxF}\}$ and $\mathcal{E}_{\tau_t} = \{\tau_s | \tau_s \notin \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}\}$. Set $\mathcal{D}_{\tau_t}$ includes higher priority tasks in the same task graph that can possibly delay the start time of τ_t. Partial preemption is considered similarly to $Delay_t^l$ formulation. For the example of Fig. 6, we precisely compute the preemption delay from τ_2 to τ_3 as $\tau_2^{maxF} - \tau_3^{maxR} = 60 - 40 = 20$, which is less that τ_2^{WCET}. For the higher priority tasks in the other task graphs, we use a similar formula as the response time analysis to compute the maximum preemption delay. The notation $\phi_{t,s}^r$ corresponds to the request phase adjustment that computes the minimum distance from the release time of τ_t to the next release time of a pre-

empting task τ_s. How to formulate the request phase $\phi_{t,s}^r$ will be explained in the next subsection.

LEMMA 2. τ_t^{minS} and τ_t^{maxS} are conservative, or $\tau_t^{minS} \leq \tau_t^{start} \leq \tau_t^{maxS}$.

PROOF. First we prove that $\tau_t^{minS} \leq \tau_t^{start}$.
Since $\tau_t^{minR} \leq \tau_t^{release}$, we need to prove $\max_{\tau_s \in \mathcal{A}_{\tau_t}} \tau_s^{minF} \leq \tau_t^{start}$.

$$\max_{\tau_s \in \mathcal{A}_{\tau_t}} \tau_s^{minF} \leq \max(\max_{\tau_s \in \Gamma_{\tau_t}} \tau_s^{minF}, \max_{\tau_s \in \mathcal{A}_{\tau_t} - \Gamma_{\tau_t}} \tau_s^{minF}).$$

Since $\mathcal{A}_{\tau_t} - \Gamma_{\tau_t} \subseteq \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{release} \geq \tau_s^{finish}\}$ for both preemptive and nonpreemptive cases, and $\tau_s^{finish} \geq \tau_s^{minF}$, $\max_{\tau_s \in \mathcal{A}_{\tau_t} - \Gamma_{\tau_t}} \tau_s^{minF} \leq \tau_t^{release}$.

$$\max_{\tau_s \in \mathcal{A}_{\tau_t}} \tau_s^{minF} \leq \max(\max_{\tau_s \in \Gamma_{\tau_t}} \tau_s^{finish}, \tau_t^{release}) = \tau_t^{start}.$$

Next, We prove that $\tau_t^{start} \leq \tau_t^{maxS} = \tau_t^{maxR} + Delay_t^l + Delay_t^h$ by showing that no task will delay the start of τ_t without being considered in $Delay_t^l$ or $Delay_t^h$ computation.

(1) Suppose actual preemption amount by lower priority tasks is larger than $Delay_t^l$. For preemptive scheduling policy, it is impossible because a lower priority task cannot preempt τ_t. For nonpreemptive scheduling policy, only one lower priority task can delay the execution of τ_t. For $\tau_s \notin \mathcal{T}_{\tau_t}$, since $\mathcal{C}_{\tau_t}$ includes all lower priority tasks, it is not possible for the preemption amount to be larger than the maximum worst case execution time of tasks in $\mathcal{C}_{\tau_t}$.
Consider $\tau_s \in \mathcal{T}_{\tau_t}$. τ_s can delay τ_t by $\tau_s^{finish} - \tau_t^{release}$ at most if τ_s starts before τ_t is released or $\tau_s^{start} < \tau_t^{release}$.

$$\tau_t^{release} + \max_{\tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{finish} - \tau_t^{release} > 0} (\tau_s^{finish} - \tau_t^{release})$$

$$\leq \tau_t^{maxR} + \max_{\tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{finish} - \tau_t^{maxR} > 0} (\tau_s^{finish} - \tau_t^{maxR})$$

$$\leq \tau_t^{maxR} + \max_{\tau_s \in \mathcal{T}_{\tau_t}, \tau_t^{maxR} < \tau_s^{maxF}} (\tau_s^{maxF} - \tau_t^{maxR}).$$

Since $\{\tau_s | \tau_s^{start} < \tau_t^{release}\} \subseteq \{\tau_s | \tau_s^{minS} < \tau_t^{maxR}\}$ and $\{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} < \tau_t^{pri}, \tau_t^{maxR} < \tau_s^{maxF}\} \cap \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} < \tau_t^{pri}, \tau_s^{minS} < \tau_t^{maxR}\} \subseteq \mathcal{B}_{\tau_t}$,

$$\tau_t^{release} + \max_{\tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{finish} - \tau_t^{release} > 0} (\tau_s^{finish} - \tau_t^{release}) \leq \tau_t^{maxR} + \max_{\tau_s \in \mathcal{B}_{\tau_t}} (\tau_s^{maxF} - \tau_t^{maxR}).$$

Since τ_s cannot be executed longer than τ_s^{WCET},

$$\tau_t^{release} + \max_{\tau_s \in \mathcal{B}_{\tau_t}} \min(\tau_s^{WCET}, \tau_s^{finish} - \tau_t^{release}) \leq \tau_t^{maxR} + Delay_t^l.$$

Finally, consider the case when all predecessors are mapped to the same processor. Since there is no time interval between the finish time of the latest predecessor task τ_p and the release time of τ_t, no lower priority task can start after τ_p finishes and before τ_t is released.

(2) Consider higher priority tasks that may preempt τ_t. For $\tau_s \in \mathcal{T}_{\tau_t}$ and $\tau_s \in \mathcal{D}_{\tau_t}$, the preemption amount is bounded by $\min(\tau_s^{WCET}, \tau_s^{maxF} - \tau_t^{maxR})$ similarly to the proof of the case $\tau_s \in \mathcal{T}_{\tau_t}$ and $\tau_s \in \mathcal{B}_{\tau_t}$ in 1.
Suppose there is a task τ_s which can preempt τ_t such that $\tau_s \in \mathcal{T}_{\tau_t}$ and $\tau_s \notin \mathcal{D}_{\tau_t}$. Then $\tau_t^{maxS} < \tau_s^{minS}$ or $\tau_s^{maxF} \leq \tau_t^{maxR}$. τ_s always starts later than τ_t since $\tau_t^{maxS} \leq$

τ_s^{minS}, or τ_s finishes earlier than τ_t since $\tau_s^{maxF} \leq \tau_t^{maxR}$. No preemption may be occurred so that there exists no such task.

For $\tau_s \notin \mathcal{T}_{\tau_t}$, the proof is trivial if the request phase is conservatively given. The conservativeness of the request phase, $\phi_{t,s}^r$ will be proven in Lemma 5.

By 1 and 2, $\tau_t^{start} \leq \tau_t^{maxS}$. Q.E.D. $\square$

The minimum finish time τ_t^{minF} is formulated as follows:

$$\tau_t^{minF} = \tau_t^{minS} + \tau_t^{BCET} + Preempt_t^B \tag{12}$$

where $Preempt_t^B$ represents the unavoidable (best-case) preemption delay that is zero for the non-preemptive scheduling policy. For the preemptive scheduling policy, $Preempt_t^B$ becomes

$$Preempt_t^B = \sum_{\tau_s \in \mathcal{F}_{\tau_t}} \tau_s^{BCET} \tag{13}$$

where $\mathcal{F}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_t^{minS} \leq \tau_s^{minS} \leq \tau_s^{maxS} \leq \tau_t^{minF}\}$. $\mathcal{F}_{\tau_t}$ is a set of higher priority tasks which always start to execute and preempt τ_t during τ_t is running from τ_t^{minS} to τ_t^{minF}.

The maximum finish time τ_t^{maxF} is formulated as follows:

$$\tau_t^{maxF} = \tau_t^{maxS} + \tau_t^{WCET} + Preempt_t^W \tag{14}$$

where $Preempt_t^W$ represents the worst-case preemption delay that is zero for the non-preemptive scheduling policy. $Preempt_t^W$ for the preemptive scheduling policy is formulated as follows:

$$Preempt_t^W = \sum_{\tau_s \in \mathcal{G}_{\tau_t}} \tau_s^{WCET} + \sum_{\tau_s \in \mathcal{E}_{\tau_t}} \left\lceil \frac{\max(0, \tau_t^{maxF} - \tau_t^{maxS} - \phi_{t,s}^s)}{\mathcal{T}_{\tau_s}^p} \right\rceil \cdot \tau_s^{WCET} \tag{15}$$

where $\mathcal{G}_{\tau_t} = \{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_t^{maxS} < \tau_s^{minS} \leq \tau_t^{maxF}\}$, indicating a set of higher priority tasks which can appear during the execution of τ_t. The notation $\phi_{t,s}^s$ is the start phase that is the the minimum distance from the start time of τ_t to the next release time of a preempting task τ_s. For tasks in the other task graphs, the maximum possible preemption delay is computed similarly to (11). The start phase will be explained in the next subsection.

LEMMA 3. τ_t^{minF} and τ_t^{maxF} are conservative, or $\tau_t^{minF} \leq \tau_t^{finish} \leq \tau_t^{maxF}$.

PROOF. (Nonpreemptive) If non-preemptive scheduling is used then no task may preempt τ_t. Hence the finish time is the sum of the start time and the execution time.

(Preemptive) First, we prove that $\tau_t^{minF} \leq \tau_t^{finish}$. Let $\mathcal{W}_{\tau_t}$ be $\{\tau_s | \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}\}$.

$\tau_t^{finish} \geq \tau_t^{start} + \tau_t^{BCET} + \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{start} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET}$ where the last term accounts for tasks that start to execute during τ_t is running, and preempt τ_t.

Since $\sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{start} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET} \leq \tau_t^{finish} - \tau_t^{start}$,

and $\sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{minS} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET} \leq \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{start} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET} + (\tau_t^{start} -$

τ_t^{minS}),

$$\tau_t^{finish} \geq \tau_t^{start} + \tau_t^{BCET} + \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{start} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET}$$

$$\geq \tau_t^{minS} + \tau_t^{BCET} + \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{minS} \leq \tau_s^{start} \leq \tau_t^{finish}} \tau_s^{BCET}$$

$$\geq \tau_t^{minS} + \tau_t^{BCET} + \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{minS} \leq \tau_s^{minS} \leq \tau_s^{maxS} \leq \tau_t^{finish}} \tau_s^{BCET}$$

since $\tau_s^{minS} \leq \tau_s^{start} \leq \tau_s^{maxS}$.

$$\tau_t^{finish} \geq \tau_t^{minF} = \tau_t^{minS} + \tau_t^{BCET} + \sum_{\tau_s \in \mathcal{W}_{\tau_t}, \tau_t^{minS} \leq \tau_s^{minS} \leq \tau_s^{maxS} \leq \tau_t^{minF}} \tau_s^{BCET}.$$

Second, we prove that $\tau_t^{finish} \leq \tau_t^{maxF}$ by contradiction.

For $\tau_s \in \mathcal{T}_{\tau_t}$, suppose there is some $\tau_s \notin \mathcal{G}_{\tau_t}$ that satisfies $\tau_s \in \mathcal{T}_{\tau_t}$ and can preempt τ_t. Then $\tau_s^{minS} \leq \tau_t^{maxS}$ or $\tau_t^{maxF} < \tau_s^{minS}$. The tasks that satisfies $(\tau_s \notin \mathcal{G}_{\tau_t}, \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{minS} \leq \tau_t^{maxS}, \tau_t^{maxF} < \tau_s^{maxF})$ belongs to the $\mathcal{D}_{\tau_t}$, so that the amount of preemption from those tasks are already included in τ_t^{maxS} and so in τ_t^{maxF}. The tasks that satisfies $(\tau_s \notin \mathcal{G}_{\tau_t}, \tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{minS} \leq \tau_t^{maxS}, \tau_s^{maxF} \leq \tau_t^{maxR})$ cannot preempt τ_t since τ_s finishes before τ_t^{maxR}. The tasks that satisfies $(\tau_s \notin \mathcal{G}_{\tau_t}, \tau_s \in \mathcal{T}_{\tau_t}, \tau_t^{maxF} < \tau_s^{minS})$ cannot preempt τ_t since τ_s starts after τ_t^{maxF}. For $\tau_s \notin \mathcal{T}_{\tau_t}$, the proof is trivial since we assume that the start phase is conservatively given. The conservativeness of the start phase, $\phi_{t,s}^s$ will be proven in Lemma 5. Q.E.D. □

After determining all time bounds of tasks, we compute the WCRT of each task graph $\mathcal{T}$ as follows:

$$\mathcal{R}_\mathcal{T} = \max_{\tau_s \in \mathcal{T}} \tau_s^{maxF} \tag{16}$$

THEOREM 1. *The HPA technique guarantees the conservativeness of every schedule time bound when it is converged.*

PROOF. Theorem 1 is proved by Lemma 1, Lemma 2, and Lemma 3. Q.E.D. □
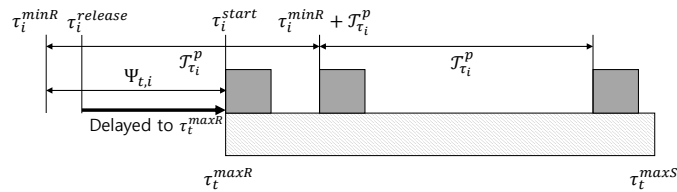
5.2. Period Shifting Computation

Fig. 8.   The worst case scenario of the preemption by τ_i to τ_t

When computing the maximum time bounds, we have to consider the worst case scenario of preemptions in the RTA analysis. Fig. 8 shows this scenario. Suppose that the target task τ_t is released at τ_t^{maxR}. The worst case preemptions of a higher priority

task τ_i occur when its start time is aligned with the release time of τ_t and the second request appears with the shortest interval from τ_t^{maxR}, followed by later requests that appear periodically from the second request. If the period shifting amount is denoted by $\Psi_{t,i}$, the next start time of τ_i will be $\max(\tau_t^{maxR} + \tau_i^{WCET}, \tau_t^{maxR} - \Psi_{t,i} + \mathcal{T}_t^p)$.
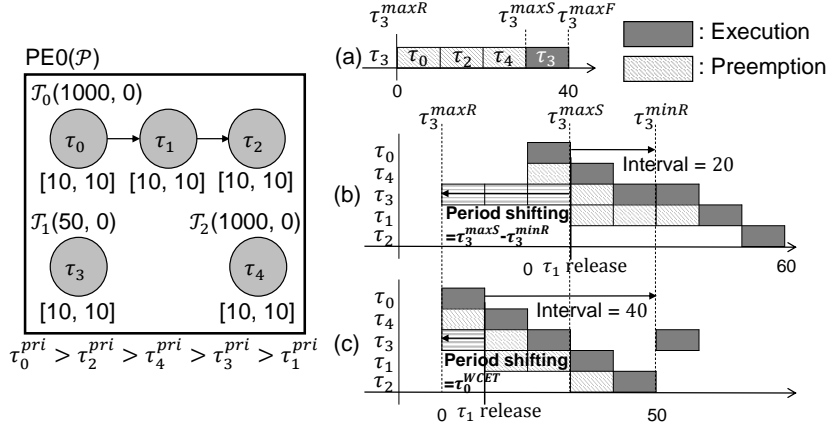

Fig. 9. An example for period shifting computation. (a) Possible preemptions to τ_3 between τ_3^{maxR} and τ_3^{maxS}, (b) estimated WCRT of $\mathcal{T}_0$ with period shifting $\tau_3^{maxS} - \tau_3^{minR}$, and (c) actual WCRT of $\mathcal{T}_0$

Recall that the Y&W method defines $\Psi_{t,i}$ as $\tau_i^{maxR} - \tau_i^{minR}$ and it fails to find the worst case behavior as shown in the example of Fig. 3. A naive revision is to define $\Psi_{t,i}$ as $\tau_i^{maxS} - \tau_i^{minR}$, by aligning the maximum start time of τ_i with the maximum release time of τ_t. The worst case bursty requests occurs when the preempting task is released at $\tau_t^{maxR} - \Psi_{t,i}$. In the example of Fig. 3, we can obtain the actual WCRT if τ_2^{maxS} is aligned at τ_1^{maxR}, which makes $\Psi_{1,2} = \tau_2^{maxS} - \tau_2^{minR} = 10$. However, it generates a loose bound for the example of Fig. 9. Fig. 9 (a) shows the worst case preemption of τ_3, considering all possible preemptions from the other tasks. From Fig. 9 (a), it is known that $\tau_3^{maxS} - \tau_3^{minR}$ is 30. If this value is used as period shifting, a loose WCRT is obtained as shown in Fig. 9 (b). The actual WCRT is computed when $\Psi_{1,3}$ is equal to τ_0^{WCET}, as shown in the Fig. 9 (c), which confirms that more elaborate preemption analysis is needed between τ_3^{maxR} and τ_3^{maxS}. For the conservative but tight bound computation of period shifting amount, we classify possible preemptions between τ_3^{maxR} and τ_3^{maxS} into three categories. The first is the preemptions that affect the start times of both τ_3 and τ_1, which is the amount of τ_4^{WCET} in Fig. 9. These preemptions can be ignored for $\Psi_{1,3}$ since those are considered in the computation of τ_1^{maxS}. The second is the preemptions from the tasks that finish earlier than τ_1 such as τ_0 in Fig. 9, which should be considered for the $\Psi_{1,3}$. The last is the preemptions from the task that starts later than τ_1 such as τ_2 that is a descendant of τ_1. These also can be ignored since they cannot appear before τ_1^{maxS}.

Based on these observations, we conservativley consider the tasks in the second category in the period shifting computation to get a tighter bound than $\tau_i^{maxS} - \tau_i^{minR}$: that is

$$\Psi_{t,i} = \tau_i^{maxR} - \tau_i^{minR} + \delta_{t,i} \qquad (17)$$

where $\delta_{t,i} = \sum_{\tau_s^{pri} > \tau_i^{pri}, \tau_s \in \mathcal{E}_{\tau_t}} \left(\left\lceil \frac{R_{\tau_t} + \delta_{t,i} + \Psi_{t,s}}{\mathcal{T}_{\tau_s}^p} \right\rceil - \left\lceil \frac{R_{\tau_t} - \phi_{t,s}^r}{\mathcal{T}_{\tau_s}^p} \right\rceil \right) \cdot \tau_s^{WCET} +$

$\sum_{\tau_s \in \mathcal{T}_{\tau_t}, \tau_s^{pri} > \tau_i^{pri}, (\tau_t^{maxR} - \delta_{t,i} \le \tau_s^{maxR} < \tau_t^{maxR} \ or \ \tau_t^{maxR} - \delta_{t,i} \le \tau_s^{maxR} - \mathcal{T}_{\tau_s}^p < \tau_t^{maxR})} \tau_s^{WCET}$, and $R_{\tau_t} = \tau_t^{maxS} - \tau_t^{maxR}$.

LEMMA 4. *Period shifting is conservative, which means that* $\left\lceil \frac{\tau_t^{maxF} - \tau_t^{maxR} + \Psi_{t,i}}{\mathcal{T}_{\tau_i}^p} \right\rceil$ *is an upper bound of preemption counts of* τ_i *in time window* $[\tau_t^{maxR}, \tau_t^{maxF}]$.

PROOF. Conservativeness of period shifting is achieved if $\delta_{t,i}$ of equation (17) includes all tasks in the second category of tasks preempting τ_i before τ_t^{maxR}.

At first, consider the contribution from $\tau_s \notin \mathcal{T}_{\tau_t}$ to $\delta_{t,i}$. Since the maximum preemption count of τ_t by $\tau_s \in \mathcal{E}_{\tau_t}$ is $\lceil \frac{t + \Psi_{t,s}}{\mathcal{T}_s^p} \rceil$ for t time duration, τ_s appears $\lceil \frac{R_{\tau_t} + \delta_{t,i} + \Psi_{t,s}}{\mathcal{T}_s^p} \rceil$ times during $R_{\tau_t} + \delta_{t,i}$. On the other hand, $\lceil \frac{R_{\tau_t} - \phi_{t,s}^r}{\mathcal{T}_s^p} \rceil$ is the number of preemptions during R_{τ_t}, which is considered in τ_t^{maxS} computation. Therefore the appearance of τ_s is bounded by $\left\lceil \frac{R_{\tau_t} + \delta_{t,i} + \Psi_{t,s}}{\mathcal{T}_{\tau_s}^p} \right\rceil - \left\lceil \frac{R_{\tau_t} - \phi_{t,s}^r}{\mathcal{T}_{\tau_s}^p} \right\rceil$ times during $\delta_{t,i}$.

Second, Consider $\tau_s \in \mathcal{T}_{\tau_t}$. If $\tau_s^{release} > \tau_t^{maxR}$ then τ_s either preempts τ_t as well as τ_i (the first category) or is scheduled later than τ_t (the third category). So we need to consider only τ_s that is released between $\tau_t^{maxR} - \delta_{t,i}$ and τ_t^{maxR}. $\tau_t^{maxR} - \delta_{t,i} < \tau_s^{release} \le \tau_s^{maxR} < \tau_t^{maxR}$. So the worst case is to consider τ_s such that $\tau_t^{maxR} - \delta_{t,i} < \tau_s^{maxR} < \tau_t^{maxR}$. If $\tau_t^{maxR} - \delta_{t,i}$ becomes negative, , the previous instance of τ_s should be considered with the following condition: $\tau_t^{maxR} - \delta_{t,i} < \tau_s^{maxR} - \mathcal{T}_{\tau_s}^p < \tau_t^{maxR}$. Q.E.D. □

Consider Fig. 3. $\delta_{1,2} = \tau_0^{WCET} = 10$ since τ_0 satisfies the condition of the second term of $\delta_{1,2}$, $\tau_1^{maxR} - \delta_{1,2} = 10 - 10 = 0 = \tau_0^{maxR}$. Hence $\Psi_{1,2} = \tau_2^{maxR} - \tau_2^{minR} + \delta_{1,2} = 0 - 0 + 10 = 10$.

Consider Fig. 9. $\delta_{1,3} = \tau_0^{WCET} = 10$. Since $\tau_1^{maxR} - \delta_{1,3} = 10 - 10 = 0 = \tau_0^{maxR}$, τ_0 is included in the second term of $\delta_{1,3}$. On the other hand, τ_2 is excluded in $\delta_{1,3}$ since $10 = \tau_1^{maxR} < \tau_2^{maxR} = 40$. $\tau_4 \notin \mathcal{T}_{\tau_1}$ has no contribution to $\delta_{t,i}$ since $\left(\left\lceil \frac{R_{\tau_1} + \delta_{1,3} + \Psi_{1,4}}{\mathcal{T}_{\tau_4}^p} \right\rceil - \left\lceil \frac{R_{\tau_1} - \phi_{1,4}^r}{\mathcal{T}_{\tau_4}^p} \right\rceil \right) = 1 - 1 = 0$. Hence $\Psi_{1,3} = \tau_3^{maxR} - \tau_3^{minR} + \delta_{1,3} = 0 - 0 + 10 = 10$.

5.3. Holistic Phase Adjustment

In this section, we describe how we combine the *period shifting* technique and the *phase adjustment* technique. There are three phase types considered in the phased adjustment technique; request phase $\phi_{t,i}^r$, start phase $\phi_{t,i}^s$, and finish phase $\phi_{t,i}^f$. The main difference between our *holistic phase adjustment* and the *phase adjustment* of the Y&W method is that the phases in our technique can be negative: If phase is negative, it acts like a period shifting.

If τ_t is a source task, request phase $\phi_{t,i}^r$ for each task $\tau_i \notin \mathcal{T}_{\tau_t}$ is initialized to $-\Psi_{t,i}$, which means that the start time of τ_t is maximally postponed by preemption of τ_i. In this way, period shifting is merged with phase adjustment so that the request phase can be negative.

If τ_t is a non-source task, the request phase $\phi_{t,i}^r$ depends on the finish phases of predecessors. If it is positive, it means the minimum distance from the release of τ_t to the next release time of a preempting task τ_i; Fig. 10 illustrates the case where task τ_t has two predecessors τ_{p_1} and τ_{p_2}.

Since the phase computation is performed per task independently, the finish phases of the predecessor tasks may be different from each other. When we compute the re-
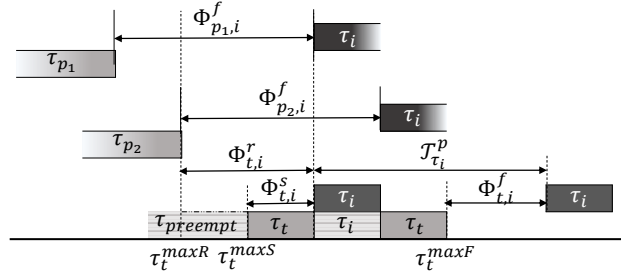
Fig. 10.   An illustrative example of holistic phase adjustment computation

quest phase $\phi_{t,i}^r$ and the predecessor tasks see different next request times of τ_i, we have to take the earliest next request time of τ_i among predecesors for the conservative computation. In Fig. 10, two predecessors τ_{p_1} and τ_{p_2} of τ_t see different next request times of τ_i: $\phi_{p_1,i}^f + \tau_{p_1}^{maxF}$ and $\phi_{p_2,i}^f + \tau_{p_2}^{maxF}$. since $\phi_{p_1,i}^f + \tau_{p_1}^{maxF} < \phi_{p_2,i}^f + \tau_{p_2}^{maxF}$, we take $\phi_{p_1,i}^f + \tau_{p_1}^{maxF}$ as the next request time of τ_i seen by τ_t, and $\phi_{t,i}^r$ becomes $(\phi_{p_1,i}^f + \tau_{p_1}^{maxF}) - \tau_t^{maxR}$, which means the distance from τ_t^{maxR} to the next request time of τ_i.

Note that the inherited phase adjustment can be negative. The negative phase, which acts like a period shifing, cannot be smaller than $-\Psi_{t,i}$. Thus we choose the maximum among the phase adjustment inherited from its predecessors and $-\Psi_{t,i}$. We formulate $\phi_{t,i}^r$ as follows;

$$\phi_{t,i}^r = \begin{cases} -\Psi_{t,i}, & \text{if } \tau_t \text{ is a aource task} \\ & \text{or } \exists_{\tau_p \in pred(\tau_t)}(\tau_p^{proc} \neq \tau_t^{proc}) \\ \max\left(-\Psi_{t,i}, \min_{\tau_p \in pred(\tau_t)}(\phi_{p,i}^f + \tau_p^{maxF}) - \tau_t^{maxR}\right), & \text{otherwise} \end{cases}$$

(18)

where $pred(\tau_t)$ is the immediate predecessor set of task τ_t. In case there is a predecessor mapped to a different processor, $\phi_{t,i}^r$ is set to $-\Psi_{t,i}$ for conservative computation.

Based on $\phi_{t,i}^r$ and τ_t^{maxS}, before computing τ_t^{maxF}, the start phase $\phi_{t,i}^s$ for each task $\tau_i \notin \mathcal{T}_{\tau_t}$ is computed as follows:

$$\phi_{t,i}^s = \begin{cases} (\phi_{t,i}^r + \tau_t^{maxR}) - \tau_t^{maxS}, & \text{if } \tau_i \notin \mathcal{E}_{\tau_t} \\ ((\phi_{t,i}^r + \tau_t^{maxR}) - \tau_t^{maxS}) \bmod \mathcal{T}_{\tau_i}^p, & \text{otherwise} \end{cases}$$

(19)

When τ_i belongs to $\mathcal{E}_{\tau_t}$, or $\tau_i \in \mathcal{E}_{\tau_t}$, the start phase is made positive by modulo operation to find the distance to the earliest future invocation of τ_i. Otherwise, The start phase $\phi_{t,i}^s$ can be negative as $\phi_{t,i}^r$.

Similarly, the finish phase $\phi_{t,i}^f$ is formulated based on $\phi_{t,i}^s$ and τ_t^{maxF} as follows:

$$\phi_{t,i}^f = \begin{cases} (\phi_{t,i}^s + \tau_t^{maxS}) - \tau_t^{maxF}, & \text{if } \tau_i \notin \mathcal{E}_{\tau_t} \text{ or } \tau_t^{proc} \in \mathcal{N} \\ ((\phi_{t,i}^s + \tau_t^{maxS}) - \tau_t^{maxF}) \bmod \mathcal{T}_{\tau_i}^p, & \text{otherwise} \end{cases}$$

(20)

Since there is no preemption during the execution of τ_t when $\tau_t^{proc} \in \mathcal{N}$, finish phase refers to the same invocation of the preempting task as the start phase. The finish phase $\phi_{t,i}^f$ is computed after τ_t^{maxF} computation and will be used for the request phases of successors. In the example of Fig. 10, $\phi_{t,i}^r$ and $\phi_{t,i}^s$ see the same invocation since the

difference $(\phi_{t,i}^r + \tau_t^{maxR}) - \tau_t^{maxS}$ is yet positive. On the other hand, $\phi_{t,i}^f$ sees the next invocation of τ_i, since the invocation of τ_i seen by $\phi_{t,i}^s$ appears during the execution of τ_t. In that case, the time difference from τ_t^{maxF} becomes negative and the modulo operation finds the positive distance to the next invocation.

LEMMA 5. *Phase adjustments $\phi_{t,i}^r$, $\phi_{t,i}^s$, and $\phi_{t,i}^f$ are conservative.*

PROOF. We prove it by induction. First, when task τ_t is a source task, $\phi_{t,i}^r$ is $-\Psi_{t,i}$ and it is conservative according to Lemma 4. When $\tau_i \notin \mathcal{E}_{\tau_t}$, $\phi_{t,i}^s$ and $\phi_{t,i}^f$ are computed directly from $\phi_{t,i}^r$ by changing the time reference from the release time to the start time (τ_t^{maxS}) and (τ_t^{maxF}) respectively in equations (19) and (20). So the consevativeness of the request phase is inherited to the start phase and the finish phase. When $\tau_i \in \mathcal{E}_{\tau_t}$, $\phi_{t,i}^s$ and $\phi_{t,i}^f$ are computed referring to the next release time of τ_i from τ_t^{maxS} and τ_t^{maxF}. Since the worst case preemption scenario to the preempted task is the periodic invocation of the preempting task after the worst case request phase, mod $\mathcal{T}_{\tau_i}^p$ operations in (19) and (20) find the closest request time of τ_i from τ_t^{maxS} and τ_t^{maxF} once the request phase is decided. It completes the initial step of the induction process.

Second, we prove the conservativeness of $\phi_{t,i}^r$, $\phi_{t,i}^s$, and $\phi_{t,i}^f$ of non-source task τ_t, assuming that for all $\tau_p \in pred(\tau_t)$, $\phi_{p,i}^f$s are conservative. If $\tau_t^{proc} \neq \tau_i^{proc}$, then $\phi_{t,i}^r$ is $-\Psi_{t,i}$ and it is conservative. Otherwise, we find the closest release time of τ_i from the finish phase $\phi_{p,i}^f$s of predecessors, by (18). Hence $\phi_{t,i}^r$ is conservative since $\phi_{p,i}^f$s are all conservative and the minimum value is chosen. The proof for $\phi_{t,i}^s$, and $\phi_{t,i}^f$ is similar to the case that task τ_t is a source task. Q.E.D. □

6. OPTIMIZATION TECHNIQUES

In this section, we describe two optimization techniques to tighten the time bounds by removing infeasible preemptions.

6.1. Exclusion Set Management

The exclusion technique manages for each task τ_i a set $\mathcal{EX}_{\tau_i}$ which includes tasks that are guaranteed to have no possibility of preempting τ_i. It is obvious that successor tasks belong to this set. If τ_i always preempts one of the predecessors of τ_s, τ_s cannot preempt τ_i since it will always be scheduled after τ_i. In addition, if τ_j is excluded by τ_i, then all $\tau_s \in \mathcal{EX}_{\tau_j}$ are also excluded by τ_i. In summary, the exclusion set $\mathcal{EX}_{\tau_i}$ becomes

$$\mathcal{EX}_{\tau_i} = \{\tau_s | \tau_s \in descendant(\tau_i) \text{ or } \tau_i \in \bigcup_{\tau_p \in ancestor(\tau_s)} (\mathcal{A}_{\tau_p} \cap \mathcal{F}_{\tau_p}) \text{ or } \tau_s \in \bigcup_{\tau_j \in \mathcal{EX}_{\tau_i}} \mathcal{EX}_{\tau_j}\}$$

(21)

where $ancestor(\tau_s)$ is a set of ancestors of τ_s and $descendant(\tau_i)$ is a set of descendants of τ_i. Since there is a cyclic dependency in (21), iterative computation is required for $\mathcal{EX}_{\tau_i}$, initially defined by $descendant(\tau_i)$. After time bound computation, it is updated using $\mathcal{A}_{\tau_i}$ and $\mathcal{F}_{\tau_i}$. Sets $\mathcal{A}_{\tau_t}$, $\mathcal{B}_{\tau_t}$, $\mathcal{D}_{\tau_t}$, $\mathcal{F}_{\tau_t}$ and $\mathcal{G}_{\tau_t}$ are modified to have an additional condition $\tau_s \notin \mathcal{EX}_{\tau_t}$. It is obvious that the exclusion technique does not affect the conservativeness of the proposed technique.

6.2. Duplicate Preemption Elimination

In our baseline technique, preemptions may occur redundantly; Fig. 11 (a) shows an elaborated example that experiences two types of duplicate preemptions. The first type of duplicate preemption may occur between tasks in the same task graph in case a

higher priority task has large release time variation. In the scheduling time bound analysis, we detect the preemption possibility by checking if a higher priority task can be released during task execution. In Fig. 11 (a), τ_5 can preempt both τ_4 and τ_8 because its release time varies between 20 and 75.

The second type of duplicate preemption occurs between tasks in different task graphs. In the phase adjustment technique, it is assumed that a preempting task preempts a predecessor task first. In the example of Fig. 11, τ_0 preempts τ_4 and phase adjustment is performed afterwards. The request phase of τ_7 to τ_0 is reset to $-\Psi_{7,0}$ since τ_7 and τ_0 are assigned to different processors, according to equation (18). Since the request phase of τ_8 is inherited from τ_7, τ_8 experiences another preemption by τ_0.

As a result, the WCRT is overestimated as illustrated in Fig. 11 (a) that contains both types of duplicate preemptions.

$$\tau_1^{pri} > \tau_2^{pri}, \tau_5^{pri} > \tau_0^{pri} > \tau_4^{pri} > \tau_8^{pri}, \tau_3^{pri} > \tau_6^{pri} > \tau_7^{pri}$$
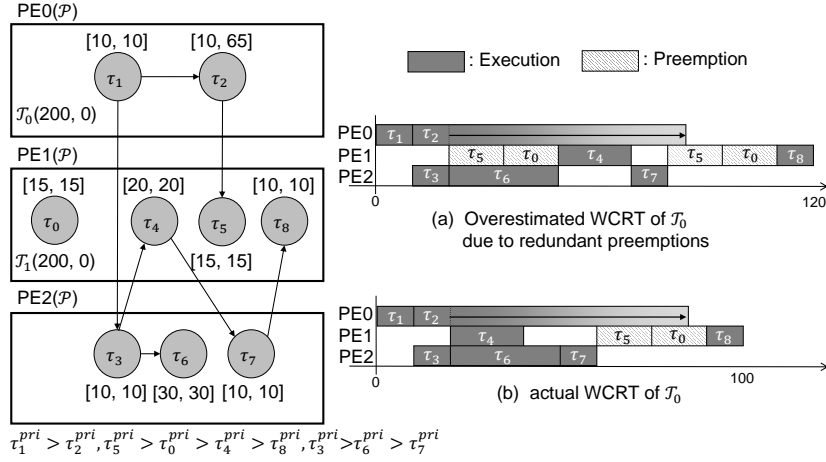
Fig. 11.   An example of common preemption elimination

To avoid duplicate preemptions, we devised an optimization that traces back the schedule and moves duplicate preemptions from ancestors to a target task. The proposed optimization heuristic is based on the abstruse fact that later preemption gives worse response time than earlier preemption when there are duplicate preemptions, which is stated in Theorem 2 below.

Algorithm 1 presents the psuedo code of the proposed heuristic. It is invoked by $RemoveDP(\tau_t)$ (lines 1-7) where τ_t is the task whose schedule time bound is computed. T_t^r represents how much the release time of τ_t is reduced by removing duplicate preemptions, and T_t^s represents how much preemption delay caused by other task graphs should be moved from predecessors to τ_t. We recursively trace back the schedule of critical path (lines 9 and 22) where τ_c and τ_{cri1} represent currently visited task and the parent task of τ_c on the critical path, respectively. We initialize T_t^r to zero and repeat $RecursiveRemoveDP$ until T_t^r is converged.

In algorithm 1, $PC_{t,i}$ means the preemption count from τ_i to τ_t. Boolean flag $detect_t$ checks if the release time of task τ_t cannot be reduced as much as the removed preemption time due to a predecessor mapped to a different PE. The boolean flag is inherited to the successors.

Let the currently visited task be τ_c in the recursive call, $RecursiveRemoveDP$. At first, we compute the amount of duplicate preemptions from tasks in the same task graph which can preempts both τ_c and τ_t (lines 13-14), where those tasks belong to set

---

ALGORITHM 1: Algorithm of duplicate preemption elimination

---

1 **Procedure** RemoveDP(τ_t)
2      $T_t^r = T_t^s = 0$
3   **repeat**
4          $T_t^s = 0$
5          $T_t^r = RecursiveRemoveDP(\tau_t, \tau_t, \phi)$
6   **until** T_t^s or T_t^r is changed
7 **Procedure** RecursiveRemoveDP($\tau_c, \tau_t, \mathcal{I}_{\tau_t}$)
8   find τ_{cri1} that satistifes $(\tau_{cri1} \in pred(\tau_c), \tau_{cri1}^{maxF} = \max_{\tau_p \in pred(\tau_c)} \tau_p^{maxF})$
9   find τ_{cri2} that satistifes $(\tau_{cri2} \in pred(\tau_c), \tau_{cri2}^{maxF} = \max_{\tau_p \in pred(\tau_c) - \{\tau_{cri1}\}} \tau_p^{maxF})$
10      $T_r = T_e = 0$
11   **if** $\tau_c \neq \tau_t$ **then**
12   T_e is increased by $\sum_{\tau_s \in \mathcal{J}_{\tau_c}} \tau_s^{WCET}$
13          $\mathcal{I}_{\tau_t} = \mathcal{I}_{\tau_t} \cup \mathcal{J}_{\tau_c}$
14   **if** $\exists_{\tau_p}(\tau_p \in pred(\tau_t), \tau_p^{maxF} = \tau_t^{maxR}, detect_p)$ **then**
15   T_e and T_t^s are increased by $\sum_{\tau_s \in \mathcal{K}_{\tau_c}} PC_{c,s} \cdot \tau_s^{WCET}$
16              $\mathcal{I}_{\tau_t} = \mathcal{I}_{\tau_t} \cup \mathcal{K}_{\tau_c}$
17   **end**
18   **end**
19   **if** τ_c is a source task **then**
20   **return** T_e
21   **end**
22      $T_r = RecursiveRemoveDP(\tau_{cri1}, \tau_t, \mathcal{I}_{\tau_t})$
23   **if** $\tau_c \neq \tau_t$ **then**
24          $T_r = \max\left(0, T_r - \sum_{\tau_s \in \mathcal{L}_{\tau_c}} \min(\tau_s^{WCET}, \tau_s^{maxF} - (\tau_c^{maxR} - T_r))\right)$
25   **end**
26   **if** τ_{cri2} is not found **then**
27   **return** $T_r + T_e$
28   **else**
29   **return** $\min(T_r, \tau_{cri1}^{maxF} - \tau_{cri2}^{maxF}) + T_e)$
30   **end**

$\mathcal{J}_{\tau_c} = \{\tau_s | \tau_s \in \mathcal{D}_{\tau_c} \cup \mathcal{G}_{\tau_c}, \tau_s \notin \mathcal{I}_{\tau_t}, \tau_s \notin \mathcal{EX}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}, \tau_t^{maxR} - T_t^r < \tau_s^{maxF}\}$. Even though preemption from the same task can be seen several times on the critical path, we need to consider only the recent preemption in this recursive function. Thus we manage task set $\mathcal{I}_{\tau_t}$ while traversing the critical path in order to consider only the recent preemption. If $detect_p$ value for τ_t is true, we also move the duplicate preemptions from tasks in the other task graphs to τ_t (lines 15-18), where those tasks belong to set $\mathcal{K}_{\tau_c} = \{\tau_s | \tau_s \notin \mathcal{T}_{\tau_t}, PC_{c,s} \neq 0, \tau_s \notin \mathcal{I}_{\tau_t}, \tau_s^{proc} = \tau_t^{proc}, \tau_s^{pri} > \tau_t^{pri}\}$. After computation of the amount of duplicate preemptions, $RecursiveRemoveDP$ is called recursively for the critical path predecessor τ_{cri1}, and we get the returned value as T_r (line 22), which is the possible release time reduction of τ_c. However, there can be additional preemptions if τ_c is released at $\tau_c^{maxR} - T_r$. We conservatively find the amount of additional preemptions and reduce T_r (line 23), where a task that incurs additional preemption belongs to set $\mathcal{L}_{\tau_c} = \{\tau_s | \tau_s \notin \mathcal{D}_{\tau_c} \cup \mathcal{G}_{\tau_c}, \tau_s \in \mathcal{T}_{\tau_c}, \tau_s \notin \mathcal{EX}_{\tau_c}, \tau_s \notin ancestor(\tau_c), \tau_s^{proc} = \tau_c^{proc}, \tau_s^{pri} > \tau_c^{pri}, \tau_s^{minS} \leq \tau_c^{maxR}, \tau_c^{maxR} - T_r < \tau_S^{maxF}\}$. For the non-preemptive scheduling policy, the condition $\tau_s^{pri} > \tau_c^{pri}$ is removed in $\mathcal{L}_{\tau_c}$. Note that release time reduction can be bounded by the other predecessors. If there is more than two predecessors of τ_c, we bound the reduced release time with the second largest

finish time among the predecessors (lines 25-26). Note that T_t^r is used in the formula of $\mathcal{J}_{\tau_c}$. Hence we set T_t^r to zero initially and repeat $RemoveDP$ until T_t^r is converged.

Refer to the example in Fig. 11. For task τ_8, we trace back the schedule of tasks τ_7, τ_4, τ_3, and τ_1 in order when we call $RemoveDP(\tau_8)$. When τ_4 is visited, We find out that $\tau_0 \in \mathcal{K}_{\tau_4}$ and $\tau_5 \in \mathcal{J}_{\tau_4}$, and both task can preempt τ_8. We remove those duplicate preemptions, then the maximum release time of τ_7 is reduced to $\tau_7^{maxR} - T_r = 40$. When returning to τ_7, We can know that $\tau_6 \in \mathcal{L}_{\tau_7}$ so that the start time of τ_7 cannot be earlier than the maximum finish time of τ_6, which is $\tau_7^{maxR} - (T_r - \tau_6^{maxF} - (\tau_7^{maxR} - T_r)) = 50$. Finally, we know that τ_8 can be released at $\tau_8^{maxR} - T_t^r = 60$ after removing all duplicate preemptions of its predecessors, which is the actual worst case.

The HPA equations need to be modified after Algorithm 1 is applied. We call $RemoveDP(\tau_t)$ after τ_t^{maxR} computation and set the reduced maximum release time $\hat{\tau}_t^{maxR}$ to $\tau_t^{maxR} - T_t^r$. Then τ_t^{maxR} is replaced by $\hat{\tau}_t^{maxR}$ in equations (9),(10), and sets $\mathcal{B}_{\tau_t}$ and $\mathcal{D}_{\tau_t}$. The terms $\tau_s^{maxF} - \tau_t^{maxR}$ and $\tau_s^{maxS} - \tau_t^{maxR} + 1 - \phi_{t,s}^r$ in the formula of $Delay_t^h$ is changed to $\tau_s^{maxF} - \hat{\tau}_t^{maxR}$ and $\max(0, \tau_s^{maxS} - \tau_t^{maxR}) + 1 - \phi_{t,s}^r$ respectively. And the request phase is not reset to the period shifting value in equation (18) when there is a predecessor mapped to different processors. The sets $\mathcal{B}_{\tau_t}$ and $\mathcal{D}_{\tau_t}$ have an additional condition $\tau_s \notin ancestor(\tau_t)$ since $\hat{\tau}_t^{maxR}$ can be smaller than the finish times of predecessors. τ_t^{maxS} is changed to have the initial value of $\hat{\tau}_t^{maxR}$ and is bounded by τ_t^{maxR}. And τ_t^{maxS} is changed to add T_t^s that is the sum of removed preemptions from the tasks in the other task graphs as follows:

$$\tau_t^{maxS} = \max(\tau_t^{maxR}, \hat{\tau}_t^{maxR} + T_t^s + Delay_t^l + Delay_t^h) \qquad (22)$$

THEOREM 2. *Algorithm 1 that removes the duplicate preemptions from the ancestors in the critical path preserves the conservativeness of the schedule time bound.*

PROOF. Refer to the electronic appendix for the proof. □

7. OVERALL HPA ALGORITHM

Now we ready to summarize the overall algorithm of the proposed technique. We compute the schedule time bounds of tasks and phase adjustment values until all time bounds are converged. The algorithm 2 shows the outermost iterative routine that integrates all computations.

ALGORITHM 2: HPA overall algorithm

---
1 **Procedure** HPA
2   period shifting $\Psi_{i,j} = \mathcal{T}_{\tau_j}^j$ for every task pair (τ_i, τ_j)
3   exclusion set $\mathcal{EX}_{\tau_i} = descendant(\tau_i)$ for every task τ_i
4   **repeat**
5   **foreach** *task τ_t in topological and priority descending order* **do**
6   compute τ_t^{minR}, τ_t^{minS}, and τ_t^{minF}
7   compute τ_t^{maxR}, $\hat{\tau}_t^{maxR}$, $\forall_{\tau_i \notin \mathcal{T}_{\tau_t}} \phi_{t,i}^r$, τ_t^{maxS}, $\forall_{\tau_i \notin \mathcal{T}_{\tau_t}} \phi_{t,i}^s$, τ_t^{maxF}, and $\forall_{\tau_i \notin \mathcal{T}_{\tau_t}} \phi_{t,i}^f$
8   **end**
9   update $\Psi_{i,j}$ for every task pair (τ_i, τ_j)
10   update $\mathcal{EX}_{\tau_i}$ for every task τ_i
11   **until** *any value is changed and $\forall_{\mathcal{T}}(\mathcal{R}_{\mathcal{T}} \le \mathcal{T}^d)$*

At first, period shifting and exclusion set are initialized (lines 2-3). Then we iteratively compute the time bounds of tasks (lines 5-8). Tasks are visited according to the

topological order first and priority descending order among independent tasks. Each time bound and phase is computed in the written order. After the time bound computation, period shifting and exclusion set set are updated (lines 9-10). This process is repeated until every value is converged. If there is a task graph that violates its deadline, we stop the iteration since it is not schedulable (line 11).

8. EXPERIMENTS

As the reference RTA technique, we implemented the Y&W method following the pseudo code in [Yen and Wolf 1998]. For comparison with MAST and SymTA/S, we use available tools; MAST suite [Harbour 2001] and pyCPA [Diemer and Axer 2012] that is a freely available compositional performance analysis tool similar to SymTA/S. The actual WCRT, which is denoted as Optimal, was obtained by an ILP-based approach [Kim et al. 2012]. The proposed HPA technique is available on-line [Choi et al. 2014].

Table I. WCRT estimation results for simple examples in Fig. 3, 4, 5, 6, and 11

	HPA	Y&W	MAST	pyCPA	Optimal
$\mathcal{T}_0$ in Fig. 3	40	35	40	50	40
$\mathcal{T}_0$ in Fig. 4	130	110	130	270	130
$\mathcal{T}_0$ in Fig. 5	140	150	140	300	140
$\mathcal{T}_0$ in Fig. 6	70	100	$\times$	210	70
$\mathcal{T}_0$ in Fig. 11	100	120	$\times$	310	100

Table I shows the comparison results for the examples shown in this paper. As discussed in Section 4, the Y&W method fails to find the WCRT for the examples of Fig. 3 and Fig. 4. MAST gives tight WCRT results for linear graphs, but no result for Fig. 6 and Fig. 11 since it supports only chain-structured graphs, and pyCPA provides highly overestimated results for all examples, since it pessimistically uses response time analysis ignoring the task dependency between processing elements. Note that the proposed HPA technique gives the optimal WCRT results for all these examples.

For extensive comparison, we generate graphs randomly; the number of task graphs varies from 3 to 5, the number of total tasks from 30 to 50, and the number of processing elements from 3 to 5. The τ^{BCET} and τ^{WCET} of each task are randomly selected in the range of [500, 1000] and [τ^{BCET}, $\tau^{BCET} \times 1.5$] respectively. The period and jitter of task graphs are randomly chosen but repaired to be schedulable if needed. Note that the problem size is too big to find optimal WCRTs with an ILP-based approach.

Table II. Comparison with the Y&W method and pyCPA for 100 random examples

	Win	Tie	Lose	Max%	Min(%)	Avg(%)
vs Y&W	193	200	1	39.42	-0.69	3.86
vs pyCPA	394	0	0	359.00	16.85	157.60

Table II shows the comparison results with the Y&W method and pyCPA. For the comparison with the Y&W method, we assume that all processing elements use the preemptive scheduling policy, and there is no jitter of input arrival. Since tasks may have multiple predecessors or successors, MAST is excluded in this comparison. The total number of task graphs is 394 in 100 randomly generated examples. The first column shows how many cases the HPA technique produces a tighter bound than the other methods. Similarly, the second and the third columns show how many cases HPA technique produces an equivalent bound and a looser bound. Columns Max, Min, and Avg indicate the maximum, minimum, and average WCRT estimation gaps between

HPA and the other approaches. HPA gives tighter bounds than the Y&W method for almost half of task graphs. A looser bound than the Y&W method was found for one task graph only with very small estimation gap, 0.69%. pyCPA provides highly over-estimated WCRT results for all task graphs, which are on average 2.57 times larger than the results of HPA.

Table III. Comparison with MAST and pyCPA for 100 random examples

	Win	Tie	Lose	Max%	Min(%)	Avg(%)
vs MAST	393	8	6	104.31	-4.48	18.95
vs pyCPA	406	1	0	295.16	0.00	44.83

For the comparison with MAST, we create another 100 random examples that are restricted to chain-structured graphs. In contrast to the previous experiment, random examples can have arbitrary mixture of preemptive and non-preemptive processing elements, and the jitter of input arrival is allowed. The number of total task graphs is 407. As shown in Table III, HPA shows remarkable performance advantage over the other methods. HPA provides tighter bounds than MAST and pyCPA in most cases.

Experiments are futher conducted to examine how the performance of the proposed HPA technique scales as the number of tasks, the number of graphs, and the execution time variation increases, compared with the other methods. The results are depicted in Fig. 12, 13, and 14, respectively. In these experiments, the default ranges of the number of processing elements, the number of task graphs, and the number of total tasks are set to [5,10], [5,7], and [30,50] respectively. For the left graphs in three figures, random examples are generated to have no jitter and only preemptive processing elements to compare with the Y&W method and pyCPA. For the right graphs, the task graphs in random examples are restricted to chain-structed graph but non-preemptive processing elements are allowed to compare with MAST and pyCPA. For all graphs in three figures, right y axis indicates the average WCRT estimation gap between HPA and pyCPA. For left(right) graphs, left y axis indicates the average WCRT estimation gap between HPA and the Y&W method(MAST). for each data point, the average value is obtained from 100 random examples.
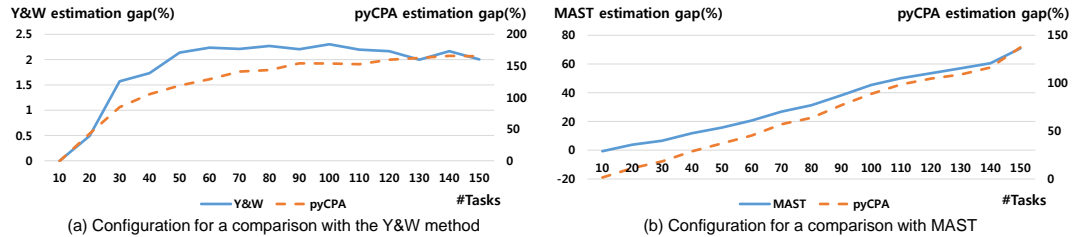

(a) Configuration for a comparison with the Y&W method

(b) Configuration for a comparison with MAST

Fig. 12.   WCRT estimation gap increases as the number of total tasks increases

Fig. 12 shows the change of estimation gap while varing the number of total tasks from 10 to 100. The estimation gap increases by the number of total tasks. For the right graph that uses chain-structed graph, the estimation gap is increased linearly, which confirms that HPA handles the effect of dependency effectively.

Fig. 13 shows the result of experiment that varies the number of task graph from 2 to 10 while fixing the number of total tasks to 100. In this experiment, the estimation

(a) Configuration for a comparison with the Y&W method

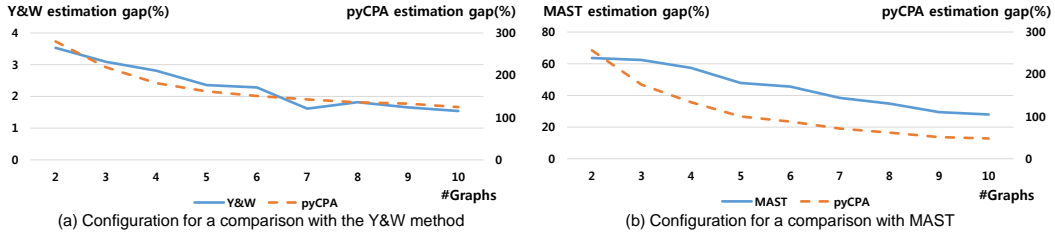(b) Configuration for a comparison with MAST

Fig. 13.   WCRT estimation gap decreases as the number of task graphs increases with the same number of tasks in total

gap decreases. It is because that as the task dependency decreases as well, the inter-application interference considered by the RTA analysis becomes dominant over intra-application interference that is computed by the schedule time bound analysis.

(a) Configuration for a comparison with the Y&W method

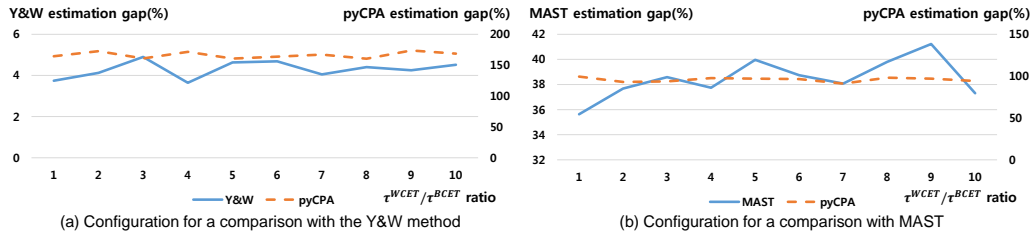(b) Configuration for a comparison with MAST

Fig. 14.   WCRT estimation gap change graphs with variation of $\frac{\tau^{WCET}}{\tau^{BCET}}$ ratio

The performance variation over the $\frac{\tau^{WCET}}{\tau^{BCET}}$ ratio from 1 to 10 is shown in Fig. 14. The figure shows that the execution time variation does not incurs any meaningful change in the performance gap.

(a) Small graph configuration
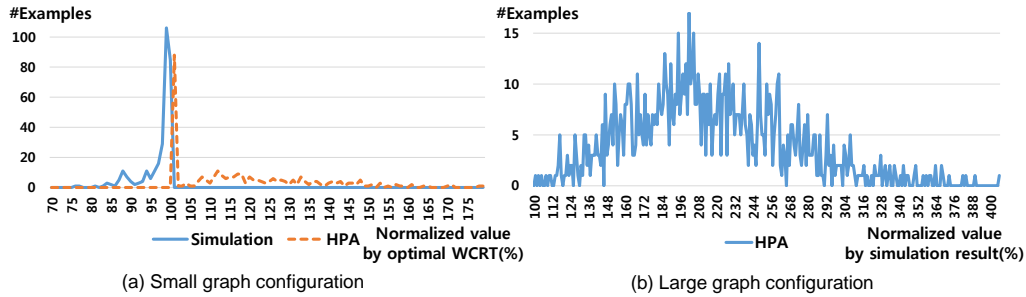
(b) Large graph configuration

Fig. 15.   Distribution graphs of the estimated WCRT from the HPA

Fig. 15 shows the distribution of the estimated WCRTs from HPA with 300 randomly generated examples that have arbitrary mixture of preemptive and non-preemptive scheduling policies and no restriction on the graph topology. Fig. 15 (a) shows the distribution with small size examples and Fig. 15 (b) with large size examples. For small(large) size examples, the ranges of the number of processing elements, the number of task graphs, and the number of total tasks are set to [2,3]([5,10]), [1,3]([5,7]), and [10,15]([60,100]), respectively. We also performed Monte-Carlo simulation to obtain the

WCRT empirically by sampling graph instances 200 million times for each example. Optimal WCRTs are found with an ILP-based approach only for small size examples. Fig. 15 (a) depicts the distribution of the WCRTs obtained from Monte-Carlo simulation and HPA, normalized by optimal WCRTs. The x axis indicates the normalized value (%) and the y axis presents the number of examples. Note that both Monte-Carlo simulation and HPA provide results close to the actual WCRT in majority cases. But Monte-Carlo simulation may not find the true WCRT despite 200 million times of simulation, which confirms the need of conservative estimation techniques. For large size examples, the distribution results from HPA are normalized by the near-WCRTs obtained from Monte-Carlo simulation in Fig. 15 (b) since optimal WCRTs cannot be obtained. Since Monte-Carlo simulation gives underestimated WCRTs, the amount of overestimation might be quite exaggerated. It shows that the estimated WCRT from HPA may have about 100% overestimation on average.
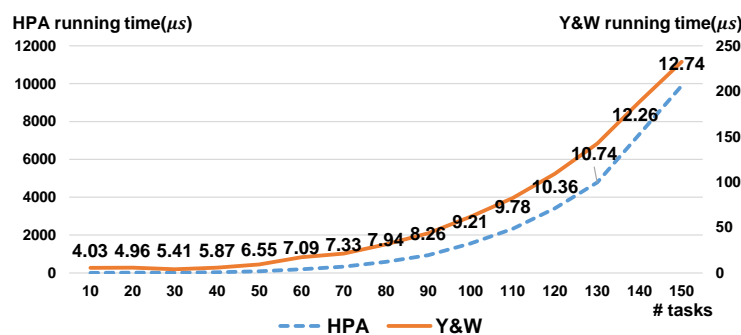

Fig. 16.   Average running time graph with average number of iteration

Finally, we measure the running time of HPA on 3.4 GHz i7 machine with 8GB main memory and the iteration number of convergence to verify the scalability of the proposed technique. We vary the number of total tasks from 10 to 150 and use 100 randomly generated examples for each number of total tasks. The average values are plotted in Fig. 16, where the left y-axis and the right y-axis indicate the running time of HPA and Y&W in micro seconds, respectively. The numbers labeled to the graph of the HPA indicate the average numbers of iterations for convergence. Every example was converged within maximum 20 iterations. Although HPA is slower than the Y&W method by an order of magnitude, both methods show similar scalability as shown in Fig. 16.

9. CONCLUSION

In this paper, we addressed a very challenging problem that is to tightly estimate the worst case response time of an application in a distributed embedded system. It is shown that a state-of-the-art technique, Y&W method, fails to find a conservative WCRT bound. Thus we propose a hybrid performance analysis (HPA) method that combines the scheduling time bound analysis and the response time analysis to consider inter-task interference between different tasks. It finds a conservative and tight WCRT bound, considering task dependency, execution time variation, arbitrary mixture of fixed-priority preemptive and non-preemptive processing elements, and input jitters. Experimental results show that it produces tighter bounds than the Y&W method, MAST, and pyCPA. Convergence and scalability of the proposed technique are confirmed empirically. Since it is an iterative technique with tens of non-linear formulas,

it is a very difficult problem to prove the convergence formally, which is left as a future work.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

REFERENCES

Aske Brekling, Michael R. Hansen, and Jan Madsen. 2008. Models and formal verification of multiprocessor system-on-chips. *J. Logic Algebraic Program.* 77(1-2). 119.

Jinwoo Kim, Hyunok Oh, Hyojin Ha, Shin-Haeng Kang, Junchul Choi, and Soonhoi Ha. 2012. An ILP-Based Worst-Case Performance Analysis Technique for Distributed Real-Time Embedded Systems. *Proceedings of the 2012 IEEE 33rd Real-Time Systems Symposium(RTSS'12)*. 363-372.

Ti-yen Yen and Wayne Wolf. 1998. Performance estimation for real-time distributed embedded systems. *Parallel and Distributed Systems, IEEE Transactions on*. 9(11). 11251136.

John Lehoczky, Lui Sha, and Ye Ding. 1989. The Rate Monotonic Scheduling Algorithm: Exact Characterization and Average Case Behavior. *Proceedings of 1989 IEEE Real-Time Systems Symposium*. 166-171.

Jinwoo Kim, Hyunok Oh, Junchul Choi, Hyojin Ha, and Soonhoi Ha. 2013. A Novel Analytical Method for Worst Case Response Time Estimation of Distributed Embedded Systems. *Proceedings of the The 50th Annual Design Automation Conference on Design Automation Conference(DAC'13)*. 1-10.

John Lehoczky. 1990. Fixed Priority Scheduling of Periodic Task Sets with Arbitrary Deadlines. *Proceedings of 11th IEEE Real-Time Systems Symposium(RTSS'90)*. 201-209.

Neil Audsley, Alan Burns, Ken Tindell, M. Richardson and Andrew J. Wellings. 1993. Applying New Scheduling Theory to Static Priority Pre-emptive Scheduling. *Software Engineering Journal*. 8(5). 284-292.

Ken Tindell and John Clark. 1994. Holistic Schedulability Analysis for Distributed Hard Real-Time Systems. *Microprocessing and Microprogramming - Parallel processing in embedded real-time systems*. 40(2-3). 117-134.

Jos Carlos Palencia and Medina Gonzlez Harbour. 1998. Schedulability Analysis for Tasks with Static and Dynamic Offsets. *Proceedings of the 1998 IEEE 19th Real-Time Systems Symposium(RTSS'98)*. 26-37.

Ken Tindell, Alan Burns, and Andrew J. Wellings. 1995. Analysis of Hard Real-Time Communications. *Real-Time Systems*. 9(2). 147-171.

Simon Schliecker and Rolf Ernst. 2010. Real-Time Performance Analysis of Multiprocessor Systems with Shared Memory. *ACM Transactions on Embedded Computing Systems*. 10(2).

Rodolfo Pellizzoni and Giuseppe Lipari. 2007. Holistic analysis of asynchronous real-time transactions with earliest deadline scheduling. *Journal of Computer and System Sciences*. 73(2). 186-206.

Ti-yen Yen and Wayne Wolf. 1995. Communication Synthesis for Distributed Embedded System. *Proceedings of IEEE/ACM International Conference on Computer-Aided Design(ICCAD'95)*. 288-294.

Paul Pop, Petru Eles, and Zebo Peng. 2000. Schedulability Analysis for Systems with Data and Control Dependencies. *Proceedings of the 12th Euromicro Conference on Real-Time Systems(ECRTS'00)*. 201-208.

Rafik Henia, Arne Hamann, Marek Jersak, Razvan Racu, Kai Richter, and Rolf Ernst. 2005. System Level Performance Analysis - The SymTA/S Approach. *IEE Proceedings - Computers and Digital Techniques*. 152(2). 148-166.

Michael Gonzlez Harbour. 2001. MAST suite. Retrieved from http://mast.unican.es/

Jonas Diemer and Philip Axer. 2012. Python implementation of Compositional Performance Analysis. Retrieved from http://pycpa.readthedocs.org/en/latest/

Junchul Choi, Hyunok Oh, and Soonhoi Ha. 2014. SWAT: A Suite of Worst Case Response Time Analysis Tools for Distributed Hard Real-Time Systems. Retrivable from http://iris.snu.ac.kr/xe/swat

Online Appendix to:
A Hybrid Performance Analysis Technique for Distributed Real-Time Embedded Systems

JUNCHUL CHOI, Seoul National University
HYUNOK OH, Hanyang University
SOONHOI HA, Seoul National University

A. PROOFS FOR OPTIMIZATION TECHNIQUES

We prove that the optimization technique of duplicate preemption elimination preserves the conservativeness of the HPA technique. It is a rather long proof. We will make several definitions, lemmas, and theorems in this section.

DEFINITION A.1. $pe(\tau_t)[x,y]$ *is defined as the sum of execution time of tasks of which priority is higher than τ_t from time x to time y. Formally,*

$$pe(\tau_t)[x,y] = \sum_{\tau_s \in S_{pe(\tau_t)[x,y]}} \min(y, \tau_s^{finish}) - \max(x, \tau_s^{start})$$

where $S_{pe(\tau_t)[x,y]} = \{\tau_s | \tau_s^{pri} > \tau_t^{pri}, (x \le \tau_s^{start} < y \text{ or } x < \tau_s^{finish} \le y)\}$ for preemptive scheduling and $S_{pe(\tau_t)[x,y]} = \{\tau_s | (\tau_s^{pri} > \tau_t^{pri} \text{ or } \tau_s^{start} < \tau_t^{start}), (x \le \tau_s^{start} < y \text{ or } x < \tau_s^{finish} \le y)\}$ for nonpreemptive scheduling.

The following two lemmas hold and their proofs are trivial.

LEMMA A.1. $0 \le pe(\tau_t)[x,y] \le y - x$.

LEMMA A.2. $pe(\tau_t)[x,z] = pe(\tau_t)[x,y] + pe(\tau_t)[y,z]$.

THEOREM A.1. *for all τ_t, if $\tau_t^{release} \le \tau_t^{maxR} - \Delta$ then it always holds that $\tau_t^{finish} \le \tau_t^{maxF} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}]$.*

PROOF. We prove it by contradiction. Suppose that τ_t^{finish} exists such that $\tau_t^{finish} > \tau_t^{maxF} - \Delta + pe(\tau_t)[\tau_t^{release}, \tau_t^{finish}]$ where $\tau_t^{release} \le \tau_t^{maxR} - \Delta$. It is obvious that $\tau_t^{maxF} \ge \tau_t^{maxR} + \tau_t^{WCET} + pe(\tau_t)[\tau_t^{maxR}, \tau_t^{maxF}]$ and $\tau_t^{finish} \le \tau_t^{release} + \tau_t^{WCET} + pe(\tau_t)[\tau_t^{release}, \tau_t^{finish}]$, by Definition A.1.

$$\tau_t^{finish} > \tau_t^{maxF} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}]$$

$$\ge \tau_t^{maxR} + \tau_t^{WCET} + pe(\tau_t)[\tau_t^{maxR}, \tau_t^{maxF}] - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}]$$

$$= \tau_t^{maxR} + \tau_t^{WCET} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxF}].$$

Thus,

$$\tau_t^{release} + \tau_t^{WCET} + pe(\tau_t)[\tau_t^{release}, \tau_t^{finish}] > \tau_t^{maxR} + \tau_t^{WCET} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxF}].$$

$$\tau_t^{release} + pe(\tau_t)[\tau_t^{release}, \tau_t^{maxR} - \Delta] + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{finish}]$$

$$> \tau_t^{maxR} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxF}].$$

$\tau_t^{release} + pe(\tau_t)[\tau_t^{release}, \tau_t^{maxR} - \Delta] > \tau_t^{maxR} - \Delta$ since $pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxF}] \geq$ $pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{finish}]$.

$\tau_t^{maxR} - \Delta < \tau_t^{release} + pe(\tau_t)[\tau_t^{release}, \tau_t^{maxR} - \Delta] \leq \tau_t^{maxR} - \Delta$ since $pe(\tau_t)[\tau_t^{release}, \tau_t^{maxR} - \Delta] \leq \tau_t^{maxR} - \Delta - \tau_t^{release}$.

It is a contradiction. Hence $\tau_t^{finish} \leq \tau_t^{maxF} - \Delta + pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}]$. **Q.E.D.** □

Theorem A.1 explains how much the maximum finish time can be reduced if the maximum release time decreases. Since the release time depends on the finish times of predecessor tasks, the theorem presents the effect of the finish time of predecessor tasks onto the finish time of the target task. This theorem will be generalized later in theorem A.2.

Hereafter, $pe(\tau_t, \Delta)$ denotes $pe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}]$ for brevity.

It is trivial that if a predecessor task always finishes early then the target task is released early. Therefore, for $\tau_s \in pred(\tau_t)$, if $\forall_{\tau_s^{finish}}(\tau_s^{finish} \leq \tau_s^{maxF} - \Delta_{\tau_s})$ then $\tau_t^{release} \leq \max_{\tau_s \in pred(\tau_t)}(\tau_s^{maxF} - \Delta_{\tau_s})$.

DEFINITION A.2. *$ancestor(\tau_t)$ is a set of ancestors of τ_t, formally $ancestor(\tau_t) = \bigcup_{i \geq 1} pred^i(\tau_t)$, where $pred^i(\tau_t) = \bigcup_{\tau_s \in pred^{i-1}(\tau_t)} pred(\tau_s)$ an d $pred^1(\tau_t) = pred(\tau_t)$.*

DEFINITION A.3. *$succ(\tau_t)$ is the immediate successor set of task τ_t, formally $succ(\tau_t) = \{\tau_c | \tau_t \in pred(\tau_c)\}$.*

DEFINITION A.4. *$descendant(\tau_t)$ is a set of descendants of τ_t, formally $descendant(\tau_t) = \bigcup_{i \geq 1} succ^i(\tau_t)$, where $succ^i(\tau_t) = \bigcup_{\tau_s \in succ^{i-1}(\tau_t)} succ(\tau_s)$ and $succ^1(\tau_t) = succ(\tau_t)$.*

DEFINITION A.5. *$path(\tau_s, \tau_d) = descendant(\tau_s) \bigcap ancestor(\tau_d)$.*

Now we will examine the effect of the early finish time of an ancestor task to the finish time of the target task. In theorem A.1, we examine the relation between the release time and the finish time which also explains the finish times between predecessor tasks and the target task. We will extend it to the effect between an ancestor task and the target task. If we apply theorem A.1 repeatedly, we can compute how much the maximum finish time reduction is inherited from an ancestor task to the target task.

THEOREM A.2. *For $\tau_a \in ancestor(\tau_t)$, if $\forall_{\tau_a^{finish}}(\tau_a^{finish} \leq \tau_a^{maxF} - \Delta_{\tau_a} + pe(\tau_a, \Delta_{\tau_a}))$ then always $\tau_t^{finish} \leq \tau_t^{maxF} - \Delta_{\tau_t} + pe(\tau_t, \Delta_{\tau_t})$, where $\Delta_{\tau_m} = \max_{\tau_i \in pred(\tau_m)}(\tau_i^{maxF}) - \max_{\tau_i \in pred(\tau_m)}(\tau_i^{maxF} - \Delta_{\tau_i} + pe(\tau_i, \Delta_{\tau_i}))$. Note that $\Delta_{\tau_m} = 0$ if there is no predecessor.*

PROOF. We will prove it by induction. First, $\tau_a^{finish} \leq \tau_a^{maxF} - \Delta_{\tau_a} + pe(\tau_a, \Delta_{\tau_a})$. Assume that $\tau_i^{finish} \leq \tau_i^{maxF} - \Delta_{\tau_i} + pe(\tau_i, \Delta_{\tau_i})$. For $\tau_m \in succ(\tau_i)$, $\tau_m^{release} = \max_{\tau_i \in pred(\tau_m)} \tau_i^{finish} \leq \max_{\tau_i \in pred(\tau_m)} \tau_i^{maxF} - \Delta_{\tau_i} + pe(\tau_i, \Delta_{\tau_i})$. Let $\Delta_{\tau_m} = \max_{\tau_i \in pred(\tau_m)}(\tau_i^{maxF}) - \max_{\tau_i \in pred(\tau_m)}(\tau_i^{maxF} - \Delta_{\tau_i} + pe(\tau_i, \Delta_{\tau_i}))$.

Since $\tau_m^{maxR} = \max_{\tau_i \in pred(\tau_m)} \tau_i^{maxF}$,

$$\tau_m^{release} = \max_{\tau_i \in pred(\tau_m)} \tau_i^{finish} \leq \max_{\tau_i \in pred(\tau_m)} \tau_i^{maxF} - \Delta_{\tau_i} + pe(\tau_i, \Delta_{\tau_i})$$

$$= \max_{\tau_i \in pred(\tau_m)} (\tau_i^{maxF}) - \Delta_{\tau_m} = \tau_m^{maxR} - \Delta_{\tau_m}.$$

By theorem A.1, $\tau_m^{finish} \leq \tau_m^{maxF} - \Delta_{\tau_m} + pe(\tau_m, \Delta_{\tau_m})$ since $\tau_m^{release} \leq \tau_m^{maxR} - \Delta_{\tau_m}$. Q.E.D. □

Theorem A.2 tells us that the contribution of the maximum finish time reduction of an ancestor task diminishes as it propagates to the child tasks. To utilize theorem A.2, we define the new reduced finish time $\hat{\tau}_t^{maxF}$ by reducing the finish time of a task τ_a by Δ as following:

DEFINITION A.6. $\hat{\tau}_{t,\Delta_{\tau_a} \leftarrow \Delta}^{maxF} = \tau_t^{maxF} - \Delta_{\tau_t} + pe(\tau_t, \Delta_{\tau_t})$, where $\Delta_{\tau_m} = \max_{\tau_i \in pred(\tau_m)} \tau_i^{maxF} - \max_{\tau_i \in pred(\tau_m)} \hat{\tau}_{i,\Delta_{\tau_a} \leftarrow \Delta}^{maxF}$ if Δ_{τ_a} is Δ.

By definition A.6 and theorem A.2, $\tau_t^{finish} \leq \hat{\tau}_t^{maxF}$ if $\forall \tau_a^{finish}, \tau_a^{finish} \leq \hat{\tau}_a^{maxF}$. To avoid the duplicate preemptions, we define a common preemptor set as $cp(\tau_a, \tau_t)$.

DEFINITION A.7. For $\tau_a \in ancestor(\tau_t)$, $cp(\tau_a, \tau_t) = \{\tau_s | \tau_s^{pri} > \max(\tau_a^{pri}, \tau_t^{pri}), \tau_s^{maxF} \geq \tau_t^{minR}, \tau_s^{minR} \leq \tau_a^{maxF}\}$.

If a preemption task can preempt both τ_a and τ_t then τ_t^{finish} is larger when it preempts τ_t rather than τ_a.

THEOREM A.3. *If a common preemptor τ_p can preempt ancestor task τ_a and target task τ_t completely, then always it produces no smaller finish time of τ_t when τ_p preempts τ_t than τ_a.*

PROOF. The reduced maximum finish time of τ_t is $\hat{\tau}_{t,\Delta_{\tau_a} \leftarrow \tau_p^{WCET}}^{maxF}$ when τ_p does not preempt τ_a but preempts τ_t. It is $\hat{\tau}_{t,\Delta_{\tau_t} \leftarrow \tau_p^{WCET}}^{maxF}$ when τ_p does not preempt τ_t but preempts τ_a. Since τ_a is an ancestor of τ_t, $\hat{\tau}_{t,\Delta_{\tau_a} \leftarrow \tau_p^{WCET}}^{maxF} \geq \hat{\tau}_{t,\Delta_{\tau_t} \leftarrow \tau_p^{WCET}}^{maxF}$. Hence, it produces no smaller finish time to preempt τ_t than τ_a. Q.E.D. □

Let us consider a case that a preemption task τ_p can preempt τ_a and τ_t, but τ_t partially. Even in this case, τ_p can always preempt τ_a fully since $\tau_p^{finish} \geq \tau_t^{release} \geq \tau_a^{finish}$. Then we may need to compare $\hat{\tau}_{t,\Delta_{\tau_a} \leftarrow \tau_p^{WCET}}^{maxF}$ with $\hat{\tau}_{t,\Delta_{\tau_t} \leftarrow maxP(\tau_t, \tau_p)}^{maxF}$ where $maxP(\tau_t, \tau_p)$ indicates the maximum partial preemption time of τ_t by τ_p. Even in this case, surprisingly, preempting τ_t always provides larger finish time than the partial preemption case, which is stated by theorem A.4.

THEOREM A.4. *If a common preemptor τ_p can preempt an ancestor task τ_a fully and the target task τ_t partially, it always produces no smaller finish time bound of τ_t when τ_p preempts τ_t rather than τ_a.*

PROOF. Assume that δ indicates the smallest partial preemption time between τ_p and τ_t. So, $\delta = \tau_p^{maxF} - \tau_t^{maxR} \leq \tau_p^{WCET}$.

(1) If τ_p preempts τ_a then τ_t finishes earlier than $\tau_t^{maxF} - \delta$.
(2) If τ_p preempts τ_t then assume that $\tau_t^{release} \leq \tau_t^{maxR} - \epsilon$. Then $\tau_t^{finish} \leq \tau_t^{maxF} - \epsilon + pe(\tau_t, \epsilon)$.
   (a) If $\epsilon \leq \tau_p^{WCET} - \delta$ then $pe(\tau_t, \epsilon) = \epsilon$, and $\tau_t^{maxF} - \epsilon + pe(\tau_t, \epsilon) = \tau_t^{maxF}$.
   (b) If $\tau_p^{WCET} - \delta < \epsilon$ then $pe(\tau_t, \epsilon) \geq \tau_p^{WCET} - \delta$. Since $\epsilon \leq \tau_p^{WCET}$ by theorem A.2, $\tau_t^{maxF} - \epsilon + pe(\tau_t, \epsilon) \geq \tau_t^{maxF} - \tau_p^{WCET} + \tau_p^{WCET} - \delta = \tau_t^{maxF} - \delta$.

Hence, in both cases 2a and 2b, the preemption of τ_t provides no smaller maximum finish time of τ_t. □

Now, we consider multiple ancestor tasks which have multiple common preemption tasks. We define a unique common preemptor set, $ucp(\tau_a, \tau_p)$, to indicate the closest ancestor with the common preemption task as following.

DEFINITION A.8.
$$ucp(\tau_a, \tau_t) = cp(\tau_a, \tau_t) - \bigcup_{\tau_m \in path(\tau_a, \tau_t)} cp(\tau_m, \tau_t).$$

Then multiple ancestors may have different common preemption tasks. To consider the cascaded scheduling effect as a whole, we redefine Δ_{τ_a}. We introduce δ_{τ_a} which denotes the reduced time by moving common preemption tasks from τ_a to τ_t. δ_{τ_a} is statically computed. Δ_{τ_a} indicates a variable used in definition A.6. The finish time of τ_m is contributed by the early finish time of ancestor tasks of τ_m and δ_{τ_m}. The total Δ_m is defined as following:

DEFINITION A.9. $\Delta_{\tau_m} = \max_{\tau_i \in pred(\tau_m)} \tau_i^{maxF} - \max_{\tau_i \in pred(\tau_m)} \hat{\tau}_i^{maxF} + \delta_{\tau_m}$, *where* $\delta_{\tau_m} = \sum_{\tau_s \in ucp(\tau_m, \tau_t)} \min(\tau_s^{WCET}, \tau_s^{maxF} - \tau_t^{maxR})$.

Finally, we define $pe(\tau_t, \Delta)$ which indicates the occupied time from $\tau_t^{maxR} - \Delta$ to τ_t^{maxR}. Although $pe(\tau_t, \Delta)$ shows the exact time, it is varying at run time. Therefore, we need to compute its bound. Although we use its bound, the previous theorems hold. Since $pe(\tau_t, \Delta) \leq \min(\Delta, maxpe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}])$ where $maxpe(\tau_t)[x, y] = \sum_{\tau_s^{pri} > \tau_t^{pri}, \tau_s^{maxF} \geq x, \tau_s^{minS} \leq y} \min(\tau_s^{WCET}, y - \tau_s^{minS}, \tau_s^{maxF} - x)$, $\min(\Delta, maxpe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}])$ is used as the bound of $pe(\tau_t, \Delta)$.

DEFINITION A.10. $pe(\tau_t, \Delta) = \min(\Delta, maxpe(\tau_t)[\tau_t^{maxR} - \Delta, \tau_t^{maxR}])$, *where* $maxpe(\tau_t)[x, y] = \sum_{\tau_s^{pri} > \tau_t^{pri}, \tau_s^{maxF} \geq x, \tau_s^{minS} \leq y} \min(\tau_s^{WCET}, y - \tau_s^{minS}, \tau_s^{maxF} - x)$.

THEOREM A.5. *Algorithm 1 that removes the duplicate preemptions from the ancestors in the critical path preserves the conservativeness of the schedule time bound.*

PROOF. By replacing variables in definition A.6 by definitions A.9 and A.10, we can compute the tighter time bound with moving the common preemption tasks from the ancestors. The conservativeness is guaranteed by Theorems A.3 and A.4 since preempting τ_t rather than ancestor τ_a produces a larger time bound. The difference between the proof and the actual technique is that *duplicate preemption elimination* technique traverses only the critical path, not all ancestors. Even though common preemptions of the ancestors not in the critical path remain, our technique is still conservative. Q.E.D. □