Partitioned Multiprocessor Fixed-Priority Scheduling of Sporadic Real-Time Tasks

Jian-Jia Chen

[1]Department of Informatics, TU Dortmund University, Germany
E-mail: jian-jia.chen@cs.uni-dortmund.de

Abstract—**Partitioned multiprocessor scheduling has been widely accepted in academia and industry to statically assign and partition real-time tasks onto identical multiprocessor systems. This paper studies fixed-priority partitioned multiprocessor scheduling for sporadic real-time systems, in which deadline-monotonic scheduling is applied on each processor. Prior to this paper, the best known results are by Fisher, Baruah, and Baker with speedup factors $4 - \frac{2}{M}$ and $3 - \frac{1}{M}$ for arbitrary-deadline and constrained-deadline sporadic real-time task systems, respectively, where M is the number of processors. We show that a greedy mapping strategy has a speedup factor $3 - \frac{1}{M}$ when considering task systems with arbitrary deadlines. Such a factor holds for polynomial-time schedulability tests and exponential-time (exact) schedulability tests. Moreover, we also improve the speedup factor to 2.84306 when considering constrained-deadline task systems. We also provide tight examples when the fitting strategy in the mapping stage is arbitrary and M is sufficiently large. For both constrained- and arbitrary-deadline task systems, the analytical result surprisingly shows that using exact tests does not gain theoretical benefits (with respect to speedup factors) if the speedup factor analysis is oblivious of the particular fitting strategy used.**

Keywords: Sporadic real-time tasks, resource augmentation, approximation, schedulability analysis.

1 Introduction

The sporadic task model has been widely adopted as the basic model for real-time systems with recurring executions [31]. A sporadic real-time task τ_i is characterized by its *minimum inter-arrival time* T_i, its timing constraint or *relative deadline* D_i, and its (worst-case) *execution time* C_i. A sporadic task defines an infinite sequence of task instances, also called *jobs*, that arrive with the minimum inter-arrival time constraint. Under the minimum inter-arrival time constraint, any two consecutive jobs of task τ_i should be temporally separated by at least T_i. When a job of task τ_i arrives at time t, the job should finish no later than its *absolute deadline* $t + D_i$. If we consider a task releases its jobs periodically, such a task model is the well-known Liu and Layland task model [30], where T_i is the *period* of the task. An input task set is said to have 1) *implicit deadlines* if the relative deadlines of sporadic tasks are equal to their minimum inter-arrival times, 2) *constrained deadlines* if the minimum inter-arrival times are no less than their relative deadlines, and (3) *arbitrary deadlines*, otherwise.

Through this paper, we only consider implicitly preemptive scheduling. That is, a job may be preempted by another job on a processor. For scheduling sporadic tasks on a processor, the preemptive earliest-deadline-first (EDF) policy is optimal [30] to meet the timing constraints. However, EDF requires to prioritize the jobs in the ready queue by using their absolute deadlines, and the overhead is in general not negligible. The industrial practice is to use fixed-priority scheduling, also supported in most real-time operating systems, in which a task is assigned with a fixed priority level. The seminal work by Liu and Layland [30] shows that rate monotonic (RM) scheduling is optimal for uniprocessor fixed-priority scheduling when considering implicit-deadline task systems. Moreover, deadline monotonic (DM) scheduling [29] is optimal for uniprocessor fixed-priority scheduling for constrained-deadline task systems. For arbitrary-deadline task systems, Audsley et al. [2] provide an optimal priority assignment algorithm to define the priority levels of the sporadic tasks for uniprocessor fixed-priority scheduling.

Testing whether a task set can be feasibly scheduled by a scheduling algorithm is called a *schedulability test*. Even though RM and DM are known to be optimal for uniprocessor fixed-priority scheduling, the exact schedulability tests for uniprocessor fixed-priority scheduling requires pseudo-polynomial time by using the exact tests by Lehoczky, Sha, and Ding [28] for constrained-deadline systems and in exponential time by Lehoczky [27] for arbitrary-deadline systems. Specifically, computing the worst-case response time of one (lowest-priority) task is shown $\mathcal{NP}$-hard by Eisenbrand and Rothvoß [20]. There have been extensive results about testing the schedulability of uniprocessor fixed-priority scheduling. The more efficient strategy is to provide only sufficient conditions that can be verified in polynomial-time, like the utilization bound [10], [26], [30], the quadratic utilization bound [7], the hyperbolic utilization bound [8], [16], the approximated request bound functions [1], [11], [21].

To quantify the performance loss due to efficient schedulability tests and assigning tasks with fixed priority levels, we will adopt the notion of speedup factors, (also known as resource augmentation factors). A fixed-priority scheduling algorithm with a speedup factor ρ guarantees to produce feasible schedules by running (each processor) ρ times as fast as in the original platform (speed), if there exists a feasible schedule (under arbitrary policies) for the task system. The speedup factors of DM scheduling, with respect to the optimal uniprocessor EDF scheduling, are $\frac{1}{\ln 2}$, 1.76322, and 2 for implicit-deadline, constrained-deadline, and arbitrary-deadline task sets [17], [19], respectively.

To schedule real-time tasks on multiprocessor platforms, there have been three widely adopted paradigms: partitioned, global, and semi-partitioned scheduling. A comprehensive survey of multiprocessor scheduling in real-time systems can be found in [18]. In this paper, we consider *partitioned scheduling*, in which the tasks are statically partitioned onto processors and all the processors are identical. That is, all the

	implicit deadlines	constrained deadlines	arbitrary deadlines
partitioned EDF	$\frac{4}{3} - \frac{1}{3M}$ [23] PTAS [24]	$3 - \frac{1}{M}$ [5] $2.6322 - \frac{1}{M}$ [13]	$4 - \frac{2}{M}$ [4] $3 - \frac{1}{M}$ [13]
partitioned DM	$\frac{7}{4}$ [10] 1.5 [25]	$3 - \frac{1}{M}$ [22] 2.84306 (*this paper*)	$4 - \frac{2}{M}$ [22] $3 - \frac{1}{M}$ (*this paper*)

TABLE I: Summary of the speedup factors in the multiprocessor partitioned scheduling problem for sporadic task systems. For more details of implicit-deadline cases, please refer to Table III in the survey [18].

jobs of a task are executed on a specific processor with fixed-priority scheduling.

However, problems on multiprocessors become $\mathcal{NP}$-complete (or worse) in the strong sense even in the simplest possible cases. For example, deciding if an implicit-deadline task set with the same period is schedulable on multiple processors is already $\mathcal{NP}$-complete in the strong sense [31]. To cope with these $\mathcal{NP}$-hardness issues, one natural approach is to focus on approximation algorithms, i.e., polynomial time algorithms that produce an approximate solution instead of an exact one. There have been many results for implicit-deadline task systems, as summarized in the survey paper [18]. But, only a few results are known for constrained-deadline and arbitrary-deadline task systems.

When considering sporadic task sets with constrained or arbitrary deadlines, the problem becomes more complicated, when EDF or fixed-priority scheduling is adopted on a processor. The recent studies in [3], [14] provide polynomial-time approximation schemes for some special cases when speeding-up is adopted for EDF scheduling. For general cases, Baruah and Fisher [4], [5] propose a simple method, denoted as *deadline-monotonic partitioning* in this paper, which 1) considers the tasks in a non-decreasing order of their relative deadlines, and 2) assigns a task (in the above order) to a processor if it can pass the schedulability condition. If there are multiple processors that are feasible for assigning a task, the deadline-monotonic partitioning algorithm by Baruah and Fisher [4], [5] uses the first-fit strategy, but the analysis works for any arbitrary fitting strategy. The (theoretical) advantage of the first-fit strategy was not shown in the literature when we consider constrained- or arbitrary-deadline task systems.

The deadline-monotonic partitioning strategy is simple, but has been shown effective in the literature [4], [5], [13], [22]. When adopting speeding-up for resource augmentation, by using EDF on a processor, the deadline-monotonic partitioning proposed by Baruah and Fisher [4], [5] has been shown with a $3 - \frac{1}{M}$ speedup factor by Chen and Chakraborty [13], where M is the given number of identical processors. Prior to this paper, for fixed-priority multiprocessor partitioned scheduling for constrained- and arbitrary-deadline task systems, the best known results are by Fisher, Baruah, and Baker [22] with speedup factors $4 - \frac{2}{M}$ and $3 - \frac{1}{M}$ for arbitrary-deadline and constrained-deadline sporadic real-time task systems, respectively. All the above results are based on a linear-approximation to efficiently and safely test the schedulability under EDF or DM scheduling to decide whether a task can be assigned on a processor.

Our Contributions: Table I summarizes the related results and the contribution of this paper for multiprocessor parti-tioned scheduling. We focus on fixed-priority multiprocessor partitioned scheduling, and improve the best known results by Fisher, Baruah, and Baker [22]. The deadline-monotonic partitioning algorithm is explored in a great detail in this paper. Our contributions are:

- We show that the deadline-monotonic partitioning algorithm has a speedup factor $3 - \frac{1}{M}$ when considering task systems with arbitrary deadlines, where M is the number of processors. Such a factor holds for polynomial-time schedulability tests and exponential-time (exact) schedulability tests. Moreover, we also improve the speedup factor to 2.84306 when considering constrained-deadline task systems by using polynomial-time and pseudo-polynomial-time (exact) schedulability tests.

- The existing results by adopting the deadline-monotonic partitioning algorithm were analyzed based on approximated schedulability tests. One of our key contributions is to answer the question: *Will it be possible to further reduce the speedup factors by using exact tests in the deadline-monotonic partitioning algorithm?* Our answer to this question is **NO!!** Using exact tests in the above algorithm does not have any chance to reduce the speedup factors if the speedup factor analysis is oblivious of the particular fitting strategy used. We show that all the speedup factor analyses in this paper are asymptotically tight with polynomial-time schedulability tests and exponential-time (or pseudo polynomial-time) schedulability tests if the speedup factor analysis is oblivious of the particular fitting strategy used. As a result, to improve the speedup factor, better fixed-priority scheduling strategies or more precise analyses for concrete fitting strategies are needed.

2 System Models and Preliminary Results

2.1 Task and Platform Model

We consider a set $\mathbf{T} = \{\tau_1, \tau_2, \ldots, \tau_N\}$ of N independent sporadic real-time tasks. A task τ_i is defined by (C_i, T_i, D_i). That is, for task τ_i, D_i is its relative deadline, T_i is its minimum inter-arrival time (period), and C_i is its (worst-case) execution time. We consider identical processors in the platform. Therefore, no matter which processor a task is assigned to, the execution and timing property remains. According to the relations of the relative deadlines and the minimum inter-arrival times of the tasks in $\mathbf{T}$, the task set can be identified to be with 1) implicit deadlines, i.e., $D_i = T_i, \forall \tau_i \in \mathbf{T}$, 2) constrained deadlines, i.e., $D_i \leq T_i, \forall \tau_i \in \mathbf{T}$, or 3) arbitrary deadlines.

For brevity, the *utilization* of task τ_i is denoted by $U_i = \frac{C_i}{T_i}$. Moreover, let Δ_i be $\max\{U_i, \frac{C_i}{D_i}\}$. For a set $\mathbf{X}$, its cardinality is denoted by $|\mathbf{X}|$.

We will consider preemptive fixed-priority scheduling on each processor. Specifically, we will only use deadline-monotonic (DM) scheduling on each processor to assign the priority levels of the tasks. That is, task τ_i is with higher priority than τ_j if $D_i < D_j$, in which the ties are broken arbitrarily. Therefore, for the rest of this paper, we index the tasks from the shortest relative deadline to the longest, i.e., $D_i \leq D_j$ if $i < j$. Note that DM priority assignment is an optimal fixed-priority scheduling for implicit-deadline and constrained-deadline task sets [29]. Although DM priority assignment is not an optimal fixed-priority assignment when we consider arbitrary-deadline task systems, it has a constant speedup factor, to be discussed in Section 2.3.

2.2   Problem Definition

Given task set $\mathbf{T}$, a *feasible task partition* on M identical processors is a collection of M subsets, says, $\mathbf{T}_1, \mathbf{T}_2, \ldots, \mathbf{T}_M$, of $\mathbf{T}$ such that

- $\mathbf{T}_m \cap \mathbf{T}_{m'} = \emptyset$ for all $m \neq m'$,
- $\cup_{m=1}^{M} \mathbf{T}_m$ is equal to the input task set $\mathbf{T}$, and
- set $\mathbf{T}_m$ can meet the timing constraints by DM scheduling on a processor m.

Without loss of generality, we can assume that $U_i \leq 100\%$ and $\frac{C_i}{D_i} \leq 100\%$, i.e., $\Delta_i \leq 100\%$, for any task τ_i; otherwise, there is clearly no feasible task partition.

2.3   Speedup Factors/Bounds

This paper focuses on the case where the arrival times of the sporadic tasks are not specified. Therefore, the approximation is for the worst cases by considering the worst-case behaviour to be feasibly scheduled by DM. If an algorithm $\mathcal{A}$ for the studied problem has a *speedup factor* ρ, it guarantees to *always produce a feasible solution by speeding each processor up to ρ times of the original speed in the platform, if task set $\mathbf{T}$ can be feasibly scheduled (not restricted to DM) on the original M identical processors*. In other words, by taking the negation of the above statement, we know that *if the algorithm $\mathcal{A}$ fails to feasibly partition the task set $\mathbf{T}$ on M identical processors, there is no feasible task partition when each processor runs $\frac{1}{\rho}$ times slower than the original platform speed.* For the rest of this paper, we use 1 to denote the original platform speed. Therefore, running the platform at speed s implies that the execution time of task τ_i becomes $\frac{C_i}{s}$. Note that speedup factors are used for quantifying the behaviour of the designed algorithm. This is useful, especially for the negation part to quantify the error the algorithm makes when it does not provide a feasible solution.

For fixed-priority scheduling, the speedup factors of DM scheduling, with respect to the optimal uniprocessor EDF scheduling, are $\frac{1}{\ln 2}$, 1.76322, and 2 for implicit-deadline, constrained-deadline, and arbitrary-deadline task sets [17], respectively. To quantify the schedulability of the input task set, we would need to know the necessary condition for being schedulable at speed s on the M processors. The necessary conditions $\max_{\tau_i \in \mathbf{T}} \Delta_i \leq s$ and $\sum_{\tau_i \in \mathbf{T}} \frac{U_i}{M} \leq s$

are pretty straightforward. As we focus on arbitrary-deadline and constrained-deadline sporadic task systems, we can also quantify the necessary condition defined by the demand. Here, we can release the first job of tasks synchronously (say, at time 0), and the subsequent job arrivals should be as rapidly as legally possible. A necessary condition to be schedulable is to ensure that the total execution time of the jobs arriving no earlier than a and with relative deadlines no later than d is no more than $M \cdot (d - a)$ for any $a < d$. This is identical to the well-known *demand bound function* $dbf(\tau_i, t)$, as in [6], of a task τ_i within any time interval with length equal to t, defined as

$$dbf(\tau_i, t) = \max\left\{0, \left\lfloor \frac{t - D_i}{T_i} \right\rfloor + 1\right\} \times C_i. \qquad (1)$$

Therefore, as a necessary condition, to ensure the schedulability on M processors, if a task set is schedulable for an algorithm on M processors, then

$$\forall t > 0, \qquad \sum_{\tau_i \in \mathbf{T}} dbf(\tau_i, t) \leq Mt. \qquad (2)$$

With the above discussions, we can conclude the following lemma for the necessary condition to be schedulable by any algorithm, which has also been utilized by Chen and Chakraborty [13].

Lemma 1: A task set is not schedulable by any multiprocessor scheduling algorithm by running the M processors at any speed s if

$$\max\left\{\max_{t > 0} \frac{\sum_{\tau_i \in \mathbf{T}} dbf(\tau_i, t)}{Mt}, \frac{\sum_{\tau_i \in \mathbf{T}} U_i}{M}, \max_{\tau_i \in \mathbf{T}} \Delta_i\right\} > s. \qquad (3)$$

For the rest of the paper, we will focus ourselves on the negation part of the speedup factor analysis. That is, we are only interested in the failure cases of the partitioning algorithm and use Lemma 1 to quantify s for showing the speedup factors. Note that Lemma 1 is also the necessary condition for global multiprocessor scheduling. It may seem that we are more pessimistic by comparing to the necessary condition of global multiprocessor scheduling instead of that of partitioned multiprocessor scheduling. However, in our tightness analysis, comparing to partitioned scheduling and global scheduling does not differ very much.

3   Deadline-Monotonic Partitioning

This section presents the deadline-monotonic partitioning strategy, proposed by Baruah and Fisher [4], [5], [22], for the multiprocessor partitioned scheduling problem. Note that such a strategy works in general for fixed-priority scheduling (RM, DM) and dynamic-priority scheduling (EDF), by adopting proper schedulability tests. The speedup factor for EDF/DM was shown to be $3 - \frac{1}{M}$ and $4 - \frac{2}{M}$ [4], [5], [22] for constrained-deadline systems and arbitrary-deadline systems, respectively. When considering EDF scheduling, Chen and Chakraborty [15] improved the speedup factor to $2.6322 - \frac{1}{M}$ and $3 - \frac{1}{M}$ for constrained-deadline systems and arbitrary-deadline systems, respectively.

For completeness, we repeat the algorithm in [4], [5], [22], in which the pseudo-code is presented in Algorithm 1. Deadline-monotonic partitioning considers the given tasks

Algorithm 1 Deadline-Monotonic Partitioning

Input: set $\mathbf{T}$ of N tasks, M processors;
1: re-index (sort) tasks such that $D_i \leq D_j$ for $i < j$;
2: $\mathbf{T}_1 \leftarrow \{\tau_1\}; \mathbf{T}_m \leftarrow \emptyset, \forall m = 2, 3, \ldots, M$;
3: **for** $k = 2$ to N **do**
4:   **if** $\exists m \in \{1, 2, \ldots, M\}$ such that $\mathbf{T}_m \cup \{\tau_k\}$ is schedulable by DM fixed-priority scheduling **then**
5:     choose $m \in \{1, 2, \ldots, M\}$ **by preference** such that $\mathbf{T}_m \cup \{\tau_k\}$ is schedulable by DM fixed-priority scheduling;
6:     assign τ_k to processor m with $\mathbf{T}_m \leftarrow \mathbf{T}_m \cup \{\tau_k\}$;
7:   **else**
8:     return "no feasible schedule is found";
9:   **end if**
10: **end for**
11: return feasible task partition $\mathbf{T}_1, \mathbf{T}_2, \ldots, \mathbf{T}_M$;

---

from the shortest relative deadline to the longest relative deadline for assignment. When a task τ_k is considered, a processor m with $m \in \{1, 2, \ldots, M\}$ is selected to assign task τ_k, where $\mathbf{T}_m$ is the set of the tasks (as a subset of $\{\tau_1, \tau_2, \ldots, \tau_{k-1}\}$), which have been assigned to processor m before considering τ_k. If there is no feasible m that can feasibly schedule τ_k and $\mathbf{T}_m$ on the processor, we return that no feasible solution is found by this algorithm.

3.1 Fitting Strategy

The fitting strategy when we consider to assign task τ_k on a processor m can be

- the *first-fit* strategy: by choosing the minimum m that is feasible;
- the *arbitrary-fit* strategy: by choosing any m that is feasible (this is also the case if the speedup factor analysis is oblivious of the particular fitting strategy used);
- the *best-fit* strategy: by choosing the index m that has the maximum *workload-index*;
- the *worst-fit* strategy: by choosing the index m that has the minimum *workload-index*.

The workload-index can be defined as the total utilization or other means. The analysis in the literature [4], [5], [15], [22] works in general by using any fitting strategy listed above, even though in several cases only the first-fit strategy was mentioned in the descriptions [4], [5], [22].

3.2 Schedulability Tests for DM

Therefore, the remaining building block is to test whether task τ_k can be feasibly scheduled on a processor m under DM scheduling. This has been widely studied in the literature. We will review some of these methods and explain their corresponding speedup factors when they are adopted in Step 4 in Algorithm 1.

Constrained Deadline: To verify the schedulability of constrained-deadline task τ_k under fixed-priority scheduling in uniprocessor systems, the time-demand analysis (TDA) [28] can be adopted. That is, if and only if

$$\exists t \text{ with } 0 < t \leq D_k \quad and \quad C_k + \sum_{\tau_i \in \mathbf{T}_m} \left\lceil \frac{t}{T_i} \right\rceil C_i \leq t, \quad (4)$$

then task τ_k is schedulable under DM scheduling, where $\mathbf{T}_m$ is the set of tasks with higher priority than task τ_k since we sort the tasks according to their relative deadlines. TDA requires

pseudo-polynomial-time complexity, as all the points that lie in $(0, D_k]$ need to be checked for Eq. (4).

Fisher, Baruah, and Baker [22] approximate the test in Eq. (4) by testing only

$$\exists t \text{ with } 0 < t \leq D_k \quad and \quad C_k + \sum_{\tau_i \in \mathbf{T}_m} \left(1 + \frac{t}{T_i}\right) C_i \leq t. \quad (5)$$

Due to the linearity of the condition in Eq. (5), the test is equivalent to the verification of whether

$$C_k + \sum_{\tau_i \in \mathbf{T}_m} \left(1 + \frac{D_k}{T_i}\right) C_i \leq D_k \quad (6)$$

for constrained-deadline systems.

We can also approximate the schedulability test by using utilization-based analysis as follows: We classify the task set $\mathbf{T}_m$ into two subsets:

- $\mathbf{T}_m^1$ consists of the higher-priority tasks with periods smaller than D_k.
- $\mathbf{T}_m^2$ consists of the higher-priority tasks with periods larger than or equal to D_k.

The following theorem has been concluded recently by using a utilization-based schedulability-test framework [16].

Theorem 1 (Chen, Huang, Liu [16]): Task τ_k in a sporadic task system with constrained deadlines is schedulable by DM scheduling algorithm on processor m if

$$\left(\frac{C'_{k,m}}{D_k} + 1\right) \prod_{\tau_j \in \mathbf{T}_m^1} (U_j + 1) \leq 2. \quad (7)$$

where $C'_{k,m}$ is $C_k + \sum_{\tau_i \in \mathbf{T}_m^2} C_i$.

Arbitrary Deadline: For arbitrary-deadline systems, the exact schedulability test is to use a *busy-window* concept to evaluate the worst-case response time [27] by using TDA. The finishing time $R_{k,h}$ of the h-th job of task τ_k in the busy window is the minimum t such that

$$hC_k + \sum_{\tau_i \in \mathbf{T}_m} \left\lceil \frac{t}{T_i} \right\rceil C_i \leq t.$$

Therefore, its response time is $R_{k,h} - (h-1)T_k$. To test the busy window length of task τ_k, the busy window of task τ_k finishes on the h-th job if $R_{k,h} \leq hT_k$. The maximum response time among the jobs in the busy window is the worst-case response time [27]. The schedulability test is exact, but this takes exponential time complexity, since the busy window length can be up to the hyper-period, which is exponential of the input size.

The approximation by Fisher, Baruah, and Baker [22] in Eq. (5) remains feasible for analyzing the arbitrary-deadline systems. This leads to test whether

$$C_k + \sum_{\tau_i \in \mathbf{T}_m} \left(1 + \frac{D_k}{T_i}\right) C_i \leq D_k \quad and \quad (8a)$$

$$U_k + \sum_{\tau_i \in \mathbf{T}_m} U_i \leq 1. \quad (8b)$$

Eq. (8b) is important in arbitrary-deadline systems to ensure

that the approximation in Eq. (8a) does not underestimate the workload after D_k.

Moreover, Bini et al. [9] provide a tighter analysis than Eq. (8a). They show that the worst-case response time of task τ_k is at most

$$\frac{C_k + \sum_{\tau_i \in \mathbf{T}_m} C_i - \sum_{\tau_i \in \mathbf{T}_m} U_i C_i}{1 - \sum_{\tau_i \in \mathbf{T}_m} U_i}.$$

Therefore, the schedulability condition in Eqs. (8a) and (8b) can be rewritten as

$$C_k + D_k \left(\sum_{\tau_i \in \mathbf{T}_m} U_i \right) + \sum_{\tau_i \in \mathbf{T}_m} C_i - \sum_{\tau_i \in \mathbf{T}_m} U_i C_i \leq D_k, \quad (9a)$$

$$U_k + \sum_{\tau_i \in \mathbf{T}_m} U_i \leq 1. \quad (9b)$$

3.3 Time Complexity and Correctness

Since we use partitioned scheduling, as long as Algorithm 1 returns a task partition, it is guaranteed to be feasible if the test in Step 4 is a sufficient schedulability test for task τ_k by using DM scheduling. The time complexity of the algorithm depends upon the time complexity of the schedulability test and the fitting strategy. Suppose that the fitting strategy requires time complexity $O(F)$ for one task and the time complexity to test whether task τ_k is schedulable on processor m is $O(H)$. Then, the overall time complexity is $O(NMH + NF)$. All the fitting strategies mentioned in Section 3.1 are in polynomial time. Since H can be polynomial, pseudo polynomial, or exponential of the input size, the time complexity is dominated by the adopted schedulability test algorithm.

The main issue here is to answer what can be guaranteed when Algorithm 1 returns failure in task partitioning. We will quantify such failures by showing that s in Eq. (3) is also sufficiently large to provide the speedup factor guarantee (by using the negation arguments).

4 Analysis for Arbitrary-Deadline Systems

This section presents our analysis for arbitrary-deadline systems. Our analysis is similar to the analysis by Fisher, Baruah, and Baker [22], but is tighter. Here, we will mainly analyze the property by using the schedulability condition in Eqs. (8a) and (8b). At the end of this section, we will explain why the analysis also works for arbitrary-deadline TDA analysis by Lehoczky [27] and response time analysis by Bini et al. [9].

Theorem 2: The speedup factor of Algorithm 1 is $3 - \frac{1}{M}$ when adopting Eqs. (8a) and (8b) for DM schedulability test under any fitting strategy.

Proof: Suppose that Algorithm 1 fails to find a feasible assignment for task τ_k due to the failure when testing Eq. (8a) or Eq. (8b). Let $\mathbf{M}_1$ be the set of processors in which Eq. (8a) fails. Let $\mathbf{M}_2$ be the set of processors in which Eq. (8a) succeeds but Eq. (8b) fails. Since task τ_k cannot be assigned on any of the M processors, we have $|\mathbf{M}_1| + |\mathbf{M}_2| = M$. By

the violation of Eq. (8a), we know that

$$|\mathbf{M}_1| C_k + \sum_{m \in \mathbf{M}_1} \sum_{\tau_i \in \mathbf{T}_m} \left(1 + \frac{D_k}{T_i} \right) C_i > |\mathbf{M}_1| D_k$$

$$\Rightarrow |\mathbf{M}_1| \frac{C_k}{D_k} + \sum_{m \in \mathbf{M}_1} \sum_{\tau_i \in \mathbf{T}_m} \left(\frac{C_i}{D_k} + U_i \right) > |\mathbf{M}_1|. \quad (10)$$

By the violation of Eq. (8b), we know that

$$|\mathbf{M}_2| U_k + \sum_{m \in \mathbf{M}_2} \sum_{\tau_i \in \mathbf{T}_m} U_i > |\mathbf{M}_2|. \quad (11)$$

Recall that Δ_k is defined as $\max\{U_k, \frac{C_k}{D_k}\}$. By Eqs. (10) and (11), we know that

$$|\mathbf{M}_1| \frac{C_k}{D_k} + |\mathbf{M}_2| U_k + \sum_{i=1}^{k-1} U_i + \sum_{m \in \mathbf{M}_1} \sum_{\tau_i \in \mathbf{T}_m} \frac{C_i}{D_k} > M$$

$$\Rightarrow \quad M\Delta_k + \sum_{i=1}^{k-1} U_i + \sum_{i=1}^{k-1} \frac{C_i}{D_k} > M.$$

$$\Rightarrow \quad (M-1)\Delta_k + \sum_{i=1}^{k} U_i + \sum_{i=1}^{k} \frac{C_i}{D_k} > M.$$

$$\Rightarrow \quad (1 - \frac{1}{M})\Delta_k + \frac{\sum_{i=1}^{k} U_i}{M} + \sum_{i=1}^{k} \frac{C_i}{M D_k} > 1.$$

$$\Rightarrow \quad (3 - \frac{1}{M}) \max \left\{ \Delta_k, \frac{\sum_{i=1}^{k} U_i}{M}, \sum_{i=1}^{k} \frac{C_i}{M D_k} \right\} > 1. \quad (12)$$

Therefore, we know that either $\Delta_k > \frac{1}{3 - \frac{1}{M}}$, or $\frac{\sum_{i=1}^{k} U_i}{M} > \frac{1}{3 - \frac{1}{M}}$, or $\sum_{i=1}^{k} \frac{C_i}{M D_k} > \frac{1}{3 - \frac{1}{M}}$. Either of the former two cases implies the unschedulability of any scheduling algorithm with speed $\frac{1}{3 - \frac{1}{M}}$. The demand bound function at time D_k is $\sum_{i=1}^{N} dbf(\tau_i, D_k) \geq \sum_{i=1}^{k} C_i$. Therefore, by Eq. (2), we know that the condition $\sum_{i=1}^{k} \frac{C_i}{M D_k} > \frac{1}{3 - \frac{1}{M}}$ implies the unschedulability of any scheduling algorithm with speed $\frac{1}{3 - \frac{1}{M}}$.

Therefore, by using Lemma 1, we reach the conclusion. ∎

The following corollaries show that the speedup factor $3 - \frac{1}{M}$ holds for any schedulability tests discussed in Section 3 for arbitrary-deadline sporadic task systems under DM scheduling.

Corollary 1: The speedup factor of Algorithm 1 is $3 - \frac{1}{M}$ when adopting Eqs. (9a) and (9b) for DM schedulability test under any fitting strategy.

Proof: By not considering the term $- \sum_{\tau_i \in \mathbf{T}_m} U_i C_i$ in Eq. (9a), the violation of Eq. (9a) leads to the same conclusion in Eq. (10). Therefore, the speedup factor remains $3 - \frac{1}{M}$. ∎

Corollary 2: The speedup factor of Algorithm 1 is $3 - \frac{1}{M}$ when adopting the exact schedulability test for DM scheduling under any fitting strategy.

Proof: If task τ_k cannot pass the exact schedulability test, it also does not pass the sufficient test by using Eqs. (8a) and (8b). Therefore, we reach the same conclusion. ∎

4.1 Tightness Analysis

The following theorem shows that the analysis in Theorem 2 is asymptotically tight even for implicit-deadline systems with first-fit strategy.

Theorem 3: The speedup factor of Algorithm 1 is at least $3 - \frac{3}{M+1} - \gamma$ when adopting Eqs. (8a) and (8b) for DM schedulability test under the first-fit strategy, where γ is an arbitrarily small positive number.

Proof: This theorem is proved by a concrete input task system with $N = 2M$ tasks. There are M light tasks with execution time $\frac{1}{3M}$ and M heavy tasks with execution time $\frac{1+\epsilon}{3}$, in which ϵ is a small positive real, i.e., $\epsilon > 0$. The M light tasks are all with period $1 - \delta$ and relative deadline $1 - \delta$ with arbitrarily small and positive $\delta \ll \epsilon$. The M heavy tasks are all with period 1 and relative deadline 1. Therefore, the $2M$ tasks are indexed such that

- $C_i = \frac{1}{3M}$, $T_i = D_i = 1 - \delta$, for $i = 1, 2, \ldots, M$, and
- $C_i = \frac{1+\epsilon}{3}$, $T_i = D_i = 1$, for $i = M+1, M+2, \ldots, 2M$.

The setting of $0 < \delta \ll \epsilon$ is just to enforce the indexing. *We will directly take $\delta \to 0$ for the rest of the proof.*[1]

By using Algorithm 1 for the above task set when adopting Eqs. (8a) and (8b) for DM schedulability test under the first-fit strategy, the M light tasks are assigned on processor 1. Then, when task τ_k with $k > M$ is considered, the condition in Eq. (8a) always fails for any of the first $k - M$ processors. Therefore, task τ_k is assigned to processor $k - M + 1$, for $k = M+1, M+2, \ldots, 2M-1$. It is then clear that task τ_{2M} cannot be assigned on any of the M processors. Therefore, Algorithm 1 returns "no feasible solution is found".

By the above setting, we have $\sum_{i=1}^{2M} U_i = \frac{1+M+M\epsilon}{3}$. By using Lemma 1, we know that the speedup factor of the above task set is at least

$$\frac{1}{\frac{1+M+M\epsilon}{3}} = 3 - \frac{3 + 3\epsilon M}{M + 1 + \epsilon M} = 3 - \frac{3}{M+1} - \gamma$$

in which the factor $\gamma = \frac{3\epsilon M^2}{(M+1)(M+1+\epsilon M)}$ becomes negligible when ϵ is sufficiently small.

Such a factor can also be shown by a concrete partitioned schedule. By the pigeonhole principle, the solution that minimizes the maximum utilization of a processor is to assign a light task and a heavy task on a processor, in which the utilization on the processor is $\frac{1+\epsilon}{3} + \frac{1}{3M}$. Therefore, the task set is not schedulable on M processors only when the speed is slower than $\frac{1+\epsilon}{3} + \frac{1}{3M}$, provided that δ is 0. As a result, we reach the same conclusion. ∎

It may seem at first glance that the speedup factor $3 - \frac{1}{M}$ in Corollary 2 is pessimistic, since we do not actually use any property in the exponential-time exact schedulability test. However, the following theorem shows that the speedup factor $3 - \frac{1}{M}$ is asymptotically tight for an arbitrary fitting strategy, for any schedulability tests used in Theorem 2, Corollary 1, and Corollary 2. As a result, to improve the speedup factor, better fixed-priority scheduling strategies or more precise analysis for concrete fitting strategies are needed.

Theorem 4: The speedup factor of Algorithm 1 is at least $3 - \frac{3}{M+1} - \gamma$ under an arbitrary fitting strategy, for any schedulability tests used in Theorem 2, Corollary 1, and Corollary 2, where γ is an arbitrarily small positive number.

---

[1]That is, the tightness examples in Theorems 3, 4, and 9 hold even when δ is 0 if the tasks are indexed in the specified order. Setting δ to an arbitrarily small positive number is more precise mathematically, but this just complicates the presentation of the proofs without any added value.

Proof: This theorem is proved by a concrete input task system with $3M$ tasks. Let δ and ϵ be very small positive real numbers, with $\delta \ll \epsilon$. There are M tasks with execution time $\frac{1}{3M}$, period ∞, and relative deadline $1 - \delta$. There are M tasks with execution time $\frac{\epsilon}{3}$, period ϵ, and relative deadline 1. There are M tasks with execution time $\frac{1+\epsilon}{3}$, period ∞, and relative deadline $1 + \delta$.

Therefore, the $3M$ tasks are indexed such that

- $C_i = \frac{1}{3M}$, $T_i = \infty$, $D_i = 1 - \delta$, for $i = 1, 2, \ldots, M$, and
- $C_i = \frac{\epsilon}{3}$, $T_i = \epsilon$, $D_i = 1$, for $i = M+1, M+2, \ldots, 2M$.
- $C_i = \frac{1+\epsilon}{3}$, $T_i = \infty$, $D_i = 1 + \delta$, for $i = 2M+1, 2M+2, \ldots, 3M$.

Again, the setting of $\delta \ll \epsilon$ is just to enforce the indexing. *We will directly take $\delta \to 0$ for the rest of the proof.* Now, we consider a feasible task assignment for the first $3M - 1$ tasks, in which

- $\tau_1, \tau_2, \ldots, \tau_{M+1}$ are assigned on processor 1, and
- τ_i and τ_{i+M-1} are assigned on processor $i - M$ for $i = M+2, M+3, \ldots, 2M$.

By using Algorithm 1 for task τ_{3M}, we know that task τ_{3M} cannot be feasibly assigned on any of the M processors since $\forall 0 < t \leq 1$ and $m = 1, 2, \ldots, M$, we have $\frac{1+\epsilon}{3} + \sum_{\tau_i \in \mathbf{T}_m} \left\lceil \frac{t}{T_i} \right\rceil C_i > t$. Therefore, Algorithm 1 returns "no feasible solution is found".

By the above setting, we know that (1) $\sum_{i=1}^{3M} dbf(\tau_i, t) = 0$ for $0 < t < 1$, (2) $\sum_{i=1}^{3M} dbf(\tau_i, 1) = \frac{2\epsilon M + M + 1}{3}$, and (3) for $1 < t$,

$$\sum_{i=1}^{3M} dbf(\tau_i, t) \leq \frac{1}{3} + (t-1)\frac{1}{3}M + \frac{1+\epsilon}{3}M + \frac{\epsilon}{3}M$$

$$= \frac{2\epsilon M + Mt + 1}{3}.$$

As a result, $\max_{t>0} \frac{\sum_{i=1}^{3M} dbf(\tau_i, t)}{Mt} = \frac{1 + 2\epsilon + \frac{1}{M}}{3}$, when ϵ is small enough. Since $\max_{\tau_i \in \mathbf{T}} \Delta_i = \frac{1+\epsilon}{3}$ and $\sum_{i=1}^{3M} \frac{U_i}{M} = \frac{1}{3}$, by Lemma 1, the speedup factor of the above task set is

$$\frac{1}{\frac{1}{3M} + \frac{2\epsilon}{3} + \frac{1}{3}} = \frac{3M}{M + 2\epsilon M + 1}$$

$$= 3 - \frac{3 + 6\epsilon M}{M + 1 + 2\epsilon M} = 3 - \frac{3}{M+1} - \gamma,$$

in which the factor $\gamma = \frac{6\epsilon M^2}{(M+1)(M+1+2\epsilon M)}$ becomes negligible when ϵ is sufficiently small. The above analysis also works for the specific task partitioning which assigns three tasks τ_i, τ_{i+M}, and τ_{i+2M} on one processor for $i = 1, 2, \ldots, M$, and assigns the priority levels by using the optimal priority assignment by Audsley et al. [2], i.e., τ_{i+M} has the lowest priority on the processor. ∎

5 Analysis for Constrained Deadlines

This section presents the analysis for constrained-deadline sporadic real-time systems. By Theorem 3, we know that the method by Fisher, Baruah, and Baker [22] leads to a speedup factor 3 when M is sufficiently large even for implicit-deadline systems. The reason is mainly due to the pessimism of Eq. (6) in the schedulability test. To get better results, we

do need better tests. A more precise strategy is to simply use the exact test for constrained-deadline systems by spending pseudo-polynomial time complexity. We have already shown (by Corollary 2 and Theorem 4) that spending more time complexity does not help in arbitrary-deadline systems if the analysis does not use the property of any specific fitting strategy. Is this also the same for constrained-deadline systems?

We will first present the analysis by using TDA as the schedulability test in Step 4 in Algorithm 1. We will conclude later that such high time complexity also does not help reduce the speedup factor, compared to the results by using the hyperbolic bound in Theorem 1.

5.1 Speedup Factor by Adopting TDA

Now, suppose that task τ_k is the first task that fails to be assigned on any of the M processors by using TDA schedulability analysis in Step 4 in Algorithm 1. For notational brevity, let $\mathbf{T}^*$ be the set $\{\tau_1, \tau_2, \ldots, \tau_{k-1}\}$ of the tasks

Therefore, we know that this leads to

$$\forall m, \forall t, \text{ with } 0 < t \le D_k, \qquad C_k + \sum_{\tau_i \in \mathbf{T}_m} \left\lceil \frac{t}{T_i} \right\rceil C_i > t.$$

By taking a summation of all the $m = 1, 2, \ldots, M$ inequalities with respect to any t, we know that the unschedulability of task τ_k by Algorithm 1 implies that

$$\forall t \text{ with } 0 < t \le D_k, \qquad MC_k + \sum_{\tau_i \in \mathbf{T}^*} \left\lceil \frac{t}{T_i} \right\rceil C_i > Mt. \tag{13}$$

Therefore, by taking the negation, we know that if

$$\exists t \text{ with } 0 < t \le D_k, \text{ and } C_k + \sum_{\tau_i \in \mathbf{T}^*} \frac{\left\lceil \frac{t}{T_i} \right\rceil C_i}{M} \le t, \tag{14}$$

then Algorithm 1 by using TDA should succeed to assign task τ_k on one of the M processors. This is basically very similar to TDA with a minor difference by dividing the higher-priority workload by M.

Testing the schedulability condition of task τ_k according to Eq. (14) can be done by using the same strategy used in the $\mathbf{k^2U}$ framework [16] to prove Theorem 1 as follows.

We classify the $k - 1$ tasks in $\mathbf{T}^*$ into two subsets.

- $\mathbf{T}^{*1}$ consists of the tasks in $\mathbf{T}^*$ with periods smaller than D_k.
- $\mathbf{T}^{*2}$ consists of the tasks in $\mathbf{T}^*$ with periods larger than or equal to D_k.

Now, let C_k' be defined as follows:

$$C_k' = C_k + \sum_{\tau_i \in \mathbf{T}^{*2}} \frac{C_i}{M}. \tag{15}$$

Suppose that we have $\kappa - 1$ tasks in $\mathbf{T}^{*1}$. Clearly, according to the definition $\kappa \ge 1$. Now, we can rewrite the condition in Eq. (14) as follows: if

$$\exists t \text{ with } 0 < t \le D_k \text{ and } C_k' + \sum_{\tau_i \in \mathbf{T}^{*1}} \frac{\left\lceil \frac{t}{T_i} \right\rceil C_i}{M} \le t, \tag{16}$$

then Algorithm 1 by using TDA should succeed to assign task τ_k on one of the M processors.

For completeness, we repeat the definition of the $\mathbf{k^2U}$ framework and the key Lemma (with some simplifications to remove individual coefficients for each task τ_i) as follows.

Definition 1: A k-point effective schedulability test is a sufficient schedulability test of a fixed-priority scheduling policy by verifying the existence of $t_j \in \{t_1, t_2, \ldots t_k\}$ with $t_1 \le t_2 \le \cdots \le t_k$ such that

$$C_k + \sum_{i=1}^{k-1} \alpha t_i U_i + \sum_{i=1}^{j-1} \beta t_i U_i \le t_j, \tag{17}$$

where $C_k > 0$, $\alpha > 0$, $U_i > 0$, and $\beta > 0$ are dependent upon the setting of the task models and task τ_i.

Lemma 2 (Chen, Huang, and Liu [16]): For a given k-point effective schedulability test, defined in Definition 1, of a scheduling algorithm, in which $0 < \alpha \ne \infty$, and $0 < \beta \ne \infty$, task τ_k is schedulable by the scheduling algorithm if the following condition holds

$$\frac{C_k}{t_k} \le \frac{\frac{\alpha}{\beta} + 1}{\prod_{j=1}^{k-1}(\beta U_j + 1)} - \frac{\alpha}{\beta}. \tag{18}$$

Proof: This comes from Lemma 1 in [16]. ∎

By adopting the $\mathbf{k^2U}$ framework [16], we can conclude the following theorem.

Theorem 5: If

$$\prod_{\tau_i \in \mathbf{T}^{*1}} \left(1 + \frac{U_i}{M}\right) \le \frac{2}{1 + \frac{C_k'}{D_k}},$$

then task τ_k is schedulable under Algorithm 1 by using TDA.

Proof: In the proof, we will reindex the tasks to satisfy the monotonicity of t_i in Definition 1. That is, the $\kappa - 1$ higher-priority tasks in $\mathbf{T}^{*1}$ are reindexed to form the corresponding sequence $\tau_1, \tau_2, \ldots, \tau_{\kappa-1}$ for ensuring that the arrival times of the last jobs no later than D_k are in a non-decreasing order. For task τ_i in $\mathbf{T}^{*1}$, we set t_i as $\left\lfloor \frac{D_k}{T_i} \right\rfloor T_i$. Now, we reindex the $\kappa - 1$ higher-priority tasks such that $t_1 \le t_2 \le \cdots \le t_{\kappa-1}$. Moreover, let t_κ be D_k.

Since $T_i < D_k$ for any task τ_i in $\mathbf{T}^{*1}$, we have $t_i \ge T_i$. Therefore, for a given t_j with $j = 1, 2, \ldots, \kappa$, the demand requested up to time t_j in Eq. (16) is at most

$$C_k + \frac{\sum_{\tau_i \in \mathbf{T}^{*2}} C_i + \sum_{\tau_i \in \mathbf{T}^{*1}} \left\lceil \frac{t_j}{T_i} \right\rceil C_i}{M}$$

$$= C_k' + \frac{\sum_{i=1}^{\kappa-1} \left\lceil \frac{t_j}{T_i} \right\rceil C_i}{M}$$

$$\le C_k' + \frac{\sum_{i=1}^{\kappa-1} \frac{t_i}{T_i} C_i}{M} + \frac{\sum_{i=1}^{j-1} C_i}{M},$$

where the inequality comes from the indexing policy defined above, i.e., $\left\lceil \frac{t_j}{T_i} \right\rceil \le \frac{t_i}{T_i} + 1$ if $j > i$ and $\left\lceil \frac{t_j}{T_i} \right\rceil \le \frac{t_i}{T_i}$ if $j \le i$.

We only apply the test for these κ different t_i values, which is equivalent to the test of the existence of t_j for $j = 1, 2, \ldots, \kappa$ such that $C_k' + \frac{\sum_{i=1}^{\kappa-1} \frac{t_i}{T_i} C_i}{M} + \frac{\sum_{i=1}^{j-1} C_i}{M} \le t_j$. This satisfies

Definition 1 (when k is κ) with $\alpha = \frac{1}{M}$ and $\beta = \frac{1}{M}$. Therefore, by using Lemma 2, if

$$\frac{C'_k}{D_k} \leq \frac{1+1}{\prod_{\tau_i \in \mathbf{T}^{*1}}(\frac{U_i}{M}+1)} - 1,$$

then task τ_k is schedulable under Algorithm 1 by using TDA. By reorganizing the above equation, we reach the conclusion. ∎

The following corollary comes from the same proof of Lemma 2 and Theorem 5, which shows that the schedulability condition also implies a lower bound of the workload $\frac{C'_k}{D_k} + \frac{\sum_{\tau_i \in \mathbf{T}^{*1}} t_i U_i}{MD_k}$.

Corollary 3: If the schedulability condition in Theorem 5 is violated, i.e., $\prod_{\tau_i \in \mathbf{T}^{*1}}(1 + \frac{U_i}{M}) > \frac{2}{1+\frac{C'_k}{D_k}}$, then

$$\frac{C'_k}{D_k} + \frac{\sum_{\tau_i \in \mathbf{T}^{*1}} \left\lfloor \frac{D_k}{T_i} \right\rfloor T_i U_i}{MD_k} > \prod_{\tau_i \in \mathbf{T}^{*1}} \left(1 + \frac{U_i}{M}\right).$$

Proof: This comes from the same proof as in Lemma 2 (Lemma 1 in [16]) by changing the objective from minimizing C^*_k to minimizing $C^*_k + \sum_{i=1}^{k-1} \alpha t_i U_i$ to enforce the unschedulability. This property has been provided by Chen, Huang, and Liu [16] to prove the speedup factor of deadline-monotonic scheduling in uniprocessor systems. Moreover, this condition has also been exploited by von der Brüggen, Chen, and Huang [32] to prove the speedup factor of non-preemptive DM scheduling with respect to non-preemptive EDF scheduling. ∎

The remaining proofs in this section require some mathematical tools, which are provided in the following lemmas.

Lemma 3: Suppose that $\sum_{\tau_i \in \mathbf{T}^{1*}} U_i > 0$ is fixed and $U_i \geq 0$ for each task τ_i. Then, $\prod_{\tau_i \in \mathbf{T}^{1*}}\left(\frac{U_i}{M}+1\right)$ is maximized when $U_1 = U_2 = \cdots = U_{|\mathbf{T}^{1*}|}$.

Proof: This can be easily proved by the fact that $\prod_{\tau_i \in \mathbf{T}^{1*}}\left(\frac{U_i}{M}+1\right)$ is a concave function with respect to the (non-negative) values of U_i. ∎

Lemma 4: The infimum $\sum_{\tau_i \in \mathbf{T}^{1*}} \frac{U_i}{M}$ to enforce $\prod_{\tau_i \in \mathbf{T}^{1*}}\left(\frac{U_i}{M}+1\right) > x$ is $\ln(x)$.

Proof: This can be derived by using Lagrange Multiplier Method to find the minimum $\sum_{\tau_i \in \mathbf{T}^{1*}} \frac{U_i}{M}$ such that $\prod_{\tau_i \in \mathbf{T}^{1*}}\left(\frac{U_i}{M}+1\right) \geq x$. By Lemma 3, it is clear that the worst case is to have all the tasks with the same utilization. Suppose that $|\mathbf{T}^{1*}|$ is n. We know that all the tasks are with utilization $M(x^{\frac{1}{n}} - 1)$, and the utilization bound $\sum_{\tau_i \in \mathbf{T}^{1*}} \frac{U_i}{M}$ is $n(x^{\frac{1}{n}} - 1)$. This converges to $\ln(x)$ when n approaches ∞. ∎

With the above discussions, we can conclude the speedup factor.

Theorem 6: The speedup factor of Algorithm 1 for constrained-deadline task systems by using TDA is $\frac{1}{W(0.5)} \approx 2.84306$, where $W(z)$ is the Lambert W function, i.e., the unique solution of $z = W(z)e^{W(z)}$.

Proof: If $\prod_{\tau_i \in \mathbf{T}^{*1}}(\frac{U_i}{M}+1) \geq 2$, we can already conclude that $\sum_{\tau_i \in \mathbf{T}^{*1}} \frac{U_i}{M} \geq \ln 2$ by using Lemma 4, and the speedup factor is $1/\ln 2 < 2.84306$ for such a case. We focus on

the other case with $\prod_{\tau_i \in \mathbf{T}^{*1}}(\frac{U_i}{M} + 1) < 2$. Suppose that σ is $\frac{2}{\prod_{\tau_i \in \mathbf{T}^{*1}}(\frac{U_i}{M}+1)} - 1$, in which $\sigma > 0$.

If τ_k is not schedulable under Algorithm 1 by TDA, then

$$\frac{C_k}{D_k} + \frac{\sum_{i=1}^{k-1} dbf(\tau_i, D_k)}{MD_k}$$

$$= \frac{C_k}{D_k} + \frac{\sum_{\tau_i \in \mathbf{T}^{*2}} C_i}{MD_k} + \frac{\sum_{\tau_i \in \mathbf{T}^{*1}} dbf(\tau_i, D_k)}{MD_k}$$

$$\geq \frac{C'_k}{D_k} + \frac{\sum_{\tau_i \in \mathbf{T}^{*1}} \left\lfloor \frac{D_k}{T_i} \right\rfloor T_i U_i}{MD_k}$$

$$>_1 \prod_{\tau_i \in \mathbf{T}^{*1}} \left(1 + \frac{U_i}{M}\right) = \frac{1+\sigma}{2},$$

where $>_1$ is by Corollary 3. Suppose that $\frac{C_k}{D_k}$ is x. Therefore, we know that

$$\sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M} > \frac{1+\sigma}{2} - x. \qquad (19)$$

Moreover, with $\prod_{\tau_i \in \mathbf{T}^{*1}}(1 + \frac{U_i}{M}) = \frac{2}{1+\sigma}$ and the fact that τ_k is not schedulable under Algorithm 1 by using TDA, by Lemma 4, we have

$$\sum_{\tau_i \in \mathbf{T}^{*1}} \frac{U_i}{M} > \ln(\frac{2}{1+\sigma}). \qquad (20)$$

For the rest of the proof, we consider two separate cases:[2]

Case 1 $x \geq \sigma$: This is an easier case. We can conclude the speedup factor by using Eq. (20)

$$\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{U_i}{M}\right\} \geq \max\left\{\sigma, \ln(\frac{2}{1+\sigma})\right\}$$

$$\geq \min_{\sigma > 0}\left\{\max\{\sigma, \ln(\frac{2}{1+\sigma})\}\right\} =_1 \frac{1}{2.66793},$$

where $=_1$ holds when $e^\sigma(1 + \sigma) = 2$.

Case 2 $x < \sigma$: There are two subcases

- **Case 2a**: If $x > \frac{1+\sigma}{4}$, we know that $x > \frac{1+\sigma}{2} - x$. Therefore, by Eq. (19), $\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M}\right\} \geq x > \frac{1+\sigma}{4}$. Hence,

$$\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{U_i}{M}, \sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M}\right\}$$

$$> \max\left\{\frac{1+\sigma}{4}, \ln(\frac{2}{1+\sigma})\right\} \geq \min_{y \geq 0} \max\left\{y, \ln\frac{1}{2y}\right\}$$

$$= W(0.5) \approx \frac{1}{2.84306},$$

where $=$ holds when $ye^y = 0.5$.

- **Case 2b**: If $x \leq \frac{1+\sigma}{4}$, we know that $x \leq \frac{1+\sigma}{2} - x$. Therefore, by Eq. (19), $\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M}\right\} >$

---

[2]These cases are concluded by the following trick: Suppose that $g_1(y)$ is an increasing function and $g_2(y)$ is a decreasing function with respect to y when $y \geq 0$. Then, we know that the lower bound on $\max\{g_1(y), g_2(y)\}$ is the intersection of these two functions, i.e., $g_1(y^*)$, where $g_1(y^*) = g_2(y^*)$ if such a value y^* exists. That is, $\max\{g_1(y), g_2(y)\} \geq \min_{y \geq 0} \max\{g_1(y), g_2(y)\} = g_1(y^*)$.

$\frac{1+\sigma}{2} - x \geq \frac{1+\sigma}{4}$. Hence,

$$\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{U_i}{M}, \sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M}\right\}$$
$$> \max\left\{\frac{1+\sigma}{4}, \ln\left(\frac{2}{1+\sigma}\right)\right\} \geq \min_{y \geq 0} \max\left\{y, \ln(\frac{1}{2y})\right\}$$
$$= W(0.5) \approx \frac{1}{2.84306},$$

where $=$ holds when $ye^y = 0.5$.

Therefore, by all the above cases, we know that

$$\max\left\{\frac{C_k}{D_k}, \sum_{\tau_i \in \mathbf{T}^*} \frac{U_i}{M}, \sum_{\tau_i \in \mathbf{T}^*} \frac{dbf(\tau_i, D_k)}{M}\right\} > W(0.5) \approx \frac{1}{2.84306},$$

which concludes the proof by applying Lemma 1. ∎

5.2 Speedup Factor by Hyperbolic Bound

This subsection further presents the speedup factor of Algorithm 1 when adopting the hyperbolic bound in Eq. (7) for testing the schedulability of DM scheduling. The speedup factor analysis in Theorem 6 for TDA schedulability analysis relies only on the violation of the schedulability condition in Theorem 5. We will show that adopting the hyperbolic bound of Eq. (7) results in the same condition in Theorem 5. Therefore, we can reach the same conclusion as in Theorem 6 by using the hyperbolic bound in polynomial time. We use the same notations, e.g., $\mathbf{T}^{*1}$, $\mathbf{T}^{*2}$, κ, etc., as used in Section 5.1.

Theorem 7: If

$$\prod_{\tau_i \in \mathbf{T}^{*1}}\left(1 + \frac{U_i}{M}\right) \leq \frac{2}{1 + \frac{C'_k}{D_k}},$$

then task τ_k is schedulable under Algorithm 1 by using the hyperbolic bound in Eq. (7), where C'_k is $C_k + \frac{\sum_{\tau_i \in \mathbf{T}^{*2}} C_i}{M}$.

Proof: We prove this by contrapositive. Suppose that task τ_k is not schedulable under Algorithm 1 by using the hyperbolic bound in Eq. (7) . Therefore, for $m = 1, 2, \ldots, M$, we have

$$\left(\frac{C_k + \sum_{\tau_i \in \mathbf{T}^2_m} C_i}{D_k} + 1\right) \prod_{\tau_i \in \mathbf{T}^1_m}(U_i + 1) > 2.$$

By multiplying the above M inequalities we reach

$$2^M < \left(\prod_{m=1}^{M}\left(\frac{C_k + \sum_{\tau_i \in \mathbf{T}^2_m} C_i}{D_k} + 1\right)\right)\left(\prod_{\tau_i \in \mathbf{T}^{*1}}(U_i + 1)\right)$$
$$\leq_1 \left(\frac{C_k + \frac{\sum_{\tau_i \in \mathbf{T}^{*2}} C_i}{M}}{D_k} + 1\right)^M \left(\prod_{\tau_i \in \mathbf{T}^{*1}}\left(\frac{U_i}{M} + 1\right)^M\right),$$

where $\leq_1$ comes from Lemma 3 (for the first part) and from the fact $(1 + U_i) \leq (1 + U_i/M)^M$ when M is a positive integer and $U_i \geq 0$ (for the second part). Therefore, we conclude that the unschedulability of task τ_k implies that

$$2 < \left(\frac{C'_k}{D_k} + 1\right)\left(\prod_{\tau_i \in \mathbf{T}^{*1}}\left(\frac{U_i}{M} + 1\right)\right).$$

By contrapositive, we reach the conclusion. ∎

Theorem 8: The speedup factor of Algorithm 1 by using the hyperbolic bound in Eq. (7) is $\frac{1}{W(0.5)} \approx 2.84306$, where $W(z)$ is the Lambert W function, i.e., $z = W(z)e^{W(z)}$.

Proof: Since the schedulability condition remains the same as in Theorem 5, the speedup factor is also the same as in Theorem 6. ∎

5.3 Tightness Analysis

We conclude this section by showing that the above speedup factor analysis is tight when M is sufficiently large under an arbitrary fitting strategy.

Theorem 9: For constrained-deadline task systems, the speedup factor of Algorithm 1 is at least $\frac{1}{W(0.5)} \approx 2.84306$ when adopting TDA or the hyperbolic bound in Eq (7) for DM schedulability test under an arbitrary fitting strategy, when M is sufficiently large.

Proof: We prove this theorem by providing a concrete task system by assuming that M is sufficiently large. There are $N = 2M^2 + 1$ tasks. Let f be $\frac{2}{\frac{1}{W(0.5)}} \approx 0.7034674$, i.e., $\ln(\frac{1}{f}) = \frac{f}{2}$. Let δ be an arbitrarily small positive number just for enforcing the indexing:

- $T_i = D_i = f + (\lceil \frac{i}{M} \rceil - 1)\frac{1-f}{M-1}$, $C_i = \frac{1-f}{M-1}$, for $i = 1, 2, \ldots, M^2$,
- $T_i = \infty$, $D_i = 1 + \delta$, $C_i = \frac{1.5f-1}{M} \approx \frac{0.0552}{M}$, for $i = M^2 + 1, M^2 + 2, \ldots, 2M^2$, and
- $T_N = \infty$, $D_N = 1 + 2\delta$, $C_N = 0.5f + \epsilon$, with $N = 2M^2 + 1$, where ϵ is a positive small number.

We again simply take δ to 0 for the rest of the proof. Moreover, $\frac{1}{M}$ is also consider negligible for the simplicity of computation, since M is assumed to be sufficiently large.

For an arbitrary fitting algorithm, consider the following task assignment by assigning task τ_{i+jM} to processor i with $i = 1, 2, \ldots, M$ for every $j = 0, 1, 2, \ldots, 2M - 1$. It is not difficult to see that the above task assignment can be achieved feasibly and results in a feasible task assignment for the first $2M^2$ tasks. The set of the tasks assigned on processor m is denoted as $\mathbf{T}_m$. Now, consider task τ_{2M^2+1}, i.e., τ_N to be assigned on processor m. The overall execution time request at time 0 is $0.5f + \epsilon + 1.5f - 1 + 1 - f = f + \epsilon$ on processor m. Therefore, it can be easily seen that task $\mathbf{T}_m \cup \{\tau_N\}$ is not schedulable under DM scheduling on processor m since the TDA test in Eq. (6) fails. As a result, task τ_N cannot be assigned on any processor.

In this input task set, the utilization of the individual task and $\frac{C_i}{D_i}$ are not more than $\frac{f}{2} + \epsilon$ for each task τ_i. Moreover, the overall task utilization is $M\sum_{i=0}^{M-1} \frac{\mu}{f+i\mu}$, where $\mu = \frac{1-f}{M-1}$. Due to the assumption that M is sufficiently large, the above total utilization is a *left Riemann sum*, i.e., $M\sum_{i=0}^{M-1} \frac{\mu}{f+i\mu} \approx M\int_0^{1-f} \frac{1}{f+x}dx = M\ln(\frac{1}{f})$. By the fact that $\ln(\frac{1}{f}) = \frac{f}{2}$, we know that the total utilization is $M\frac{f}{2}$ when M is sufficiently large.

Now, we examine $\max_{t>0} \frac{\sum_{i=1}^{N} dbf(\tau_i, t)}{Mt}$. By definition, $\max_{t>0} \frac{\sum_{i=1}^{N} dbf(\tau_i, t)}{Mt} > \max_{t>0} \frac{\sum_{i=1}^{N-1} dbf(\tau_i, t)}{Mt}$. By the con-
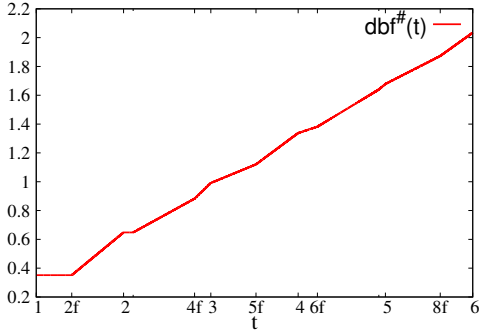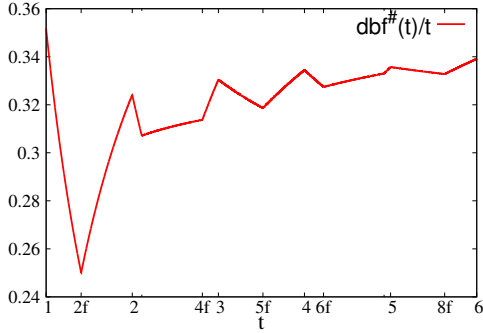
(a) $dbf^\sharp(t)$

(b) $dbf^\sharp(t)/t$

Fig. 1: Functions $dbf^\sharp(t)$ and $\frac{dbf^\sharp(t)}{t}$ for the proof of Theorem 9 with t in the range of $[1,6]$.

(a) $1 \le t \le 30$

(b) $30 \le t \le 4000$

Fig. 2: $\frac{dbf^\sharp(t)}{t}$ when t is a positive integer number (the curve is just for visualization) for the proof of Theorem 9, and the reference point $\frac{f}{2} \approx 0.3517337$

struction of the task set, we have (1) $\frac{\sum_{i=1}^{N-1} dbf(\tau_i,t)}{Mt} = 0$ if $0 < t < 1$, and (2) $\frac{\sum_{i=1}^{N-1} dbf(\tau_i,1)}{M} = \frac{f}{2}$ if $t = 1$. Therefore, we know that the speedup factor for this task set is purely dominated by $\max_{t>0} \frac{\sum_{i=1}^{N-1} dbf(\tau_i,t)}{Mt}$.

However, proving that $\frac{\sum_{i=1}^{N-1} dbf(\tau_i,t)}{Mt} \le \frac{f}{2}$ if $t > 1$ is pretty complicated. The proof involves quite some mathematical derivations, and is left in Appendix A in [12]. It should be clear that we can ignore the $M-1$ duplicated copies of the tasks by considering only the tasks assigned on one processor (before considering τ_N). The complete proof is Appendix A in [12]. The proof strategy first makes an over-approximation, denoted as $dbf^\sharp(t)$, of the sum $\sum_{\tau_i \in \mathbf{T}_m} dbf(\tau_i,t)$ of the demand bound functions at time t (after removing the $M-1$ duplicated copies). Based on such an over-approximation, it can be shown that $\max_{t>0} \frac{dbf^\sharp(t)}{t}$ happens when t is a positive integer for $\mathbf{T}_m$. It can then be proved that the maximum $\frac{dbf^\sharp(t)}{t}$ happens when $t = 1$ or $t = \infty$, in which both leads $\frac{dbf^\sharp(t)}{t}$ to $\frac{f}{2}$.

Figure 1 draws the functions $dbf^\sharp(t)$ and $\frac{dbf^\sharp(t)}{t}$ in the range of $[1,6]$. Figure 2 provides an illustrative view of $\frac{dbf^\sharp(t)}{t}$ from $t = 1,2,\ldots,4000$. It can also be easily shown by assigning τ_N to one processor alone, we can find a task partition that requires a speedup factor asymptotically equal to $\frac{2}{f}$ when M is sufficiently large. ∎
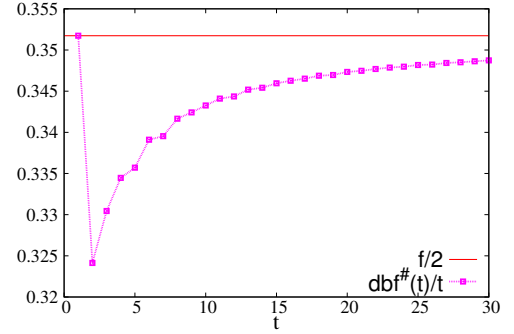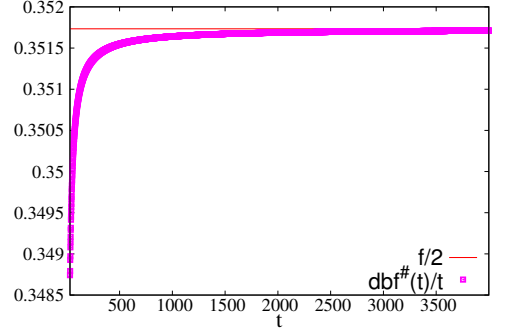
6 Concluding Remarks

This paper provides detailed analysis for the deadline-monotonic partitioning algorithm proposed by Fisher, Baruah, and Baker [22] for multiprocessor partitioned fixed-priority scheduling, by using exact schedulability tests and approximated schedulability tests. It may seem at first glance that using exact schedulability tests is more precise, but the proof shows that such exact tests are with the same speedup factors as approximated tests. We show that the deadline-monotonic partitioning algorithm has a speedup factor $3-\frac{1}{M}$ when considering task systems with arbitrary deadlines. Such a factor holds for polynomial-time schedulability tests and exponential-time schedulability tests. Moreover, we also improve the speedup factor to 2.84306 when considering constrained-deadline task systems.

The speedup factor analyses in this paper are asymptotically tight if the analysis is oblivious of the particular fitting strategy used. In all the tightness analyses, we only take Lemma 1, which also implicitly implies the reference to optimal global scheduling. The tasks are designed on purpose, e.g., M^2 tasks (instead of M tasks) with period ∞ in the proof of Theorem 9, to show that such factors also hold (asymptotically or with minor changes) for optimal partitioned scheduling. However, this does not limit the potential to have better speedup factors by adopting better fixed-priority scheduling strategies or more precise analysis for concrete fitting strategies.

References

[1] K. Albers and F. Slomka. An event stream driven approximation for the analysis of real-time systems. In *ECRTS*, pages 187–195, 2004.

[2] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A. J. Wellings. Applying new scheduling theory to static priority pre-emptive scheduling. *Software Engineering Journal*, 8(5):284–292, 1993.

[3] S. Baruah. The partitioned EDF scheduling of sporadic task systems. In *Real-Time Systems Symposium (RTSS)*, pages 116 –125, 2011.

[4] S. K. Baruah and N. Fisher. The partitioned multiprocessor scheduling of sporadic task systems. In *RTSS*, pages 321–329, 2005.

[5] S. K. Baruah and N. Fisher. The partitioned multiprocessor scheduling of deadline-constrained sporadic task systems. *IEEE Trans. Computers*, 55(7):918–923, 2006.

[6] S. K. Baruah, A. K. Mok, and L. E. Rosier. Preemptively scheduling hard-real-time sporadic tasks on one processor. In *IEEE Real-Time Systems Symposium*, pages 182–190, 1990.

[7] E. Bini. The quadratic utilization upper bound for arbitrary deadline real-time tasks. *IEEE Trans. Computers*, 64(2):593–599, 2015.

[8] E. Bini, G. C. Buttazzo, and G. M. Buttazzo. Rate monotonic analysis: the hyperbolic bound. *Computers, IEEE Transactions on*, 52(7):933–942, 2003.

[9] E. Bini, T. H. C. Nguyen, P. Richard, and S. K. Baruah. A response-time bound in fixed-priority scheduling with arbitrary deadlines. *IEEE Trans. Computers*, 58(2):279–286, 2009.

[10] A. Burchard, J. Liebeherr, Y. Oh, and S. H. Son. New strategies for assigning real-time tasks to multiprocessor systems. pages 1429–1442, 1995.

[11] S. Chakraborty, S. Künzli, and L. Thiele. Approximate schedulability analysis. In *IEEE Real-Time Systems Symposium*, pages 159–168, 2002.

[12] J. Chen. Partitioned multiprocessor fixed-priority scheduling of sporadic real-time tasks. *Computing Research Repository (CoRR)*, abs/1505.04693, http://arxiv.org/abs/1505.04693, 2015.

[13] J.-J. Chen and S. Chakraborty. Resource augmentation bounds for approximate demand bound functions. In *IEEE Real-Time Systems Symposium*, pages 272 – 281, 2011.

[14] J.-J. Chen and S. Chakraborty. Partitioned packing and scheduling for sporadic real-time tasks in identical multiprocessor systems. In *ECRTS*, pages 24–33, 2012.

[15] J.-J. Chen and S. Chakraborty. Resource augmentation for uniprocessor and multiprocessor partitioned scheduling of sporadic real-time tasks. *Real-Time Systems*, 49(4):475–516, 2013.

[16] J.-J. Chen, W.-H. Huang, and C. Liu. k2U: A general framework from k-point effective schedulability analysis to utilization-based tests. In *Real-Time Systems Symposium (RTSS)*, 2015.

[17] R. Davis, T. Rothvo, S. Baruah, and A. Burns. Quantifying the sub-optimality of uniprocessor fixed priority pre-emptive scheduling for sporadic tasksets with arbitrary deadlines. In *Real-Time and Network Systems (RTNS)*, pages 23–31, 2009.

[18] R. I. Davis and A. Burns. A survey of hard real-time scheduling for multiprocessor systems. *ACM Comput. Surv.*, 43(4):35, 2011.

[19] R. I. Davis, T. Rothvoß, S. K. Baruah, and A. Burns. Exact quantification of the sub-optimality of uniprocessor fixed priority pre-emptive scheduling. *Real-Time Systems*, 43(3):211–258, 2009.

[20] F. Eisenbrand and T. Rothvoß. Static-priority real-time scheduling: Response time computation is NP-Hard. In *Proceedings of the 29th IEEE Real-Time Systems Symposium, RTSS 2008, Barcelona, Spain, 30 November - 3 December 2008*, pages 397–406, 2008.

[21] N. Fisher and S. K. Baruah. A fully polynomial-time approximation scheme for feasibility analysis in static-priority systems with arbitrary relative deadlines. In *ECRTS*, pages 117–126, 2005.

[22] N. Fisher, S. K. Baruah, and T. P. Baker. The partitioned scheduling of sporadic tasks according to static-priorities. In *18th Euromicro Conference on Real-Time Systems, ECRTS'06, 5-7 July 2006, Dresden, Germany, Proceedings*, pages 118–127, 2006.

[23] R. L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

[24] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *J. ACM*, 34(1):144–162, 1987.

[25] A. Karrenbauer and T. Rothvoß. A 3/2-approximation algorithm for rate-monotonic multiprocessor scheduling of implicit-deadline tasks. In *Workshop of Approximation and Online Algorithms WAOA*, pages 166–177, 2010.

[26] C.-G. Lee, L. Sha, and A. Peddi. Enhanced utilization bounds for qos management. *IEEE Trans. Computers*, 53(2):187–200, 2004.

[27] J. P. Lehoczky. Fixed priority scheduling of periodic task sets with arbitrary deadlines. In *RTSS*, pages 201–209, 1990.

[28] J. P. Lehoczky, L. Sha, and Y. Ding. The rate monotonic scheduling algorithm: Exact characterization and average case behavior. In *IEEE Real-Time Systems Symposium*, pages 166–171, 1989.

[29] J. Y.-T. Leung and J. Whitehead. On the complexity of fixed-priority scheduling of periodic, real-time tasks. *Perform. Eval.*, 2(4):237–250, 1982.

[30] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20(1):46–61, 1973.

[31] A. K. Mok. Fundamental design problems of distributed systems for the hard-real-time environment. Technical report, Cambridge, MA, USA, 1983.

[32] G. von der Bruggen, J.-J. Chen, and W. Huang. Schedulability and optimization analysis for non-preemptive static priority scheduling based on task utilization and blocking factors. In *ECRTS*, pages 90–101, 2015.