# Task Scheduling with Self-Suspensions in Soft Real-Time Multiprocessor Systems* 

Cong Liu and James H. Anderson<br>Department of Computer Science, University of North Carolina at Chapel Hill


#### Abstract

In work on multiprocessor real-time systems, task scheduling with self-suspensions is a relatively unexplored topic. In this paper, soft real-time sporadic task systems are considered that include self-suspending tasks. Conditions are presented for guaranteeing bounded deadline tardiness in such systems under global EDF or FIFO multiprocessor scheduling. These conditions enable many soft real-time task systems with self-suspending tasks to be scheduled with little or no utilization loss.


## 1 Introduction

In many real-time systems, tasks interact with external devices that introduce self-suspension delays. Examples of such devices include solid-state and magnetic disks and network cards. Delays introduced by such devices can be moderate (e.g., roughly $15 \mu s$ per read and $200 \mu s$ per write for NAND Flash) or quite lengthy (e.g., roughly 15 ms for magnetic disks) [3,5]. Unfortunately, such delays quite negatively impact schedulability in real-time systems if deadline misses cannot be tolerated [12].

In this paper, we consider whether, on multiprocessor platforms, such negative impacts can be ameliorated if task deadlines are soft. Our focus on multiprocessors is motivated by the advent of multicore platforms. There is currently great interest in providing operating-system support to enable real-time workloads to be hosted on such platforms. Many such workloads can be expected to include self-suspending tasks. Moreover, in many settings, such workloads can be expected to have soft timing constraints. The soft timing constraint considered in this paper pertains to implicit-deadline sporadic task systems and requires that deadline tardiness be bounded.

All multiprocessor scheduling algorithms follow either a partitioning or global-scheduling approach (or some combination of the two). Under partitioning, tasks are statically assigned to processors, while under global scheduling, they

[^0]

Figure 1: Example task system with low utilization.
may migrate. Under partitioning schemes, constraints on overall utilization are required to ensure timeliness even if bounded deadline tardiness can be tolerated. On the other hand, a variety of global-scheduling approaches are capable of ensuring bounded tardiness in sporadic systems (without self-suspending tasks), as long as the system is not overutilized [6].

Unfortunately, if tasks may self-suspend, then bounded tardiness cannot be guaranteed even on uniprocessors without constraining the system in some way. Consider, for example, the uniprocessor task system, scheduled by the earliest-deadline-first (EDF) algorithm, shown in Fig. 1. This system consists of two tasks, each of which releases a new job every 10 time units that executes for 0.5 time unit, then suspends for 9 time units, and then executes for another 0.5 time unit. Tardiness grows unboundedly in this system, even though its total utilization is small.

Motivated by the observations above, we consider in this paper whether it is possible to specify reasonable constraints under which bounded deadline tardiness is guaranteed under global scheduling algorithms, for implicitdeadline sporadic tasks systems with self-suspending tasks. We focus specifically on two global scheduling algorithms that are capable of ensuring bounded tardiness for ordinary (suspension-less) sporadic systems with no utilization loss [2,7], namely, the global earliest-deadline-first (GEDF) algorithm and the global first-in first-out (GFIFO) algorithm.

Related work. To our knowledge, self-suspensions have not been considered before in the context of global real-time scheduling algorithms.

In work pertaining to uniprocessors (and by extension multiprocessors scheduled via partitioning), several schedulability tests have been presented for analyzing tasks with self-suspensions. These include a utilization-based test for

EDF [1] and response-time-bound tests for fixed-priority systems [4, 8-10] and EDF-scheduled systems [10]. On a more negative note, Ridouard et. al [13] have shown that the feasibility problem for hard real-time, independent tasks with self-suspensions is NP-hard in the strong sense. Ridouard and Richard [12] have also shown that no optimal on-line algorithm exists for task systems with selfsuspensions. In fixed-priority systems, scheduling penalties associated with self-suspensions can be lessened by using a technique called the period enforcer [11], which forces suspensions to occur more predictably.
Contributions. We show that GEDF's and GFIFO's ability to guarantee bounded tardiness in sporadic systems with self-suspending tasks hinges upon a task parameter that we call the "maximum suspension ratio," denoted $\xi_{\max }^{H}$, with range $[0,1]$. We present a general tardiness bound, which is applicable to either GEDF or GFIFO, that expresses tardiness as a function of $\xi_{\max }^{H}$ and other task parameters. This bound shows that task systems consisting of both self-suspending tasks and ordinary computational tasks that do not suspend can be supported with bounded tardiness if $\xi_{\text {max }}^{H}<1-\frac{U_{\text {sum }}^{s}+U_{L}^{c}}{m}$, where $m$ is the number of processors, $U_{s u m}^{s}$ is the total utilization of selfsuspending tasks, and $U_{L}^{c}$ is the total utilization of the $m-1$ computational tasks of highest utilization. The task model assumed in obtaining this result is very general and allows self-suspensions within a task's jobs to interleave arbitrarily with computation. We show via a counterexample that task systems that violate our utilization constraint may have unbounded tardiness. To assess the impact of this constraint, we present schedulability experiments that compare our analysis to a common approach for analyzing systems with self-suspensions wherein suspensions are merely treated as computation. In these experiments, our approach proved to be superior (i.e., could guarantee bounded tardiness for more systems) in most of the tested scenarios.
Organization. The rest of this paper is organized as follows. Sec. 2 describes our system model. In Sec. 3, our tardiness bound is derived and evaluated. Sec. 4 concludes.

## 2 System Model

We consider the problem of scheduling a set $\tau=$ $\left\{T_{1}, \ldots, T_{n}\right\}$ of $n$ independent sporadic tasks on $m \geq 2$ identical processors. Each task is released repeatedly, with each such invocation called a job. Jobs alternate between computation and suspension phases. We assume that each job of $T_{l}$ executes for at most $e_{l}$ time units (across all of its execution phases) and suspends for at most $s_{l}$ time units (across all of its suspension phases). We place no restrictions on how these phases interleave (a job can even begin
or end with a suspension phase). The $j^{t h}$ job of $T_{l}$, denoted $T_{l, j}$, is released at time $r_{l, j}$ and has a deadline at time $d_{l, j}$. Associated with each task $T_{l}$ is a period $p_{l}$, which specifies both the minimum time between two consecutive job releases of $T_{l}$ and the relative deadline of each such job, i.e., $d_{l, j}=r_{l, j}+p_{l}$. The utilization of a task $T_{l}$ is defined as $u_{l}=e_{l} / p_{l}$, and the utilization of the task system $\tau$ as $U_{\text {sum }}=\sum_{T_{i} \in \tau} u_{i}$. We require $e_{l}+s_{l} \leq p_{l}, u_{l} \leq 1$, and $U_{\text {sum }} \leq m$; otherwise, tardiness can grow unboundedly (if $e_{l}+s_{l}>p_{l}$ or $u_{l}>1$, then $T_{l}$ 's tardiness grows unboundedly; if $U_{\text {sum }}>m$, then the system is overloaded, which implies that tardiness grows unboundedly for at least one task).

A common case for real-time workloads is that both selfsuspending tasks and computational tasks (which do not suspend) co-exist. To reflect this, we let $U_{\text {sum }}^{s}$ denote the total utilization of all self-suspending tasks, and $U_{s u m}^{c}$ denote the total utilization of all computational tasks.

Successive jobs of the same task are required to execute in sequence. If a job $T_{i, j}$ completes at time $t$, then its response time is $t-r_{i, j}$ and its tardiness is $\max \left(0, t-d_{i, j}\right)$. A task's tardiness is the maximum tardiness of any of its jobs. Note that, when a job of a task misses its deadline, the release time of the next job of that task is not altered.

Unless stated otherwise, we henceforth assume that each job of any task $T_{l}$ executes for exactly $e_{l}$ time units. By Claim 4, given in an appendix, any tardiness bound derived for systems that meet this restriction applies to other systems as well.

So that our analysis can be more accurately applied in settings where a task's total suspension time varies from job to job, we assume that a fixed parameter $H(H \geq 1)$ is specified and that $S_{i}^{H}$ denotes the maximum total selfsuspension length for any $H(H \geq 1)$ consecutive jobs of task $T_{i}$. Note that if $H=1$, then a maximum per-job total suspension length is being assumed.

Under GEDF (GFIFO), released jobs are prioritized by their deadlines (release times). So that our results can be applied to both algorithms, we consider a generic scheduling algorithm (GSA) where each job is prioritized by some time point between its release time and deadline. Specifically, for any job $T_{i, j}$, we define a priority value $\rho_{i, j}=r_{i, j}+\kappa \cdot p_{i}$, where $0 \leq \kappa \leq 1$. Lower priority values denote higher priorities, and ties are assumed to be broken in favor of tasks with smaller indices. Note that GEDF and GFIFO are special cases of GSA where $\kappa$ is set to 1 and 0 , respectively.

## 3 A Tardiness Bound for GSA

We derive a tardiness bound for GSA by comparing the allocations to a task system $\tau$ in a processor sharing (PS) schedule and an actual GSA schedule of interest for $\tau$, both on $m$


Figure 2: Illustration of Defs. 1-3.
processors, and quantifying the difference between the two. We analyze task allocations on a per-task basis.

The time interval $\left[t_{1}, t_{2}\right)$, where $t_{2}>t_{1}$, consists of all time instances $t$, where $t_{1} \leq t<t_{2}$, and is of length $t_{2}-t_{1}$. For any time $t>0$, the notation $t^{-}$is used to denote the time $t-\varepsilon$ in the limit $\varepsilon \rightarrow 0+$, and the notation $t^{+}$is used to denote the time $t+\varepsilon$ in the limit $\varepsilon \rightarrow 0+$.

Definition 1. A task $T_{i}$ is active at time $t$ if there exists a job $T_{i, v}$ such that $r_{i, v} \leq t<d_{i, v}$.

Definition 2. Job $T_{i, v}$ is pending at time $t$ if $t \geq r_{i, v}$ and $T_{i, v}$ has not completed all of its execution phases by $t$. Note that $T_{i, v}$ is not pending at $t$ if it has completed all its execution phases by $t$ but not all of its suspension phases.

Definition 3. Job $T_{i, v}$ is enabled at $t$ if $t \geq r_{i, v}, T_{i, v}$ has not completed by $t$, and its predecessor (if any) has completed by $t$. A job is considered to be completed if it has finished its last phase (be it suspension or computation).

The above three definitions are illustrated in Fig. 2.
Let $A\left(T_{i, j}, t_{1}, t_{2}, S\right)$ denote the total allocation to the job $T_{i, j}$ in an arbitrary schedule $S$ in $\left[t_{1}, t_{2}\right)$. Then, the total time allocated to all jobs of $T_{i}$ in $\left[t_{1}, t_{2}\right)$ in $S$ is given by

$$
A\left(T_{i}, t_{1}, t_{2}, S\right)=\sum_{j \geq 1} A\left(T_{i, j}, t_{1}, t_{2}, S\right) .
$$

Consider a PS schedule $P S$. In such a schedule, $T_{i}$ executes with the rate $u_{i}$ when it is active. (Note that suspensions are not considered in the PS schedule.) Thus, if $T_{i}$ is active throughout $\left[t_{1}, t_{2}\right)$, then

$$
\begin{equation*}
A\left(T_{i, j}, t_{1}, t_{2}, P S\right)=\left(t_{2}-t_{1}\right) u_{i} . \tag{1}
\end{equation*}
$$

The difference between the allocation to a job $T_{i, j}$ up to time $t$ in a PS schedule and an arbitrary schedule $S$, denoted the lag of job $T_{i, j}$ at time $t$ in schedule $S$, is defined by

$$
\begin{equation*}
\operatorname{lag}\left(T_{i, j}, t, S\right)=A\left(T_{i, j}, 0, t, P S\right)-A\left(T_{i, j}, 0, t, S\right) \tag{2}
\end{equation*}
$$

The concept of lag is important because, if lags remain bounded, then tardiness is bounded as well. The $L A G$ for a
finite job set $J$ at time $t$ in the schedule $S$ is defined by

$$
\begin{gather*}
L A G(J, t, S)=\sum_{T_{i, j} \in J} \operatorname{lag}\left(T_{i, j}, t, S\right) \\
=\quad \sum_{T_{i, j} \in J}\left(A\left(T_{i, j}, 0, t, P S\right)-A\left(T_{i, j}, 0, t, S\right)\right) . \tag{3}
\end{gather*}
$$

Our tardiness-bound derivation follows a format originally presented in [2] and focuses on a given task system $\tau$. We order the jobs in $\tau$ based on their priorities: $T_{i, v} \prec T_{a, b}$ iff $\rho_{i, v}<\rho_{a, b}$ or $\left(\rho_{i, v}=\rho_{a, b}\right) \wedge(i<a)$. Let $T_{l, j}$ be a job of a task $T_{l}$ in $\tau, t_{d}=d_{l, j}$, and $S$ be a GSA schedule for $\tau$ with the following property.
(P) The tardiness of every job $T_{i, k}$ such that $T_{i, k} \prec T_{l, j}$ is at most $x+e_{i}+s_{i}$ in $S$, where $x \geq 0$.

Our objective is to determine the smallest $x$ such that the tardiness of $T_{l, j}$ is at most $x+e_{l}+s_{l}$. This would by induction imply a tardiness of at most $x+e_{i}+s_{i}$ for all jobs of every task $T_{i}$, where $T_{i} \in \tau$. We assume that $T_{l, j}$ finishes after $t_{d}$, for otherwise, its tardiness is trivially zero. The steps for determining the value for $x$ are as follows.

1. Determine an upper bound on the work pending for tasks in $\tau$ that can compete with $T_{l, j}$ after $t_{d}$. This is dealt with in Lemmas 1-3 in Sec. 3.1.
2. Determine a lower bound on the amount of work pending for tasks in $\tau$ that can compete with $T_{l, j}$ after $t_{d}$, required for the tardiness of $T_{l, j}$ to exceed $x+e_{l}+s_{l}$. This is dealt with in Lemma 4 in Sec. 3.2.
3. Determine the smallest $x$ such that the tardiness of $T_{l, j}$ is at most $x+e_{l}+s_{l}$, using the above upper and lower bounds. This is dealt with in Theorem 1 in Sec. 3.3.

Definition 4. We categorize jobs based on the relationship between their priorities and deadlines and those of $T_{l, j}$ :

$$
\begin{aligned}
& \mathbf{d}=\left\{T_{i, v}:\left(T_{i, v} \preceq T_{l, j}\right) \wedge\left(d_{i, v} \leq t_{d}\right)\right\} ; \\
& \mathbf{D}=\left\{T_{i, v}:\left(T_{i, v} \prec T_{l, j}\right) \wedge\left(d_{i, v}>t_{d}\right)\right\} .
\end{aligned}
$$

dis the set of jobs with deadlines at most $t_{d}$ with priority at least that of $T_{l, j}$. These jobs do not execute beyond $t_{d}$ in the PS schedule. Note that $T_{l, j}$ is in $\mathbf{d}$. $\mathbf{D}$ is the set of jobs that have higher priorities than $T_{l, j}$ and deadlines greater than $t_{d}$. Note that jobs not in $\mathbf{d} \cup \mathbf{D}$ have lower priority than those in $\mathbf{d} \cup \mathbf{D}$ and thus do not affect the scheduling of jobs in d $\cup \mathbf{D}$. For simplicity, we will henceforth assume that no job not in $\mathbf{d} \cup \mathbf{D}$ executes in either the PS or GSA schedule. Let $D_{C I}$ be the set of tasks with jobs in D. D consists of carryin jobs, which have a release time before $t_{d}$ and a deadline after $t_{d}$. Exactly one such job exists for each task in $D_{C I}$. (Note that $\mathbf{D}$ is empty under GEDF because jobs with later deadlines have lower priorities.)

Definition 5. A time instant $t$ is busy for a job set $J$ if all $m$ processors execute a job in $J$ at $t$. A time interval is busy for $J$ if each instant within it is busy for $J$.

The following claim follows from the definition of $L A G$.
Claim 1. If $L A G\left(\boldsymbol{d}, t_{2}, S\right)>L A G\left(\boldsymbol{d}, t_{1}, S\right)$, where $t_{2}>$ $t_{1}$, then $\left[t_{1}, t_{2}\right)$ is non-busy for $\boldsymbol{d}$. In other words, LAG for $\boldsymbol{d}$ can increase only throughout a non-busy interval.

An interval could be non-busy for $\mathbf{d}$ for two reasons:

1. There are not enough enabled non-suspended jobs in $\mathbf{d}$ to occupy all available processors. Such an interval is called non-busy non-displacing.
2. There are enabled non-suspended jobs in d that are not scheduled (because jobs in $\mathbf{D}$ occupy one or more processors).

Definition 6. Let $\delta_{k}$ be the amount of execution time consumed by a carry-in job $T_{k, v}$ by time $t_{d}$.

Definition 7. Let $B\left(\mathbf{D}, t_{d}, S\right)$ be the amount of work due to jobs in $\mathbf{D}$ that can compete with $T_{l, j}$ after $t_{d}$.

Since $\mathbf{d} \cup \mathbf{D}$ includes all jobs of higher priority than $T_{l, j}$, the competing work for $T_{l, j}$ is given by the sum of (i) the amount of work pending at $t_{d}$ for jobs in $\mathbf{d}$, and (ii) the amount of work $B\left(\mathbf{D}, t_{d}, S\right)$ demanded by jobs in $\mathbf{D}$ that competes with $T_{l, j}$ after $t_{d}$. Since jobs from d have deadlines at most $t_{d}$, they do not execute in the PS schedule beyond $t_{d}$. Thus, the work pending for them is given by $L A G\left(\mathbf{d}, t_{d}, S\right)$. Therefore, the competing work for $T_{l, j}$ after $t_{d}$ is given by $L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right)$. Let

$$
\begin{equation*}
Z=L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right) \tag{4}
\end{equation*}
$$

A summary of the terms defined so far, as well as some additional terms defined later, is presented in Table 1.

### 3.1 Upper Bound

In this section, we determine an upper bound on $Z$.
Definition 8. Let $t_{n}$ be the end of the latest non-busy nondisplacing interval for $\mathbf{d}$ before $t_{d}$, if any; otherwise, $t_{n}=0$.

The following two lemmas have been proved previously for both GEDF [2] and GFIFO [7] for ordinary sporadic task systems without self-suspensions. Note that the value of $L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right)$ depends only on allocations in the PS schedule $P S$ and allocations to jobs in $\mathbf{d} \cup \mathbf{D}$ in the actual schedule $S$ by time $t_{d}$. The PS schedule is not impacted by self-suspensions. Also, Property (P) alone is sufficient for determining how much work any job in $\mathbf{d} \cup \mathbf{D}$ other than $T_{l, j}$ completes before $t_{d}$. For these reasons, Lemmas 1 and 2 continue to hold for task systems with selfsuspensions. For completeness, proofs are given in the appendix.

| $m$ | Number of processors |
| :---: | :---: |
| $n$ | Number of tasks |
| $S\left(T_{l, j}\right)$ | Start time of job $T_{l, j}$ |
| $F\left(T_{l, j}\right)$ | Finish time of job $T_{l, j}$ |
| $t_{d}$ | Deadline of job $T_{l, j}$ |
| $E_{\text {sum }}^{s}$ | Total execution cost of all selfsuspending tasks in $\tau$ |
| $E_{\text {sum }}$ | Total execution cost of all tasks in $\tau$ |
| $S_{\text {sum }}^{s}$ | Total suspension length of all tasks in $\tau$ |
| $u_{\text {max }}^{s}$ | Maximum utilization of any selfsuspending task in $\tau$ |
| $U_{L}^{c}$ | Sum of the $\min (m-1, c)$ largest computational task utilizations, where $c$ is the number of computational tasks |
| $E_{L}^{c}$ | Sum of the $\min (m-1, c)$ largest computational task execution costs |
| $S_{i}^{H}$ | Maximum total self-suspension length for any $H(H>1)$ consecutive jobs of task $T_{i}$ |
| $S_{\text {max }}^{H}$ | Maximum total self-suspension length for any $H(H>1)$ consecutive jobs of any task in $\tau$ |
| $\xi_{i}^{H}$ | Suspension ratio of $T_{i}$ |
| $\xi_{\text {max }}^{H}$ | Maximum suspension ratio |
| $\delta_{k}$ | Amount of execution time consumed by a carry-in job $T_{k, v}$ by time $t_{d}$ |
| $B\left(\mathbf{D}, t_{d}, S\right)$ | Amount of work due to jobs in $\mathbf{D}$ that can compete with $T_{l, j}$ after $t_{d}$ |
| W | Amount of work due to jobs in $\mathbf{d} \cup \mathbf{D}$ that can compete with $T_{l, j}$ at or after $t_{d}+y$, including the work due for $T_{l, j}$ |
| Z | Amount of competing work for $T_{l, j}$ after $t_{d}$ |
| $t_{s}$ | Earliest non-busy instant in $\left[t_{d}, t_{d}+y\right)$ |
| $t_{p}$ | Finish time of $T_{l, j}$ 's predecessor, if it exists; otherwise $(j=1), t_{p}=0$ |

Table 1: Summary of notation.

Lemma 1. $L A G\left(\boldsymbol{d}, t_{d}, S\right) \leq \operatorname{LAG}\left(\boldsymbol{d}, t_{n}, S\right)+$ $\sum_{T_{k} \in D_{C I}} \delta_{k}\left(1-u_{k}\right)$, where $t \in\left[0, t_{d}\right]$.
Lemma 2. $\operatorname{lag}\left(T_{i}, t, S\right) \leq u_{i} \cdot x+e_{i}+u_{i} \cdot s_{i}$ for any task $T_{i}$ and $t \in\left[0, t_{d}\right]$.

Lemma 3 below upper bounds $L A G\left(\mathbf{d}, t_{n}, S\right)$.
Definition 9. Let $E_{\text {sum }}^{s}$ be the total execution cost of all self-suspending tasks in $\tau$. Let $E_{\text {sum }}$ be the total execution cost of all tasks in $\tau$. Let $S_{\text {sum }}^{s}$ be the total suspension length of all tasks in $\tau$. Let $u_{\max }^{s}$ be the maximum utilization of any self-suspending task in $\tau$.

Definition 10. Let $U_{L}^{c}$ be the sum of the $\min (m-1, c)$ largest computational task utilizations, where $c$ is the num-
ber of computational tasks. Let $E_{L}^{c}$ be the sum of the $\min (m-1, c)$ largest computational task execution costs.
Lemma 3. $L A G\left(\boldsymbol{d}, t_{n}, S\right) \leq\left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+$ $E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s}$.
Proof. By summing individual task lags at $t_{n}$, we can bound $L A G\left(\mathbf{d}, t_{n}, S\right)$. If $t_{n}=0$, then $L A G\left(\mathbf{d}, t_{n}, S\right)=0$, so assume $t_{n}>0$. Consider the set of tasks $\beta=\left\{T_{i}: \exists T_{i, v}\right.$ in $\mathbf{d}$ such that $T_{i, v}$ is enabled at $\left.t_{n}^{-}\right\}$. Given that the instant $t_{n}^{-}$is non-busy non-displacing, at most $m-1$ computational tasks in $\beta$ have jobs executing at $t_{n}^{-}$. Due to suspensions, however, $\beta$ may contain more than $m-1$ tasks. In the worst case, all suspending tasks in $\tau$ have a suspended enabled job at $t_{n}^{-}$and $\min (m-1, c)$ computational tasks have an enabled job executing at $t_{n}^{-}$. If task $T_{i}$ does not have an enabled job at $t_{n}^{-}$, then $\operatorname{lag}\left(T_{i}, t_{n}, S\right) \leq 0$. Therefore, by (3), we have

$$
\begin{aligned}
L A G\left(\mathbf{d}, t_{n}, S\right) & =\sum_{T_{i}: T_{i, v}^{w} \in \mathbf{d}} \operatorname{lag}\left(T_{i}, t_{n}, S\right) \\
& \leq \sum_{T_{i} \in \beta} \operatorname{lag}\left(T_{i}, t_{n}, S\right) \\
& \{\text { by Lemma } 2\} \\
& \leq \sum_{T_{i} \in \beta}\left(u_{i} \cdot x+e_{i}+u_{i} \cdot s_{i}\right) \\
& \leq\left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+E_{L}^{c} \\
& +u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s}
\end{aligned}
$$

The demand placed by jobs in $\mathbf{D}$ after $t_{d}$ is $B\left(\mathbf{D}, t_{d}, S\right)=$ $\sum_{T_{k} \in D_{C I}}\left(e_{k}-\delta_{k}\right)$. Thus, by (4) and Lemmas 1 and 3, we have the following upper bound:

$$
\begin{align*}
Z \leq & \left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s} \\
& +\sum_{T_{k} \in D_{C I}}\left(\delta_{i}\left(1-u_{k}\right)+\left(e_{k}-\delta_{k}\right)\right) \\
\leq & \left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s} \\
& +E_{\text {sum }} . \tag{5}
\end{align*}
$$

### 3.2 Lower Bound

Lemma 4, given below, establishes a lower bound on $Z$ that is necessary for the tardiness of $T_{l, j}$ to exceed $x+e_{l}+s_{l}$.
Definition 11. If job $T_{i, v}$ is enabled and not suspended at time $t$ but does not execute at $t$, then it is preempted at $t$.

Definition 12. If $T_{i, v}$ 's first phase is an execution (suspension) phase and it begins executing (a suspension) for the first time at $t$, then $t$ is called its start time, denoted $S\left(T_{i, v}\right)$. If $T_{i, v}$ 's last phase (be it execution or suspension) completes at time $t^{\prime}$, then $t^{\prime}$ is called its finish time, denoted $F\left(T_{i, v}\right)$.
Definition 13. Let $S_{\text {max }}^{H}=\max \left\{S_{1}^{H}, S_{2}^{H}, \ldots, S_{n}^{H}\right\}$. $\left(S_{i}^{H}\right.$ was defined earlier in Sec. 2.)

Definition 14. Let $\xi_{i}^{H}=\frac{S_{\max }^{H}}{S_{\max }^{H}+H \cdot e_{i}}$ be the suspension ratio of $T_{i}$. Let $\xi_{\max }^{H}=\max \left\{\xi_{1}, \xi_{2}, \ldots, \xi_{n}\right\}$ be the maximum suspension ratio.

Lemma 4. If the tardiness of $T_{l, j}$ exceeds $x+e_{l}+s_{l}$, then $Z>\left(1-\xi_{\text {max }}^{H}\right) \cdot m x-(m-1) e_{l}-m \cdot s_{l}-n \cdot\left(S_{\text {max }}^{H}+2 S_{\text {max }}^{1}\right)$.

Proof. We prove the contrapositive: we assume that

$$
\begin{align*}
Z \leq & \left(1-\xi_{\max }^{H}\right) \cdot m x-(m-1) e_{l} \\
& -m \cdot s_{l}-n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right) \tag{6}
\end{align*}
$$

holds and show that the tardiness of $T_{l, j}$ cannot exceed $x+$ $e_{l}+s_{l}$. Let $\eta_{l}$ be the amount of work $T_{l, j}$ performs by time $t_{d}$ in $S$. Define $y$ as follows.

$$
\begin{equation*}
y=\left(1-\xi_{\max }^{H}\right) \cdot x+\frac{\eta_{l}}{m} \tag{7}
\end{equation*}
$$

Let $W$ be the amount of work due to jobs in $\mathbf{d} \cup \mathbf{D}$ that can compete with $T_{l, j}$ at or after $t_{d}+y$, including the work due for $T_{l, j}$. We consider two cases.

Case 1. $\left[t_{d}, t_{d}+y\right)$ is a busy interval for $\mathbf{d} \cup \mathbf{D}$. In this case, the amount of work due to jobs in $\mathbf{d} \cup \mathbf{D}$ performed within $\left[t_{d}, t_{d}+y\right]$ is $m y$, and hence, $W=Z-m y$. Thus, by (6) and (7), $W \leq\left(1-\xi_{\text {max }}^{H}\right) \cdot m x-(m-1) e_{l}-m \cdot s_{l}-n$. $\left(S_{\max }^{H}+2 S_{\max }^{1}\right)-m y=\left(1-\xi_{\max }^{H}\right) \cdot m x-(m-1) e_{l}-m$. $s_{l}-n \cdot\left(S_{\max }^{H}+2 S_{\text {max }}^{1}\right)-\left(1-\xi_{\text {max }}^{H}\right) \cdot m x-\eta_{l}<0$. Since $T_{l, j}$ can suspend for at most $s_{l}$ time units after $t_{d}+y$ (and at least one task executes while it is not suspended), the amount of work performed by the system for jobs in $\mathbf{d} \cup \mathbf{D}$ during the interval $\left[t_{d}+y, F\left(T_{l, j}\right)\right)$ is at least $F\left(T_{l, j}\right)-t_{d}-y-s_{l}$. Hence, $F\left(T_{l, j}\right)-t_{d}-y-s_{l} \leq W<0$. Therefore, the tardiness of $T_{l, j}$ is $F\left(T_{l, j}\right)-t_{d}<y+s_{l}=\left(1-\xi_{\text {max }}^{H}\right)$. $x+\frac{\eta_{l}}{m}+s_{l} \leq x+e_{l}+s_{l}$.

Case 2. $\left[t_{d}, t_{d}+y\right)$ is a non-busy interval for $\mathbf{d} \cup \mathbf{D}$. Let $t_{s} \geq t_{d}$ be the earliest non-busy instant in $\left[t_{d}, t_{d}+y\right)$. Job $T_{l, j}$ cannot become enabled until its predecessor (if it exists) completes. Let $t_{p}$ be the finish time of $T_{l, j}$ 's predecessor (i.e., $\left.T_{i, j-1}\right)$, if it exists; otherwise $(j=1)$, let $t_{p}=0$. We consider three subcases.

Subcase 2.1. $t_{p} \leq t_{s}$ and $T_{l, j}$ is not preempted after $t_{s}$. In this case, $T_{l, j}$ performs its remaining execution and suspension phases in sequence without preemption after $t_{s}$ (note that, by Def. 11, $T_{l, j}$ is not considered to be preempted when it is suspended). Thus, because $t_{s}<t_{d}+y$, by (7), the tardiness of $T_{l, j}$ is at most $t_{s}+e_{l}-\eta_{l}+s_{l}-t_{d}<t_{d}+y+e_{l}-$ $\eta_{l}+s_{l}-t_{d}=\left(1-\xi_{\max }^{H}\right) \cdot x+\frac{\eta_{l}}{m}+e_{l}-\eta_{l}+s_{l} \leq x+e_{l}+s_{l}$.

The claim below will be used in the next two subcases.
Claim 2. The amount of work due to $\boldsymbol{d} \cup \boldsymbol{D}$ performed within $\left[t_{1}, t_{2}\right)$, where $S\left(T_{l, j}\right) \leq t_{1}<t_{2} \leq F\left(T_{l, j}\right)$, is at least $m\left(t_{2}-t_{1}\right)-(m-1) e_{l}-m \cdot s_{l}$.


Figure 3: Subcase 2.2.
Proof. Within $\left[t_{1}, t_{2}\right)$, all intervals during which $T_{l, j}$ is preempted are busy, and $T_{l, j}$ can execute for at most $e_{l}$ time. Within intervals where $T_{l, j}$ executes, at least one processor is occupied by $T_{l, j}$. Thus, at most $m-1$ processors are idle while $T_{l, j}$ executes (for at most $e_{l}$ time units) in $\left[t_{1}, t_{2}\right.$ ). Also, all processors can be idle while $T_{l, j}$ is suspended and this happens for at most $s_{l}$ time units in $\left[t_{1}, t_{2}\right)$.

Subcase 2.2. $t_{p} \leq t_{s}$ and $T_{l, j}$ is preempted after $t_{s}$. Let $t_{1}$ be the earliest time when $T_{l, j}$ is preempted after $t_{s}$, and let $t_{2}$ be the last time $T_{l, j}$ resumes execution after being preempted. (A finite number of jobs have higher priority than $T_{l, j}$, so $t_{2}$ exists.) Then, as shown in Fig. 3, $T_{l, j}$ executes or suspends within $\left[t_{s}, t_{1}\right)$. Also, because $T_{l, j}$ is preempted at $t_{1}, t_{1}$ is busy with respect to $\mathbf{d} \cup \mathbf{D}$. Within $\left[t_{1}, t_{2}\right), T_{l, j}$ could be repeatedly preempted. All such intervals during which $T_{l, j}$ is preempted must be busy in order for the preemption to happen. Note that $F\left(T_{l, j}\right) \leq t_{2}+e_{l}-\eta_{l}+s_{l}$. Thus, if $t_{2} \leq y+t_{d}$, then $F\left(T_{l, j}\right) \leq y+t_{d}+e_{l}-\eta_{l}+s_{l}$, which by (7) implies that $T_{l, j}$ 's tardiness is $F\left(T_{l, j}\right)-t_{d} \leq$ $y+e_{l}-\eta_{l}+s_{l} \leq\left(1-\xi_{\max }^{H}\right) \cdot x+e_{l}+s_{l} \leq x+e_{l}+s_{l}$, as required. If $t_{2}>t_{d}+y$, then by Claim 2, the amount of work due to $\mathbf{d} \cup \mathbf{D}$ performed within $\left[t_{s}, t_{d}+y\right)$ is at least $m\left(t_{d}+y-t_{s}\right)-(m-1) e_{l}-m \cdot s_{l}$. Because $\left[t_{d}, t_{s}\right)$ is busy, the work due to $\mathbf{d} \cup \mathbf{D}$ performed within $\left[t_{d}, t_{d}+y\right)$ is thus at least $m y-(m-1) e_{l}-m \cdot s_{l}$. Hence, the amount of work that can compete with $T_{l, j}$ (including work due to $T_{l, j}$ ) at or after $t_{d}+y$ is

$$
\begin{aligned}
W \leq & Z-\left(m y-(m-1) e_{l}-m \cdot s_{l}\right) \\
& \{\operatorname{by}(6)\} \\
\leq & \left(1-\xi_{\max }^{H}\right) \cdot m x-(m-1) e_{l}-m \cdot s_{l}- \\
& n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)-\left(m y-(m-1) e_{l}-m \cdot s_{l}\right) \\
= & \left(1-\xi_{\max }^{H}\right) \cdot m x-n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)-m y \\
& \{\operatorname{by}(7)\} \\
= & -n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)-\eta_{l} \\
\leq & 0 .
\end{aligned}
$$

Therefore, the tardiness of $T_{l, j}$ is $F\left(T_{l, j}\right)-t_{d} \leq y+$ $W \leq y=\left(1-\xi_{m a x}^{H}\right) \cdot x+\frac{\eta_{l}}{m}<x+e_{l}+s_{l}$.

Subcase 2.3: $t_{p}>t_{s}$. The earliest time $T_{l, j}$ can commence its first phase (be it an execution or suspension


Figure 4: Subcase 2.3.
phase) is $t_{p}$, as shown in Fig. 4. If fewer than $m$ tasks have enabled jobs in $\mathbf{d} \cup \mathbf{D}$ at any time instant within $\left[t_{s}, t_{p}\right)$, then $T_{l, j}$ will begin its first phase at $t_{p}$ and finish by time $t_{p}+e_{l}+s_{l}$. (Note that the number of enabled jobs in $\mathbf{d} \cup \mathbf{D}$ does not increase after $t_{d}$.) By Property ( P ) (applied to $T_{l, j}$ 's predecessor), $t_{p} \leq t_{d}-p_{l}+x+e_{l}+s_{l} \leq t_{d}+x$. Thus, the tardiness of $T_{l, j}$ is $F\left(T_{l, j}\right)-t_{d} \leq t_{p}+e_{l}+s_{l}-t_{d} \leq$ $x+e_{l}+s_{l}$.

The remaining possibility (which requires a much lengthier argument) is: $t_{p}>t_{s}$ and at least $m$ tasks have enabled jobs in $\mathbf{d} \cup \mathbf{D}$ at each time instant within $\left[t_{s}, t_{p}\right)$. In this case, given that at least $m$ tasks have enabled jobs in $\mathbf{d} \cup \mathbf{D}$ at $t_{s}, t_{s}$ is non-busy due to suspensions.

Let $W^{\prime}$ be the amount of work due to $\mathbf{d} \cup \mathbf{D}$ performed during $\left[t_{s}, t_{p}\right)$. Let $I$ be the total idle time in $\left[t_{s}, t_{p}\right)$, where the idle time at each instant is the number of idle processors at that instant. Then, $W^{\prime}+I=m \cdot\left(t_{p}-t_{s}\right)$. The following claim will be used to complete the proof of Subcase 2.3.

Claim 3. $W^{\prime} \geq\left(1-\xi_{\text {max }}^{H}\right) \cdot m\left(t_{p}-t_{s}\right)-n \cdot\left(S_{\text {max }}^{H}+2 S_{\text {max }}^{1}\right)$.
Proof. We begin by dividing the interval $\left[t_{s}, t_{p}\right)$ into subintervals on a per-processor basis. The subintervals on processor $k$ are denoted $\left[I T_{i}^{(k)}, E T_{i}^{(k)}\right.$ ), where $1 \leq i \leq q_{k}$, $I T_{i}^{(k)}=t_{s}, I T_{i+1}^{(k)}=E T_{i}^{(k)}$, and $E T_{q_{k}}^{(k)}=t_{p}$, as illustrated in Fig 5. With each such subinterval $\left[I T_{i}^{(k)}, E T_{i}^{(k)}\right)$, we associate a unique task, denoted $T_{i}^{(k)}$. We assume that during $\left[I T_{i}^{(k)}, E T_{i}^{(k)}\right), T_{i}^{(k)}$ executes only on processor $k$, and $E T_{i}^{(k)-}$ is the last time $T_{i}^{(k)}$ is enabled within $\left[t_{s}, t_{p}\right)$. Thus, if $E T_{i}^{(k)}<t_{p}$, then the last job of $T_{i}^{(k)}$ to be enabled within $\left[t_{s}, t_{p}\right.$ ) finishes its last phase (be it execution or suspension) at time $E T_{i}^{(k)-}$; if $E T_{i}^{(k)}=t_{p}$, then $T_{i}^{(k)}$ has enabled jobs throughout $\left[I T_{i}^{(k)}, t_{p}\right)$. Note that it is possible that $T_{i}^{(k)}$ executes or suspends within $\left[t_{s}, t_{p}\right)$ prior to $I T_{i}^{(k)}$. We call the subinterval $\left[I T_{i}^{(k)}, E T_{i}^{(k)}\right.$ ) the presence interval of $T_{i}^{(k)}$. The fact that a unique task can be associated with each subinterval follows from the assumption that at least $m$ tasks have jobs in $\mathbf{d} \cup \mathbf{D}$ that are enabled


Figure 5: Presence intervals within $\left[t_{s}, t_{p}\right]$.
at each time instant in $\left[t_{s}, t_{p}\right) .{ }^{1}$ (Note that multiple jobs of $T_{i}^{(k)}$ may execute during its presence interval.) We let $\lambda^{(k)}$ denote the set of all tasks that have presence intervals on processor $k$.

We now upper-bound the idleness on processor $k$ by bounding its idleness within one of its presence intervals. For conciseness, we denote this interval and its corresponding task as $[I T(T), E T(T))$ and $T$, respectively. If processor $k$ is idle at any time in $[I T(T), E T(T))$, then some job of $T$ is suspended at that time. Thus, the total suspension time of jobs of $T$ in $[I T(T), E T(T)$ ), denoted $I(T)$, upperbounds the idle time on processor $k$ in $[I T(T), E T(T))$.

Task $T$ may have multiple jobs that are enabled within its presence interval. Such a job is said to fully execute in the presence interval if it starts its first phase (be it execution or suspension) within the presence interval and also completes all of its execution phases in that interval (note that it may not complete all of its suspension phases). A job is said to partially execute in the presence interval if it starts its first phase (be it execution or suspension) before the presence interval or completes its last execution phase after the presence interval. Note that at most two jobs of $T$ could partially execute in its presence interval (namely, the first and last jobs to be enabled in that interval). We now prove that $I(T)$, and hence the idleness within $[I T(T), E T(T))$ on processor $k$, is at most $\xi_{\text {max }}^{H} \cdot(E T(T)-I T(T))+S_{\text {max }}^{H}+2 S_{\text {max }}^{1}$ (see Def. 13). Depending on the number of jobs of $T$ that execute during $T$ 's presence interval, we have two cases.

Case 1. $T$ has at most $H$ jobs that fully execute in its presence interval. (Additionally, $T$ may have at most two jobs that partially execute in its presence interval.) In this case, $I(T)$ is clearly at most $S_{\text {max }}^{H}+2 S_{\text {max }}^{1}$.

Case 2. T has more than $H$ jobs that fully execute in its presence interval. (Again, $T$ may have at most two jobs that partially execute in its presence interval.) In this case, the jobs of $T$ that are enabled in its presence interval can be divided into $n_{T}$ sets, where one set contains fewer than $H$ fully executed jobs plus at most two partially-executed jobs and each of the remaining $n_{T}-1$ sets contains exactly $H$ fully-executed jobs. Let $\theta$ denote the union of the latter $n_{T}-1$ job sets. Without loss of generality, we assume that

[^1]$\theta$ contains jobs that are enabled consecutively. The total suspension time for the first job set defined above is clearly at most $S_{\text {max }}^{H}+2 S_{\text {max }}^{1}$. We complete this case by showing that the total suspension time for all jobs in $\theta$ is at most $\xi_{\text {max }}^{H} \cdot(E T(T)-I T(T))$.

To ease the analysis, let $I T^{\prime}(T)$ be the start time of the first enabled job in $\theta$, and $E T^{\prime}(T)=\min (E T(T), F T)$, where $F T$ is the finish time of the last enabled job in $\theta$. Then $E T^{\prime}(T)-I T^{\prime}(T) \leq E T(T)-I T(T)$. Also, $E T^{\prime}(T)-I T^{\prime}(T)=I(\theta)+\Delta(\theta)+E(\theta)$, where $I(\theta)$ is the total suspension time of all jobs in $\theta$ within $\left[I T^{\prime}(T), E T^{\prime}(T)\right), \Delta(\theta)$ is the total preemption time of all jobs in $\theta$ within $\left[I T^{\prime}(T), E T^{\prime}(T)\right)$, and $E(\theta)$ is the total execution time of all jobs in $\theta$ within $\left[I T^{\prime}(T), E T^{\prime}(T)\right)$. (Recall that, by Def. 11, a suspended task is not considered to be preempted.) Given that $\theta$ contains $\left(n_{T}-1\right) H$ fully-executed jobs, $I(\theta) \leq\left(n_{T}-1\right) \cdot S_{T}^{H}$. Moreover, given our assumption that each job executes for the corresponding task's worst-case execution time, $E(\theta)=\left(n_{T}-1\right) \cdot H \cdot e(T)$, where $e(T)$ is the worst-case execution time of $T$. Thus,

$$
\begin{aligned}
I(\theta)= & \frac{I(\theta)}{E T^{\prime}(T)-I T^{\prime}(T)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
= & \frac{I(\theta)}{I(\theta)+\Delta(\theta)+E(\theta)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
\leq & \frac{I(\theta)}{I(\theta)+E(\theta)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
& \left\{\text { because } I(\theta) \leq\left(n_{T}-1\right) \cdot S_{T}^{H}\right\} \\
\leq & \frac{\left(n_{T}-1\right) \cdot S_{T}^{H}}{\left(n_{T}-1\right) \cdot S_{T}^{H}+E(\theta)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
& \left\{\text { because } E(\theta)=\left(n_{T}-1\right) \cdot H \cdot e(T)\right\} \\
\leq & \frac{\left(n_{T}-1\right) \cdot S_{T}^{H} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right)}{\left(n_{T}-1\right) \cdot S_{T}^{H}+\left(n_{T}-1\right) \cdot H \cdot e(T)} \\
= & \frac{S_{T}^{H}}{S_{T}^{H}+H \cdot e(T)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
& \{\text { by Def. } 13\} \\
\leq & \frac{S_{\text {max }}^{H}}{S_{\text {max }}^{H}+H \cdot e(T)} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
& \{\text { by Def. } 14\} \\
= & \xi_{\text {max }}^{H} \cdot\left(E T^{\prime}(T)-I T^{\prime}(T)\right) \\
\leq & \xi_{\text {max }}^{H} \cdot(E T(T)-I T(T)) .
\end{aligned}
$$

This concludes the proof of Case 2 of Claim 3.
Given that a task can be identified with only one presence interval and there are at most $n$ tasks, the idleness within [ $t_{s}, t_{p}$ ) on $m$ processor satisfies

$$
\begin{aligned}
I \leq & n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right) \\
& +\sum_{T \in \lambda^{(1)} \cup \ldots \cup \lambda^{(k)}} \xi_{\max }^{H} \cdot(E T(T)-I T(T)) \\
= & \xi_{\max }^{H} \cdot m\left(t_{p}-t_{s}\right)+n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right) .
\end{aligned}
$$

Thus, $W^{\prime}=m\left(t_{p}-t_{s}\right)-I \geq\left(1-\xi_{\text {max }}^{H}\right) \cdot m\left(t_{p}-t_{s}\right)-n$. $\left(S_{\max }^{H}+2 S_{\max }^{1}\right)$. This completes the proof of Claim 3.

We now complete the proof of Subcase 3.2 (and thereby Lemma 4). As shown in Fig. 4, $\left[t_{d}, t_{s}\right)$ and $\left[t_{p}, S\left(T_{l, j}\right)\right)$ are busy for $\mathbf{d} \cup \mathbf{D}$. By Claim 2, the amount of work due to $\mathbf{d} \cup \mathbf{D}$ performed in $\left[S\left(T_{l, j}\right), F\left(T_{l, j}\right)\right)$ is at least $m\left(F\left(T_{l, j}\right)-\right.$ $\left.S\left(T_{l, j}\right)\right)-(m-1) e_{l}-m \cdot s_{l}$. By Claim 3, the amount of work due to $\mathbf{d} \cup \mathbf{D}$ performed in $\left[t_{s}, t_{p}\right)$ is at least $\left(1-\xi_{\text {max }}^{H}\right)$. $m\left(t_{p}-t_{s}\right)-n \cdot\left(S_{\text {max }}^{H}-2 S_{\text {max }}^{1}\right)$. By summing over all of these subintervals, we can lower-bound the amount of work due to $\mathbf{d} \cup \mathbf{D}$ performed in $\left[t_{d}, F\left(T_{l, j}\right)\right)$, i.e., $Z$ :

$$
\begin{align*}
Z \geq & m\left(t_{s}-t_{d}\right)+\left(1-\xi_{\max }^{H}\right) \cdot m\left(t_{p}-t_{s}\right) \\
& -n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)+m\left(S\left(T_{l, j}\right)-t_{p}\right) \\
& +m\left(F\left(T_{l, j}\right)-S\left(T_{l, j}\right)\right)-(m-1) e_{l}-m \cdot s_{l} . \tag{8}
\end{align*}
$$

By (6) and (8), we therefore have

$$
\begin{array}{ll} 
& \left(1-\xi_{\max }^{H}\right) \cdot m x-(m-1) e_{l}-m \cdot s_{l} \\
& -n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right) \\
\geq & m\left(t_{s}-t_{d}\right)+\left(1-\xi_{\max }^{H}\right) \cdot m\left(t_{p}-t_{s}\right) \\
& -n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)+m\left(S\left(T_{l, j}\right)-t_{p}\right) \\
& +m\left(F\left(T_{l, j}\right)-S\left(T_{l, j}\right)\right)-(m-1) e_{l}-m \cdot s_{l}
\end{array}
$$

which gives,

$$
F\left(T_{l, j}\right)-t_{d} \leq\left(1-\xi_{\max }^{H}\right) \cdot x+\xi_{\max }^{H} \cdot\left(t_{p}-t_{s}\right)
$$

According to Property (P) (applied to $T_{l, j}$ 's predecessor), $t_{p}-t_{s} \leq t_{p}-t_{d} \leq x-p_{l}+e_{l}+s_{l} \leq x$. Therefore, $F\left(T_{l, j}\right)-t_{d} \leq\left(1-\xi_{\text {max }}^{H}\right) \cdot x+\xi_{\text {max }}^{H} \cdot x<x+e_{l}+s_{l}$.

### 3.3 Determining $x$

Setting the upper bound on $L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right)$ in (5) to be at most the lower bound in Lemma 4 will ensure that the tardiness of $T_{l, j}$ is at most $x+e_{l}+s_{l}$. The resulting inequality can be used to determine a value for $x$. By (5) and Lemma 4, this inequality is $\left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+$ $E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s}+E_{\text {sum }} \leq\left(1-\xi_{\max }^{H}\right) \cdot m x-(m-$ 1) $e_{l}-m \cdot s_{l}-n \cdot\left(S_{\text {max }}^{H}+2 S_{\text {max }}^{1}\right)$.

Let $V=E_{\text {sum }}^{s}+E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s}+E_{\text {sum }}+(m-$ 1) $e_{l}+m \cdot s_{l}+n \cdot\left(S_{\max }^{H}+2 S_{\max }^{1}\right)$. Solving for $x$, we have

$$
\begin{equation*}
x \geq \frac{V}{\left(1-\xi_{\max }^{H}\right) \cdot m-U_{s u m}^{s}-U_{L}^{c}} \tag{9}
\end{equation*}
$$

$x$ is well-defined provided $U_{s u m}^{s}+U_{L}^{c}<\left(1-\xi_{\max }^{H}\right) \cdot m$. If this condition holds and $x$ equals the right-hand side of (9), then the tardiness of $T_{l, j}$ will not exceed $x+e_{l}+s_{l}$. A value for $x$ that is independent of the parameters of $T_{l}$ can be obtained by replacing $(m-1) e_{l}+m \cdot s_{l}$ with $\max _{l}((m-$ 1) $\left.e_{l}+m \cdot s_{l}\right)$ in $V$.

Theorem 1. With $x$ as defined in (9), the tardiness of any task $T_{l}$ scheduled under $G S A$ is at most $x+e_{l}+s_{l}$, provided $U_{\text {sum }}^{s}+U_{L}^{c}<\left(1-\xi_{\max }^{H}\right) \cdot m$.

For GFIFO and GEDF, the bound in Theorem 1 can be improved.

Corollary 1. For GFIFO, Theorem 1 holds with $V$ replaced by $V-E_{\text {sum }}+\sum_{p_{i}>p_{l}} e_{i}$ in the numerator of (9).

Proof. Under GFIFO, $D_{C I}$ consists of carry-in jobs that are released before $r_{l, j}$ and have deadlines later than $t_{d}$, which implies that these jobs have periods greater than $p_{l}$. Thus, the upper bound in (5) can be refined to obtain $L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right) \leq\left(U_{\text {sum }}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+$ $E_{L}^{c}+u_{\max }^{s} \cdot S_{\text {sum }}^{s}+\sum_{p_{i}>p_{l}} e_{i}$. Using this upper bound to solve for $x$, the corollary follows.

Corollary 2. For GEDF, Theorem 1 holds with $V$ replaced by $V-E_{\text {sum }}$ in the numerator of (9).

Proof. Under GEDF, the demand placed by jobs in D after $t_{d}$ is zero because $\mathbf{D}=\emptyset$. Thus, under GEDF, $L A G\left(\mathbf{d}, t_{d}, S\right)+B\left(\mathbf{D}, t_{d}, S\right) \leq\left(U_{s u m}^{s}+U_{L}^{c}\right) \cdot x+E_{\text {sum }}^{s}+$ $E_{L}^{c}+u_{\text {max }}^{s} \cdot S_{\text {sum }}^{s}$. Using this upper bound to solve for $x$, the corollary follows.

### 3.4 A Counterexample

Previous research has shown that every sporadic task system for which $U_{\text {sum }} \leq m$ without self-suspensions has bounded tardiness under GEDF and GFIFO [2,7]. We now show that it is possible for a task system containing self-suspending tasks to have unbounded tardiness under GEDF or GFIFO if the utilization constraint in Theorem 1 is violated.

Consider a two-processor task set $\tau$ that consists of three self-suspending tasks: $T_{1}=((3$ (exec.), 7 (susp.)), 10 (period) $), T_{2}=((1$ (exec.), 8 (susp.), 1 (exec.)), 10 (period)), and $T_{3}=((1$ (exec.), 8 (susp.), 1 (exec.)), 10 (period)). For this system, $\xi_{\text {max }}^{H}=0.8($ assuming $H=1)$ and $U_{\text {sum }}^{s}+U_{L}^{c}=$ 0.7. Thus, $\left(1-\xi_{\text {max }}^{H}\right) \cdot m=0.4<U_{\text {sum }}^{s}+U_{L}^{c}$, which violates the condition stated in Theorem 1. Fig. 6 shows the tardiness of each task in this system under GFIFO/GEDF by job index assuming each job is released as early as possible. We have verified analytically that the tardiness growth rate seen in Fig. 6 continues indefinitely.

### 3.5 Experimental Evaluation

In this section, we describe experiments conducted using randomly-generated task sets to evaluate the applicability of the tardiness bound in Theorem 1. Our goal is to examine how restrictive the theorem's utilization cap is, and to compare it with another commonly-used approach, which we call SuspToComp, wherein all suspension phases are treated


Figure 6: Tardiness growth rates in counterexample.

|  |  | short suspensions $\xi_{\max }=0.05$ | moderate suspensions $\xi_{\max }=0.2$ | long suspensions $\xi_{\max }=0.5$ |
| :---: | :---: | :---: | :---: | :---: |
| light | min: avg: max: | $\begin{gathered} 2.6 \mu s \\ 197 \mu \mathrm{~s} \\ 526 \mu \mathrm{~s} \end{gathered}$ | $\begin{gathered} 12 \mu \mathrm{~s} \\ 938 \mu \mathrm{~s} \\ 2.5 \mathrm{~ms} \end{gathered}$ | $50 \mu \mathrm{~s}$ 3.75 ms 10 ms |
| medium | min: <br> avg: <br> max: | $263 \mu \mathrm{~s}$ $789 \mu s$ <br> 1.6 ms | 1.25 ms 3.75 ms 7.5 ms | 5 ms 15 ms 30 ms |
| heavy | min: <br> avg: <br> max: | $\begin{aligned} & 789 \mu \mathrm{~s} \\ & 2.2 \mathrm{~ms} \\ & 4.2 \mathrm{~ms} \\ & \hline \end{aligned}$ | 3.75 ms 10.3 ms 20 ms | $\begin{gathered} 15 \mathrm{~ms} \\ 41.25 \mathrm{~ms} \\ 80 \mathrm{~ms} \\ \hline \end{gathered}$ |

Table 2: Per-job suspension-length ranges.
as computation phases. From [2,7], tardiness is bounded under SuspToComp provided $U_{\text {sum }} \leq m$ and $U_{L} \leq m$, where $U_{L}$ is the sum of the $\min (m-1, n)$ largest task utilizations. (Note that, under SuspToComp, treating suspensions as computation causes utilizations to be higher.)

In our experiments, task sets were generated as follows. Task periods were uniformly distributed over [ $50 \mathrm{~ms}, 100 \mathrm{~ms}$ ]. Per-task utilizations were distributed differently for each experiment using three uniform distributions: [0.001,0.1] (light), [0.1,0.3] (medium), and [0.3,0.8] (heavy). Task execution costs were calculated from periods and utilizations. We varied $U_{\text {sum }}^{s}$ as follows: $U_{\text {sum }}^{s}=$ $0.1 \cdot U_{\text {sum }}$ (suspensions are relatively infrequent), $U_{\text {sum }}^{s}=$ $0.4 \cdot U_{\text {sum }}$ (suspensions are moderately frequent), and $U_{\text {sum }}^{s}=0.7 \cdot U_{\text {sum }}$ (suspensions are frequent). Moreover, we varied $\xi_{\max }$ as follows: 0.05 (suspensions are short), 0.2 (suspensions are moderate), and 0.5 (suspensions are long). Table. 2 shows suspension-length ranges generated by these parameters. We also varied $U_{\text {sum }}$ within $\{1,2, \ldots$, $8\}$. For each combination of $\left(u_{\max }, U_{\text {sum }}^{s}, U_{\text {sum }}\right), 1,000$ task sets were generated for an eight-processor system. For each generated system, soft real-time schedulability (i.e., the ability to ensure bounded tardiness) was checked under SuspToComp and using the condition stated in Theorem 1. In doing so, system overheads were ignored (factoring overheads into our analysis is beyond the scope of this paper).

The schedulability results that were obtained are shown in Fig. 8 (the organization of which is explained in the figure's caption). Each curve plots the fraction of the generated


Figure 7: Average tardiness under LA, as computed via Theorem 1.
task sets the corresponding approach successfully scheduled, as a function of total utilization. As Fig. 8 shows, our approach proved to be superior, sometimes by a substantial margin, in all tested scenarios summarized in the first two rows of graphs. However, in many of the scenarios summarized in the third row of grows, SuspToComp proved to be superior. In these scenarios, task utilizations are high and suspensions are long or frequent. Our analysis is negatively impacted in such cases because $U_{L}^{c}$ tends to be large when utilizations are high, and $\xi_{\max }$ tends to be large when suspensions are long. It is worth noting, however, that our approach allows certain tasks to be designated as computational tasks. Thus, the SuspToComp approach is really a special case of our approach. It would be interesting to investigate intermediate choices between the two extremes of modeling all versus no suspensions as computation.

In addition to schedulability, the magnitude of tardiness, as computed using the bound in Theorem 1, is of importance. Fig. 7 depicts the average of the computed bounds for each of the tested scenarios in our experimental framework for the case where $U_{\text {sum }}=m$ and $U_{s}^{\text {sum }}=0.1 \cdot U_{\text {sum }}$ (that is, for each scenario in this case, an average of all bounds for all tasks in all schedulable task sets is plotted). As can be seen, tardiness is reasonable if task utilizations are low and suspensions are short. However, as either task utilizations or suspension lengths increase, tardiness increases, as an examination of the bound in Theorem 1 suggests should be the case. With large task utilizations and long suspension lengths, tardiness is quite high, perhaps unacceptably so, even though these systems are deemed to be schedulable.

## 4 Conclusion

We have derived a tardiness bound that can be applied to globally-scheduled sporadic task systems that include self-


Figure 8: Soft real-time schedulability results. In the first (respectively, second and third) row of graphs, light (respectively, medium and heavy) per-task utilizations are assumed. In the first (respectively, second and third) column of graphs, relatively infrequent (respectively, moderately frequent and frequent) suspensions are assumed. Each graph gives three curves per tested approach for the cases of short, moderate, and long suspensions, respectively. The label "LA-s(m/l)" indicates the approach of this paper assuming short (moderate/long) suspensions. Similar "SC" labels are used for SuspToComp.
suspending tasks. This bound is applicable to a class of global algorithms that includes GEDF and GFIFO. The derived tardiness bound requires overall utilization to be constrained. We have shown via a counterexample that utilization constraints are fundamental. We also presented schedulability experiments that suggest that our constraint is quite liberal in many systems and is often less pessimistic than modeling suspensions as computation.

## References

[1] U. Devi. An improved schedulability test for uniprocessor periodic task systems. In Proc. of the 15th Euromicro Conf. on Real-Time Systems, pp. 23-30, 2003.
[2] U. Devi and J. Anderson. Tardiness bounds under global EDF scheduling on a multiprocessor. In Proc. of the 26th IEEE Real-Time Systems Symp., pp. 330-341, 2005.
[3] W. Kang, S. Son, J. Stankovic, and M. Amirijoo. I/Oaware deadline miss ratio management in real-time embed-
ded databases. In Proc. of the 28th IEEE Real-Time Systems Symp., pp. 277-287, 2007.
[4] I. Kim, K. Choi, S. Park, D. Kim, and M. Hong. Real-time scheduling of tasks that contain the external blocking intervals. In Proc. of the 2nd Int'l Workshop on Real-Time Computing Systems and Applications, pp. 54-59, 1995.
[5] S. Lee and B. Moon. Design of flash-based DBMS: An inpage logging approach. In Proc. of the 2007 ACM SIGMOD Conf. on Management of Data, pp. 55-66, 2007.
[6] H. Leontyev and J. Anderson. Generalized tardiness bounds for global multiprocessor scheduling. In Proc. of the 28 th IEEE Real-Time Systems Symp., pp. 413-422, 2007.
[7] H. Leontyev and J. Anderson. Tardiness bounds for FIFO scheduling on multiprocessors. In Proc. of the 19th Euromicro Conf. on Real-Time Systems, pp. 71-80, 2007.
[8] J. W. S. Liu. Real-Time Systems. Prentice Hall, 2000.
[9] J. C. Palencia and M. Gonzlez Harbour. Schedulability analysis for tasks with static and dynamic offsets. In Proc. of the 19th IEEE Real-Time Systems Symp., pp. 26-37, 1998.
[10] J. C. Palencia and M. Gonzlez Harbour. Response time analysis of EDF distributed real-time systems. In J. Embedded Comput., Vol.1, pp. 225-237, 2005.
[11] R. Rajkumar. Dealing with Suspending Periodic Tasks. IBM Thomas J. Watson Research Center, 1991.
[12] F. Ridouard and P. Richard. Worst-case analysis of feasibility tests for self-suspending tasks. In Proc. of the 14th Real-Time and Network Systems, pp. 15-24, 2006.
[13] F. Ridouard, P. Richard, and F. Cottet. Negative results for scheduling independent hard real-time tasks with selfsuspensions. In Proc. of the 25th IEEE Real-Time Systems Symp., pp. 47-56, 2004.

## 5 Appendix

In this appendix, we prove Claim 4 and Lemmas 1 and 2. In proving Claim 4, we lift the restriction, given earlier in Sec. 2, that each job executes for the corresponding task's worst-case execution time. To state this claim, some additional terminology is required.

We say that a sporadic task system $\tau$ is concrete if the release time (and hence deadline) and actual execution cost and suspension time of every job of each task is fixed. Two concrete task systems are compatible if they have the same jobs with the same release times (they can have different actual execution and suspension times). A concrete task system $\tau$ is maximal if the actual execution time of any job equals the corresponding task's worst-case execution time.

Claim 4. For any concrete task system $\tau$, there exists a compatible maximal concrete task system $\tau^{\prime}$ such that, for any job $T_{i, k}$, its response time in the GSA schedule for $\tau^{\prime}$ is at least its response time in the GSA schedule for $\tau$.

Proof. The existence of the desired maximal concrete system is demonstrated via a construction method in which computation phases are ranked as follows: (i) if $T_{i, k} \prec$ $T_{x, y}$, then all computation phases of $T_{i, k}$ are ranked before all computation phases of $T_{x, y}$; (ii) earlier computation phases of $T_{i, k}$ are ranked before later computation phases of $T_{i, k}$. Let $\epsilon$ be a positive value that is small enough so that within any interval $[t, t+\epsilon)$, each processor schedules exactly one job or no job, and the actual and worst-case execution cost of any computation phase is a multiple of $\epsilon$.

Consider a computation phase $C$ of a job $T_{i, k}$ that is not maximal. We show that the length of $C$ can be increased by $\epsilon$ by adding to the end of $C$ a piece of computation $\rho$ of length $\epsilon$. In so doing, it may be necessary to reduce the length of a lower-ranked computation phase by $\epsilon$ and to reduce the length of a subsequent suspension phase (if any) of $T_{i, k}$. By inducting over all computation phases in rank order, and by iteratively increasing any non-maximal execution time by $\epsilon$, we can obtain a compatible concrete task system that is maximal. The construction method will ensure that no job's response time is reduced.

Let $[t, t+\epsilon$ ) denote the time interval where $\rho$ should be added to the schedule (according to GSA). If, before adding $\rho$, task $T_{i}$ is scheduled within $[t, t+\epsilon$ ), then the computation phase of $T_{i}$ executing at that time, call it $C^{\prime}$, is ranked lower than $C$. In this case, we can accommodate $\rho$ by reducing $C^{\prime}$ in length by $\epsilon .^{2}$ If $C$ and $C^{\prime}$ are separated by a suspension phase, then the length of that suspension phase must be defined to be zero.

In the rest of the proof, we consider the other possibility: before adding $\rho, T_{i}$ is not scheduled within $[t, t+\epsilon$ ) (and hence, it is not scheduled in $\left[t^{\prime}, t+\epsilon\right.$ ), where $t^{\prime}$ is the completion time of $C$ ). In this case, if there is an idle processor in $[t, t+\epsilon$ ), then $\rho$ can be scheduled there without modifying the length of any lower-ranked computation phase. On the other hand, if there is no idle processor, then, as $\rho$ should be scheduled in $[t, t+\epsilon)$, there must be a computation phase ranked lower than $C$ scheduled then. We can accommodate $\rho$ and allow it to be scheduled in $[t, t+\epsilon$ ) by reducing the length of that lower-ranked computation phase by $\epsilon$. If $C$ is followed by a suspension phase, then, once $\rho$ has been added to the schedule, it may be necessary to reduce the length of that suspension phase. In particular, if, before adding $\rho, T_{i, k}$ was suspended in $[t, t+\epsilon$ ), then the length of that suspension phase must be reduced so that it starts as $t+\epsilon$.

Note that the construction method used in this proof strongly exploits the fact that, in our task model, suspension phases are upper-bounded, and hence, can be reduced.

[^2]Theorem 1 shows that, for any maximal concrete task system, the tardiness of any task $T_{l}$ scheduled under GSA is at most $x+e_{l}+s_{l}$, with $x$ as defined in (9). By Claim 4, the same is true for any non-maximal concrete task system.

Let $A\left(J, t_{1}, t_{2}, S\right)$ denote the total time allocated to all jobs in the job set $J$ in $\left[t_{1}, t_{2}\right)$ in the schedule $S$.

Lemma 1. $L A G\left(\mathbf{d}, t_{d}, S\right) \leq L A G\left(\mathbf{d}, t_{n}, S\right)+$ $\sum_{T_{k} \in D_{C I}} \delta_{k}\left(1-u_{k}\right)$, where $t \in\left[0, t_{d}\right]$.

Proof. By (3), we have

$$
\begin{align*}
L A G\left(\mathbf{d}, t_{d}, S\right) \leq & L A G\left(\mathbf{d}, t_{n}, S\right)+A\left(\mathbf{d}, t_{n}, t_{d}, P S\right) \\
& -A\left(\mathbf{d}, t_{n}, t_{d}, S\right) \tag{10}
\end{align*}
$$

We split $\left[t_{n}, t_{d}\right)$ into $z$ non-overlapping intervals $\left[t_{p_{i}}, t_{q_{i}}\right), 1 \leq i \leq z$, such that $t_{n}=t_{p_{1}}, t_{q_{i-1}}=t_{p_{i}}$, and $t_{q_{z}}=t_{d}$. Each interval $\left[t_{p_{i}}, t_{q_{i}}\right)$ is either busy or nonbusy displacing for $\mathbf{d}$, by the selection of $t_{n}$. We assume that the intervals are defined so that for each non-busy displacing interval $\left[t_{p_{i}}, t_{q_{i}}\right)$, if a task in $D_{C I}$ executes in $\left[t_{p_{i}}, t_{q_{i}}\right)$ then it executes continuously throughout $\left[t_{p_{i}}, t_{q_{i}}\right)$; we let $\alpha_{i}$ denote the set of such tasks.

We now bound the difference between the work performed in the PS schedule and the GSA schedule $S$ across each of these intervals $\left[t_{p_{i}}, t_{q_{i}}\right)$. The sum of these bounds will give us a bound on the total allocation difference throughout $\left[t_{n}, t_{d}\right)$. Depending on the nature of the interval $\left[t_{p_{i}}, t_{q_{i}}\right)$, two cases are possible.

Case 1. $\left[t_{p_{i}}, t_{q_{i}}\right)$ is busy. Since in $S$ all processors are occupied by jobs in $\mathbf{d}$, we have $A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, P S\right)-$ $A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, S\right) \leq U_{\text {sum }}\left(t_{q_{i}}, t_{p_{i}}\right)-m\left(t_{q_{i}}-t_{p_{i}}\right) \leq 0$.

Case 2. $\left[t_{p_{i}}, t_{q_{i}}\right)$ is non-busy displacing. The cumulative utilization of all tasks $T_{k} \in \alpha_{i}$ is $\sum_{T_{k} \in \alpha_{i}} u_{k}$. The carry-in jobs of these tasks do not belong to $\mathbf{d}$, by the definition of $\mathbf{d}$. Therefore, the allocation of jobs in d during $\left[t_{p_{i}}, t_{q_{i}}\right)$ in $P S$ is $A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, P S\right) \leq\left(t_{q_{i}}-t_{p_{i}}\right)\left(m-\sum_{T_{k} \in \alpha_{i}} u_{k}\right)$. All processors are occupied at every time instant in the interval $\left[t_{p_{i}}, t_{q_{i}}\right)$, because it is displacing. Thus, $A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, S\right)=$ $\left(t_{q_{i}}-t_{p_{i}}\right)\left(m-\left|\alpha_{i}\right|\right)$. Therefore, the allocation difference for jobs in $\mathbf{d}$ throughout the interval is

$$
\begin{align*}
& A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, P S\right)-A\left(\mathbf{d}, t_{p_{i}}, t_{q_{i}}, S\right) \\
\leq & \left(t_{q_{i}}-t_{p_{i}}\right)\left(\left(m-\sum_{T_{k} \in \alpha_{i}} u_{k}\right)-\left(m-\left|\alpha_{i}\right|\right)\right) \\
= & \left(t_{q_{i}}-t_{p_{i}}\right) \sum_{T_{k} \in \alpha_{i}}\left(1-u_{k}\right) \tag{11}
\end{align*}
$$

For each task $T_{k}$ in $D_{C I}$, the sum of the lengths of the intervals $\left[t_{p_{i}}, t_{q_{i}}\right)$ in which the carry-in job of $T_{k}$ executes continuously is at most $\delta_{k}$. Thus, summing the allocation
differences for all the intervals $\left[t_{p_{i}}, t_{q_{i}}\right)$ given by (11), we have

$$
\begin{align*}
& A\left(\mathbf{d}, t_{n}, t_{d}, P S\right)-A\left(\mathbf{d}, t_{n}, t_{d}, S\right) \\
\leq & \sum_{i=1}^{z} \sum_{T_{k} \in D_{C I}}\left(t_{q_{i}}-t_{p_{i}}\right)\left(1-u_{k}\right) \\
\leq & \sum_{T_{k} \in D_{C I}} \delta_{k}\left(1-u_{k}\right) \tag{12}
\end{align*}
$$

Setting this value into (10), we get $L A G\left(\mathbf{d}, t_{d}, S\right) \leq$ $L A G\left(\mathbf{d}, t_{n}, S\right)+A\left(\mathbf{d}, t_{n}, t_{d}, P S\right)-A\left(\mathbf{d}, t_{n}, t_{d}, S\right) \leq$ $L A G\left(\mathbf{d}, t_{n}, S\right)+\sum_{T_{k} \in D_{C I}} \delta_{k}\left(1-u_{k}\right)$.

Lemma 2. $\operatorname{lag}\left(T_{i}, t, S\right) \leq u_{i} \cdot x+e_{i}+u_{i} \cdot s_{i}$ for any task $T_{i}$ and $t \in\left[0, t_{d}\right]$.

Proof. Let $d_{i, k}$ be the deadline of the earliest pending job of $T_{i}, T_{i, k}$, in the schedule $S$ at time $t$. If such a job does not exist, then $\operatorname{lag}\left(T_{i}, t, S\right)=0$, and the lemma holds trivially. Let $\gamma_{i}$ be the amount of work $T_{i, k}$ performs before $t$.

By the selection of $T_{i, k}$, we have

$$
\begin{aligned}
\operatorname{lag}\left(T_{i}, t, S\right) & =\sum_{h \geq k} \operatorname{lag}\left(T_{i, h}, t, S\right) \\
& =\sum_{h \geq k}\left(A\left(T_{i, h}, 0, t, P S\right)-A\left(T_{i, h}, 0, t, S\right)\right)
\end{aligned}
$$

Given that no job executes before its release time, $A\left(T_{i, h}, 0, t, S\right)=A\left(T_{i, h}, r_{i, h}, t, S\right)$. Thus,

$$
\begin{align*}
\operatorname{lag}\left(T_{i}, t, S\right)= & A\left(T_{i, k}, r_{i, h}, t, P S\right)-A\left(T_{i, k}, r_{i, k}, t, S\right) \\
& +\sum_{h>i}\left(A\left(T_{i, h}, r_{i, h}, t, P S\right)\right. \\
& \left.-A\left(T_{i, h}, r_{i, h}, t, S\right)\right) \tag{13}
\end{align*}
$$

By the definition of $P S, A\left(T_{i, k}, r_{i, h}, t, P S\right) \leq e_{i}$, and $\sum_{h>k} A\left(T_{i, h}, r_{i, h}, t, P S\right) \leq u_{i} \cdot \max \left(0, t-d_{i, k}\right)$. By the selection of $T_{i, k}, A\left(T_{i, k}, r_{i, k}, t, S\right)=\gamma_{i}$, and $\sum_{h>k} A\left(T_{i, h}, r_{i, h}, t, S\right)=0$. By setting these values into (13), we have

$$
\begin{equation*}
\operatorname{lag}\left(T_{i}, t, S\right) \leq e_{i}-\gamma_{i}+u_{i} \cdot \max \left(0, t-d_{i, k}\right) \tag{14}
\end{equation*}
$$

There are two cases to consider.
Case 1. $d_{i, k} \geq t . \quad$ In this case, (14) implies $\operatorname{lag}\left(T_{i}, t, S\right) \leq e_{i}-\gamma_{i}$, which implies $\operatorname{lag}\left(T_{i}, t, S\right) \leq u_{i} \cdot x+e_{i}+u_{i} \cdot s_{i}$.

Case 2. $d_{i, k}<t$. In this case, because $t \leq t_{d}$ and $d_{l, j}=$ $t_{d}, T_{i, k}$ is not the job $T_{l, j}$. Thus, by Property (P), $T_{i, k}$ has tardiness at most $x+e_{i}+s_{i}$, so $t+e_{i}-\gamma_{i} \leq d_{i, k}+x+e_{i}+s_{i}$. Thus, $t-d_{i, k} \leq x+\gamma_{i}+s_{i}$. Setting this value into (14), we have $\operatorname{lag}\left(T_{i}, t, S\right) \leq u_{i} \cdot x+e_{i}+u_{i} \cdot s_{i}$.


[^0]:    *Work supported by AT\&T, IBM, Intel, and Sun Corps.; NSF grants CNS 0834270, CNS 0834132, and CNS 0615197; and ARO grant W911NF-06-1-0425.

[^1]:    ${ }^{1}$ As time increases from $t_{s}$ to $t_{p}$, whenever a presence interval ends on processor $k$, a task exists that can be used to define the next presence interval on processor $k$. It can also be assumed, without loss of generality, that this task executes only on processor $k$ during its presence interval.

[^2]:    ${ }^{2}$ If $C^{\prime}$ is of length $\epsilon$ and is followed by a suspension phase, then we can avoid altering the length of that suspension phase by assuming that $C^{\prime}$ executes for zero time at time $t+\epsilon$. Note that $C^{\prime}$ 's execution time will be increased in a subsequent induction step. A similar comment applies to the argument in the next paragraph.

