

Deferrable Scheduling for Maintaining Real-Time Data Freshness: Algorithms, Analysis, and Results

Ming Xiong, *Member, IEEE*, Song Han, *Student Member, IEEE*,
Kam-Yiu Lam, and Deji Chen, *Member, IEEE*

Abstract—The periodic update transaction model has been used to maintain the freshness (or temporal validity) of real-time data. Period and deadline assignment has been the main focus of past studies, such as the More-Less scheme [25], in which update transactions are guaranteed by the Deadline Monotonic scheduling algorithm [16] to complete by their deadlines. In this paper, we propose a *deferrable scheduling* algorithm for fixed-priority transactions, a novel approach for minimizing update workload while maintaining the temporal validity of real-time data. In contrast to prior work on maintaining data freshness periodically, update transactions follow an aperiodic task model in the deferrable scheduling algorithm. The deferrable scheduling algorithm exploits the semantics of temporal validity constraint of real-time data by judiciously deferring the sampling times of update transaction jobs as late as possible. We present a theoretical estimation of its processor utilization and a sufficient condition for its schedulability. Our experimental results verify the theoretical estimation of the processor utilization. We demonstrate through the experiments that the deferrable scheduling algorithm is an effective approach and it significantly outperforms the More-Less scheme in terms of reducing processor workload.

Index Terms—Deferrable scheduling, real-time databases, temporal validity, fixed-priority scheduling.

1 INTRODUCTION

REAL-TIME and embedded systems are applied in many application domains that require timely processing of a massive amount of real-time data. Examples of real-time data include sensor data in sensor networks, positions of aircraft in air traffic control systems [14], and vehicle velocity in adaptive cruise control applications [6]. Such real-time data are typically managed in a real-time database system (RTDBS). Those data values are used to model the current status of entities in a system environment. However, real-time data are different from traditional data in that they have time semantics in which sampled values are valid only for a certain time interval [19], [18], [23]. The concept of *temporal validity* is used to define the correctness of real-time data [19]. A real-time data object is *fresh* (or *temporally valid*) if its value truly reflects the current status of the corresponding entity in the system environment. Each real-time data object is associated with a *validity interval* as the lifespan of the current data value defined

based on the dynamic properties of the data object. A new data value needs to be installed into the database before the validity interval of the old value expires, that is, the old one becomes temporally invalid. Otherwise, the RTDBS cannot detect and respond to environmental changes in a timely manner. In recent years, there has been a tremendous amount of work devoted to this area [5], [1], [12], [14], [30], [19], [20], [21], [22], [26], [11], [25], [8].

To maintain temporal validity, *sensor update transactions*, which capture the latest status of the entities in the system environment, are generated to refresh the values of the real-time data periodically [19], [14], [25]. A sensor update transaction has an infinite number of periodic jobs, which have fixed-length periods and relative deadlines. The update problem for periodic update transactions consists of two parts [25]: 1) *the determination of the sampling periods and deadlines of update transactions* and 2) *the scheduling of update transactions*. Prior work has proposed two approaches for minimizing the update workload while maintaining real-time data freshness. As explained in [19], [14], a simple method for maintaining the temporal validity of real-time data is to use the *Half-Half (HH)* scheme in which the update period for a real-time data object is set to be half the validity interval of the object. To further reduce the update workload, the *More-Less (ML)* scheme is proposed and studied in [2], [25].

This paper presents *Deferrable Scheduling for Fixed-Priority transactions (DS-FP)*, a novel algorithm for maintaining real-time data freshness, with the objective being to minimize the update workload [27], [28]. We study the problem of data freshness maintenance for firm real-time update transactions in a single-processor RTDBS. Distinct from the past work of *HH* and *ML*, which have a fixed

- M. Xiong is with Bell Labs, Alcatel-Lucent, 600 Mountain Avenue, Room 2D-512, Murray Hill, NJ 07974. E-mail: xiong@research.bell-labs.com.
- S. Han is with the Department of Computer Sciences, University of Texas at Austin, Austin, TX 78712. E-mail: shan@cs.utexas.edu.
- K.-Y. Lam is with the Department of Computer Science, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon, Hong Kong. E-mail: cskylan@cityu.edu.hk.
- D. Chen is with Emerson Process Management, 12301 Research Blvd., Research Park Plaza, Bldg. III, Austin, TX 78759. E-mail: deji.chen@emerson.com.

Manuscript received 14 Nov. 2006; revised 1 Oct. 2007; accepted 18 Dec. 2007; published online 9 Jan. 2008.

Recommended for acceptance by A. Zomaya.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TC-0440-1106.

Digital Object Identifier no. 10.1109/TC.2008.16.

period and relative deadline for each transaction, *DS-FP* adopts an *a*periodic task model. In contrast to *ML*, in which a relative deadline is always equivalent to the worst-case response time of a transaction, *DS-FP* dynamically assigns relative deadlines to transaction jobs by deferring the sampling time of a transaction job as much as possible while still guaranteeing the temporal validity of real-time data. The deferral of a job's sampling time results in a shorter relative deadline than its worst-case response time, which, in turn, increases the separation of two consecutive jobs. Thus, the deferral of sampling time lends itself to a reduced processor workload produced by update transactions. We prove that *DS-FP* outperforms *ML* in terms of schedulability and present a sufficient condition for the schedulability of a set of transactions under *DS-FP*. We also analyze the average processor utilization under *DS-FP*. Our experimental study of *DS-FP* demonstrates that it is an effective algorithm for reducing the workload of real-time update transactions. It also verifies the accuracy of our theoretical estimation of average processor utilization under *DS-FP* and demonstrates the effectiveness of the *DS-FP* algorithms.

The rest of this paper is organized as follows: Section 2 reviews the existing approaches for real-time data freshness maintenance. In Section 3, we propose the *DS-FP* algorithm. Our detailed discussion on *DS-FP* includes an analysis of its schedulability and nonoptimality, as well as an estimation of its average processor utilization. Section 4 presents the performance studies and Section 5 briefly describes the related work. Finally, we conclude our study in Section 6 and present open questions for *DS-FP*.

2 BACKGROUND: DATA FRESHNESS MAINTENANCE

Real-time data, whose state may become invalid with the passage of time, need to be refreshed by sensor update transactions generated by intelligent sensors that sample the values of real-world entities. To monitor the states of entities faithfully, real-time data must be refreshed before they become invalid. The actual length of the temporal validity interval of a real-time data object is application dependent. For example, real-time data with validity interval requirements are discussed in [19], [20], [18]. One of the important design goals of RTDBSs is to guarantee that real-time data remain fresh, that is, they are always valid.

2.1 Temporal Validity for Data Freshness

As real-time data values change continuously with time, the correctness of a real-time data object X_i depends on the difference between the real-time status $S(E_i)$ of the real-world entity E_i and the current sampling value $Val(X_i)$ of X_i .

Definition 2.1. A real-time data object X_i at time t is temporally valid (or temporally consistent) if (for its update job, $J_{i,j}$, finished last before t) the sampling time $r_{i,j}$ plus the validity interval length (or validity length) \mathcal{V}_i of the data object is not less than t , that is, $r_{i,j} + \mathcal{V}_i \geq t$ [21], [19], [1].

A data value for real-time data object X_i sampled at any time t will be valid for \mathcal{V}_i following that t up to $(t + \mathcal{V}_i)$. Next, we review existing approaches that adopt a periodic task model for sensor update transactions.

TABLE 1
Symbols and Definitions

Symbol	Definition
X_i	Real-time data object i
τ_i	Update transaction updating X_i
$J_{i,j}$	The $(j + 1)^{th}$ job of τ_i ($i = 1, \dots, m, j = 0, 1, 2, \dots$)
$R_{i,j}$	Response time of $J_{i,j}$
C_i	Computation time of transaction τ_i
\mathcal{V}_i	Validity (interval) length of X_i
$f_{i,j}$	Finishing (Completion) time of $J_{i,j}$
$r_{i,j}$	Release (Sampling) time of $J_{i,j}$
$d_{i,j}$	Absolute deadline of $J_{i,j}$
P_i	Period of transaction τ_i in <i>ML</i>
D_i	Relative deadline of transaction τ_i in <i>ML</i>
$P_{i,j}$	Separation of jobs (i.e., $r_{i,j+1} - r_{i,j}$) in <i>DS-FP</i>
$D_{i,j}$	Relative deadline of $J_{i,j}$ in <i>DS-FP</i>
\bar{P}_i	Average period of transaction τ_i in <i>DS-FP</i>
\bar{D}_i	Average relative deadline of transaction τ_i in <i>DS-FP</i>
\bar{U}_{DS}	Average processor utilization in <i>DS-FP</i>
$\Theta_i(a, b)$	Total cumulative processor demands from higher-priority transactions received by τ_i in interval $[a, b)$

2.2 Half-Half and More-Less

In this section, traditional approaches for maintaining temporal validity, namely, the *Half-Half* (*HH*) and *More-Less* (*ML*) approaches, are reviewed.

In this paper, $\mathcal{T} = \{\tau_i\}_{i=1}^m$ refers to a set of periodic update transactions $\{\tau_1, \tau_2, \dots, \tau_m\}$, and $\mathcal{X} = \{X_i\}_{i=1}^m$ refers to a set of real-time data objects. We assume that τ_i has a higher priority than τ_j for $i < j$, unless specified otherwise. All real-time data objects are assumed to be kept in the main memory. Associated with X_i ($1 \leq i \leq m$) is a validity interval of length \mathcal{V}_i : Transaction τ_i ($1 \leq i \leq m$) updates the corresponding data object X_i . Because each update transaction updates a different data object, no concurrency control is considered for update transactions. We assume that a sensor always samples the value of a real-time data object at the beginning of its period and the system is *synchronous* (that is, all of the first jobs of update transactions are initiated at the same time), unless stated otherwise. For convenience, let $d_{i,j}$, $f_{i,j}$, and $r_{i,j}$ denote the absolute deadline, completion (finishing) time, and sampling (release) time of job $J_{i,j}$ of τ_i , respectively. We also assume that jitter between the sampling time and the release time of a job is zero for convenience of presentation (readers are referred to Section 3.3 for how jitters can be handled). Formal definitions of the frequently used symbols are given in Table 1. Deadlines of update transactions are firm deadlines. The goal of *HH* and *ML*, which adopt a *periodic* task model, is to determine period P_i and relative deadline D_i so that all of the update transactions are schedulable and the CPU workload resulting from periodic update transactions is minimized.

Both *HH* and *ML* assume a simple execution semantics for periodic transactions: A transaction must be executed once every period. However, there is no guarantee as to when a job of a periodic transaction is actually executed within a period. Throughout this paper, we assume that the scheduling algorithms are preemptive and we ignore all

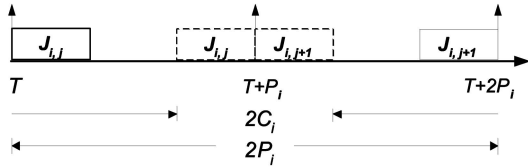


Fig. 1. Extreme execution cases of $J_{i,j}$ and $J_{i,j+1}$.

preemption overhead. For convenience, we use the terms transaction and task interchangeably.

Half-Half. In *HH*, the period and relative deadline of an update transaction are each typically set to be half the data validity length [19], [14]. In Fig. 1, the farthest distance of two consecutive jobs of τ_i (based on the sampling time $r_{i,j}$ of job $J_{i,j}$ and the deadline $d_{i,j+1}$ of its next job) is $2P_i$. If $2P_i \leq \mathcal{V}_i$, then the validity of real-time object X_i is guaranteed as long as jobs of τ_i meet their deadlines. Unfortunately, this approach incurs an unnecessarily high CPU workload of update transactions in the RTDBSs compared to *More-Less*.

More-Less. Consider the worst-case response time for any job of a periodic transaction τ_i , where the response time is the difference between the transaction initiation time ($I_i + KP_i$) and the transaction completion time, where I_i is the offset within the period and K is a natural number.

Lemma 2.1. For a set of periodic transactions $\mathcal{T} = \{\tau_i\}_{i=1}^m$ ($D_i \leq P_i$) with transaction initiation time ($I_i + KP_i$) ($K = 0, 1, 2, \dots$), the worst-case response time for any job of τ_i occurs for the first job of τ_i when $I_1 = I_2 = \dots = I_m = 0$ [16].

For $I_i = 0$ ($1 \leq i \leq m$), the transactions are *synchronous*. A time instant after which a transaction has the worst-case response time is called a *critical instant*. For example, time 0 is a critical instant for all of the transactions if those transactions are *synchronous*.

To minimize the update workload and guarantee temporal validity, *ML* uses *Deadline Monotonic (DM)* [16] to schedule periodic update transactions [2], [25]. There are three constraints to follow for τ_i ($1 \leq i \leq m$):

- **Validity constraint.** The sum of the period and the relative deadline of transaction τ_i is always less than or equal to \mathcal{V}_i , that is, $P_i + D_i \leq \mathcal{V}_i$, as shown in Fig. 2.
- **Deadline constraint.** The period of an update transaction is assigned to be more than half the validity length of the object to be updated, while its corresponding relative deadline is less than half the validity length of the same object. For τ_i to be schedulable, D_i must be greater than or equal to C_i , the worst-case execution time of τ_i , that is, $C_i \leq D_i \leq P_i$.
- **Schedulability constraint.** For a given set of update transactions, the *DM* scheduling algorithm [16] is used to schedule the transactions. Consequently, $\sum_{j=1}^i (\lceil \frac{D_i}{P_j} \rceil \cdot C_j) \leq D_i$ ($1 \leq i \leq m$) if τ_j has a higher priority than τ_i for $i > j$.

ML assigns priorities to transactions based on *Shortest Validity First (SVF)*, that is, in the *inverse* order of validity

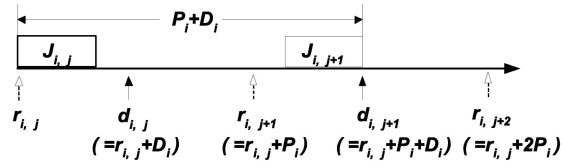


Fig. 2. Illustration of the *ML* scheme.

length, and ties are resolved in favor of transactions with less slack (that is, $\mathcal{V}_i - C_i$ for τ_i). It assigns deadlines and periods to τ_i as follows:

$$D_i = f_{i,0}^{ml} - r_{i,0}^{ml}, \quad (1)$$

$$P_i = \mathcal{V}_i - D_i, \quad (2)$$

where $f_{i,0}^{ml}$ and $r_{i,0}^{ml}$ are the finishing and sampling times of the first job of τ_i under *ML*, respectively. Note that, in a synchronous system, $r_{i,0}^{ml} = 0$ and the first job's response time is the worst-case response time in *ML*. In this paper, superscript *ml* is used to distinguish the finishing and sampling times in *ML* from those in *DS-FP*.

3 DEFERRABLE SCHEDULING

All schedulers discussed in this paper are *work-conserving* for released jobs. In other words, the scheduler never idles the processor while there is a job awaiting execution (that is, after it is released). Next, we introduce the *Deferrable Scheduling algorithm for Fixed Priority* transactions (*DS-FP*). Section 3.1 presents the intuition of the algorithm and Section 3.2 describes the details of the algorithm. Section 3.3 compares it with *ML*. Section 3.4 provides an estimation of *DS-FP*'s average processor utilization. Section 3.5 discusses whether the algorithm is optimal.

3.1 Intuition of DS-FP

In *ML*, D_i is determined by the first job's response time, which is the *worst-case* response time of all of the jobs of τ_i . Thus, *ML* is pessimistic on the deadline and period assignment in the sense that it uses a periodic task model that has a fixed period and relative deadline for each task, and the relative deadline is equivalent to the worst-case response time. It should be noted that the *validity constraint* can always be satisfied as long as $P_i + D_i \leq \mathcal{V}_i$. However, the processor workload is minimized only if $P_i + D_i = \mathcal{V}_i$. Otherwise, P_i can always be increased to reduce processor workload as long as $P_i + D_i < \mathcal{V}_i$. Given the release time $r_{i,j}$ of job $J_{i,j}$ and the deadline $d_{i,j+1}$ of job $J_{i,j+1}$ ($j \geq 0$),

$$d_{i,j+1} \leq r_{i,j} + \mathcal{V}_i \quad (3)$$

guarantees that the *validity constraint* can be satisfied, as depicted in Fig. 3. Correspondingly, the following equation follows directly from (3):

$$(r_{i,j+1} - r_{i,j}) + (d_{i,j+1} - r_{i,j+1}) \leq \mathcal{V}_i. \quad (4)$$

If $r_{i,j+1}$ is shifted onward to $r'_{i,j+1}$ along the time line in Fig. 3, it does not violate (4) and $J_{i,j+1}$ can still be completed by its deadline. This shift can be achieved, for example, in the *ML* schedule if preemption to $J_{i,j+1}$ from higher priority

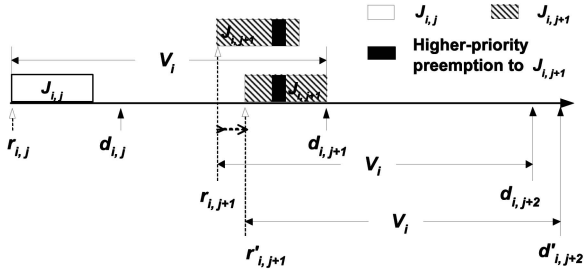


Fig. 3. Illustration of *DS-FP* scheduling ($r_{i,j+1}$ is shifted to $r'_{i,j+1}$).

transactions in $[r_{i,j+1}, d_{i,j+1}]$ is less than the worst-case preemption to the first job of τ_i . Thus, temporal validity can still be guaranteed as long as $J_{i,j+1}$ is completed by its deadline $d_{i,j+1}$.

The intention of *DS-FP* is to defer the sampling time $r_{i,j+1}$ of $J_{i,j}$'s subsequent job as late as possible while still guaranteeing the *validity constraint*. Note that the sampling time of a job is also its release time, that is, the time that the job is ready to execute, as we assume zero cost for sampling and no arrival jitter for a job for convenience of presentation.

The deferral of job $J_{i,j+1}$'s release time reduces the relative deadline of the job if its absolute deadline is fixed as in (3). For example, although $r_{i,j+1}$ is deferred to $r'_{i,j+1}$ in Fig. 3, it still has to be completed by its deadline $d_{i,j+1}$ in order to satisfy the *validity constraint* (3). Thus, its relative deadline $D_{i,j+1}$ becomes $d_{i,j+1} - r'_{i,j+1}$, which is less than $d_{i,j+1} - r_{i,j+1}$. The deadline of $J_{i,j+1}$'s subsequent job $J_{i,j+2}$ can be further deferred to $(r'_{i,j+1} + \mathcal{V}_i)$ to satisfy the *validity constraint*. Consequently, the processor utilization for the completion of three jobs $J_{i,j}$, $J_{i,j+1}$, and $J_{i,j+2}$ then becomes $\frac{3C_i}{2\mathcal{V}_i - (d_{i,j+1} - r'_{i,j+1})}$. It is less than the utilization $\frac{3C_i}{2\mathcal{V}_i - (d_{i,j+1} - r_{i,j+1})}$ required for the completion of the same amount of work in *ML*.

Definition 3.1. Let $\Theta_i(a, b)$ denote the total cumulative processor demands made by all jobs of a higher priority transaction τ_j for $\forall j$ ($1 \leq j \leq i-1$) during the time interval $[a, b]$ from a schedule \mathcal{S} produced by a fixed-priority scheduling algorithm. Then,

$$\Theta_i(a, b) = \sum_{j=1}^{i-1} \theta_j(a, b),$$

where $\theta_j(a, b)$ is the total processor demands made by all jobs of a single transaction τ_j during $[a, b]$.

Next, we discuss how much a job's release time can be deferred. We shall use $r_{i,j+1}$ instead of $r'_{i,j+1}$ to denote the final deferred release time. According to the fixed-priority scheduling theory, $r_{i,j+1}$ can be derived backward from its deadline $d_{i,j+1}$ as follows:

$$r_{i,j+1} = d_{i,j+1} - R_{i,j+1}(r_{i,j+1}, d_{i,j+1}), \quad (5)$$

$$R_{i,j+1}(r_{i,j+1}, d_{i,j+1}) = \Theta_i(r_{i,j+1}, d_{i,j+1}) + C_i, \quad (6)$$

where $R_{i,j+1}(r_{i,j+1}, d_{i,j+1})$ denotes the response time of $J_{i,j+1}$ in the time interval $[r_{i,j+1}, d_{i,j+1}]$. Note that the schedule of all higher priority jobs that are released prior to $d_{i,j+1}$ needs

TABLE 2
Parameters and Results for Example 3.1

i	C_i	\mathcal{V}_i	$\frac{C_i}{\mathcal{V}_i}$	<i>ML</i>		<i>Half-Half</i>
				P_i	D_i	$P_i(D_i)$
1	1	5	0.2	4	1	2.5
2	2	10	0.2	7	3	5
3	2	20	0.1	14	6	10

to be computed before $\Theta_i(r_{i,j+1}, d_{i,j+1})$ is computed. This computation can be invoked using a recursive process from jobs of lower priority transactions to higher priority transactions. Nevertheless, it does not require that a *complete* schedule of all jobs should be constructed offline before the task set is executed. Indeed, the computation of job deadlines and their corresponding release times is performed online while the transactions are being scheduled. We only need to compute the first jobs' response times when the system starts. Upon the completion of job $J_{i,j}$, the deadline of its next job $d_{i,j+1}$ is first derived from (3) and, then, the corresponding release time $r_{i,j+1}$ is derived from (5). If $\Theta_i(r_{i,j+1}, d_{i,j+1})$ cannot be computed due to incomplete schedule information of release times and absolute deadlines from higher priority transactions, *DS-FP* computes their schedule information online until it can gather enough information to derive $r_{i,j+1}$. Job $J_{i,j}$'s *DS-FP* scheduling information (for example, release time, deadline, and bookkeeping information) can be discarded after it is completed and no lower priority transactions need its information for deriving their schedules. This process is called *garbage collection* in *DS-FP*.

Let $S_J(t)$ denote the set of jobs of all transactions whose deadlines have been computed by time t . Also, let $LSD_i(t)$ denote the latest scheduled deadline of τ_i at t , that is, the maximum of all $d_{i,j}$ for jobs $J_{i,j}$ of τ_i whose deadlines have been computed by t . Then,

$$LSD_i(t) = \max_{J_{i,j} \in S_J(t)} \{d_{i,j}\} \quad (j \geq 0). \quad (7)$$

Given job $J_{k,j}$, whose scheduling information has been computed at time t , and $\forall i$ ($i > k$), if

$$LSD_i(t) \geq d_{k,j}, \quad (8)$$

then the information of $J_{k,j}$ can be garbage collected.

Example 3.1. Suppose that there are three update transactions whose parameters are shown in Table 2. The resulting periods and deadlines in *HH* and *ML* are shown in the same table. The utilizations of *HH* and *ML* are $U_{ml} \approx 0.68$ and $U_{hh} = 1.00$, respectively. Figs. 4a and 4b depict the schedules produced by *ML* and *DS-FP*, respectively. It can be observed from both schedules that the release times of transaction jobs $J_{3,1}$, $J_{2,3}$, $J_{2,4}$ are shifted from times 14, 21, and 28 in *ML* to 18, 22, and 30 in *DS-FP*, respectively.

The *DS-FP* algorithm is described in Section 3.2.

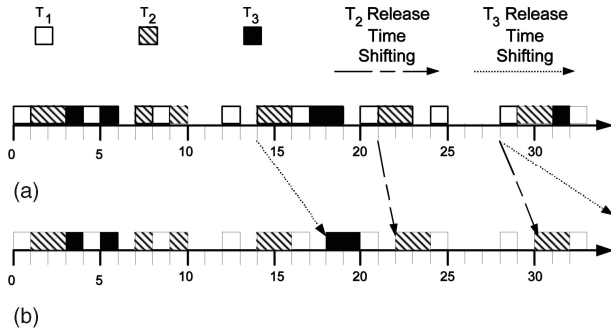


Fig. 4. Comparing ML and DS-FP schedules. (a) An ML schedule. (b) A DS-FP schedule.

3.2 Deferrable Scheduling Algorithm

This section presents *DS-FP*, a *fixed-priority* scheduling algorithm. Transaction priority assignment policy in *DS-FP* is the same as in *ML*, that is, *SFV*. Given an update transaction set \mathcal{T} , it is assumed that τ_i has a priority higher than τ_j if $i < j$, as we let $\mathcal{V}_i \leq \mathcal{V}_j$. Algorithm 3.1 presents the *DS-FP* algorithm. For convenience of presentation, garbage collection is omitted in the algorithm. There are two cases for the *DS-FP* algorithm: 1) At the system initialization time, lines 13-20 iteratively calculate the first job's response time for τ_i and the first job's deadline is set as its response time (line 21) and 2) upon the completion of τ_i 's job $J_{i,k}$ ($1 \leq i \leq m, k \geq 0$), the deadline of its next job $J_{i,k+1}$, $d_{i,k+1}$, is derived at line 27 so that the farthest distance of $J_{i,k}$'s sampling time and $J_{i,k+1}$'s finishing time is bounded by the validity length \mathcal{V}_i (3). Finally, the sampling time of $J_{i,k+1}$, $r_{i,k+1}$, is derived backward from its deadline by accounting for the interferences from higher priority transactions (line 29).

Algorithm 3.1: the *DS-FP* algorithm

Input: a set of update transactions $\mathcal{T} = \{\tau_i\}_{i=1}^m$ ($m \geq 1$) with known $\{C_i\}_{i=1}^m$ and $\{\mathcal{V}_i\}_{i=1}^m$.

Output: construct a partial schedule \mathcal{S} if \mathcal{T} is feasible; otherwise, reject.

```

1 case (|system initialization time|):
2  $t \leftarrow 0$ ; // Initialization
3 //  $LSD_i$ —Latest Scheduled Deadline of  $\tau_i$ 's jobs.
4  $LSD_i \leftarrow 0, \forall i (1 \leq i \leq m)$ ;
5  $\ell_i \leftarrow 0, \forall i (1 \leq i \leq m)$ ;
6 //  $\ell_i$  is the latest scheduled job of  $\tau_i$ 
7 for  $i = 1$  to  $m$  do
8 // Schedule finish time for  $\tau_i$ .
9  $r_{i,0} \leftarrow 0$ ;
10  $f_{i,0} \leftarrow C_i$ ;
11 // Calculate higher priority (HP) preemptions.
12  $oldHPPreempt \leftarrow 0$ ; // initial HP preemptions
13  $hpPreempt \leftarrow CalcHPPreempt(i, 0, 0, f_{i,0})$ ;
14 while ( $hpPreempt > oldHPPreempt$ ) do
15 // Accounting for the interferences of HP tasks
16  $f_{i,0} \leftarrow r_{i,0} + hpPreempt + C_i$ ;
17 if ( $f_{i,0} > \mathcal{V}_i - C_i$ ) then abort fi;
18  $oldHPPreempt \leftarrow hpPreempt$ ;
19  $hpPreempt \leftarrow CalcHPPreempt(i, 0, 0, f_{i,0})$ ;
20 od
21  $d_{i,0} \leftarrow f_{i,0}$ ;

```

```

22 od
23 return;
24
25 case (upon the completion of  $J_{i,k}$ ):
26 //Schedule release time for  $J_{i,k+1}$ .
27  $d_{i,k+1} \leftarrow r_{i,k} + \mathcal{V}_i$ ; // get the next deadline for  $J_{i,k+1}$ 
28 //  $r_{i,k+1}$  is also the sampling time for  $J_{i,k+1}$ 
29  $r_{i,k+1} \leftarrow ScheduleRT(i, k+1, C_i, d_{i,k+1})$ ;
30 return;

```

Algorithm 3.2: *ScheduleRT*($i, k, C_i, d_{i,k}$)

Input: $J_{i,k}$, with C_i and $d_{i,k}$.

Output: $r_{i,k}$.

```

1  $oldHPPreempt \leftarrow 0$ ; // initial HP preemptions
2  $hpPreempt \leftarrow 0$ ;
3  $r_{i,k} \leftarrow d_{i,k} - C_i$ ;
4 // Calculate HP preemptions backwards from  $d_{i,k}$ .
5  $hpPreempt \leftarrow CalcHPPreempt(i, k, r_{i,k}, d_{i,k})$ ;
6 while ( $hpPreempt > oldHPPreempt$ ) do
7 // Accounting for the interferences of HP tasks
8  $r_{i,k} \leftarrow d_{i,k} - hpPreempt - C_i$ ;
9 if ( $r_{i,k} < d_{i,k-1}$ ) then abort fi;
10  $oldHPPreempt \leftarrow hpPreempt$ ;
11  $hpPreempt \leftarrow GetHPPreempt(i, k, r_{i,k}, d_{i,k})$ ;
12 od
13 return  $r_{i,k}$ ;

```

Algorithm 3.3: *CalcHPPreempt*(i, k, t_1, t_2)

Input: $J_{i,k}$, and a time interval $[t_1, t_2]$.

Output: total cumulative processor demands from HP transactions τ_j ($1 \leq j \leq i-1$) during $[t_1, t_2]$.

```

1  $\ell_i \leftarrow k$ ; // Record the latest scheduled job of  $\tau_i$ .
2  $d_{i,k} \leftarrow t_2$ ;
3  $LSD_i \leftarrow t_2$ ;
4 if ( $i = 1$ )
5 then // No preemptions from HP tasks.
6 return 0;
7 elsif ( $LSD_{i-1} \geq LSD_i$ )
8 then // Get preemptions from  $\tau_j$  ( $\forall j, 1 \leq j < i$ )
9 // because  $\tau_j$ 's schedule is complete before  $t_2$ .
10 return  $GetHPPreempt(i, k, t_1, LSD_i)$ ;
11 fi
12 //build  $\mathcal{S}$  up to or exceeding  $t_2$  for  $\tau_j (1 \leq j < i)$ .
13 for  $j = 1$  to  $i-1$  do
14 while ( $d_{j,\ell_j} < LSD_i$ ) do
15  $d_{j,\ell_j+1} \leftarrow r_{j,\ell_j} + \mathcal{V}_j$ ;
16  $r_{j,\ell_j+1} \leftarrow ScheduleRT(j, \ell_j + 1, C_j, d_{j,\ell_j+1})$ ;
17  $\ell_j \leftarrow \ell_j + 1$ ;
18  $LSD_j \leftarrow d_{j,\ell_j}$ ;
19 od
20 end
21 return  $GetHPPreempt(i, k, t_1, LSD_i)$ ;

```

Function *ScheduleRT*($i, k, C_i, d_{i,k}$) (Algorithm 3.2) calculates the release time $r_{i,k}$ with known computation time C_i and deadline $d_{i,k}$. It starts with release time $r_{i,k} = d_{i,k} - C_i$, then iteratively calculates $\Theta_i(r_{i,k}, d_{i,k})$, which is the total cumulative processor demands made by all higher priority jobs of $J_{i,k}$ during the interval $[r_{i,k}, d_{i,k})$, and adjusts $r_{i,k}$ by

accounting for the interferences from higher priority transactions (lines 5 to 12). The computation of $r_{i,k}$ continues until the interferences from higher priority transactions do not change in an iteration. In particular, line 9 detects any infeasible schedule. A schedule becomes infeasible under *DS-FP* if $r_{i,k} < d_{i,k-1}$ ($k > 0$), that is, the release time of $J_{i,k}$ becomes earlier than the deadline of its preceding job $J_{i,k-1}$. Function $GetHPPreempt(i, k, t_1, t_2)$ scans the interval $[t_1, t_2)$, adds up the total preemptions from $\tau_j (\forall j, 1 \leq j \leq i-1)$, and returns $\Theta_i(t_1, t_2)$, the cumulative processor demands of τ_j during $[t_1, t_2)$ from schedule \mathcal{S} that has been built.

Function $CalcHPPreempt(i, k, t_1, t_2)$ (Algorithm 3.3) calculates $\Theta_i(t_1, t_2)$, the total cumulative processor demands made by all higher priority jobs of $J_{i,k}$ during the interval $[t_1, t_2)$. Line 7 ensures that $(\forall j, 1 \leq j < i)$ τ_j 's schedule is completely built before time t_2 . This is because τ_i 's schedule cannot be completely built before t_2 unless the schedules of its higher priority transactions are complete before t_2 . In this case, the function simply returns an amount of higher priority preemptions for τ_i during $[t_1, t_2)$ by invoking $GetHPPreempt(i, k, t_1, t_2)$, which returns $\Theta_i(t_1, t_2)$. If any higher priority transaction τ_j ($j < i$) does not have a complete schedule during $[t_1, t_2)$, its schedule \mathcal{S} up to or exceeding t_2 is built on the fly (lines 14-19). This enables the computation of $\Theta_i(t_1, t_2)$. The latest scheduled deadline of τ_i 's job LSD_i indicates the latest deadline of τ_i 's jobs that have been computed.

The *worst-case* complexity of $ScheduleRT$ is $O(m \cdot \mathcal{V}_m^2)$, assuming that $\frac{\mathcal{V}_m}{\mathcal{V}_1}$ is a constant. An important property of $ScheduleRT(i, k, C_i, d_{i,k})$ terminating at time $t = d_{i,k}$ is that the latest scheduled deadline of τ_l ($LSD_l(t)$) is not larger than that of τ_j ($LSD_j(t)$) if τ_l does not have a priority higher than τ_j ($l \geq j$). This is proven in the following lemma.

Lemma 3.1. *Given a synchronous update transaction set \mathcal{T} and $ScheduleRT(i, k, C_i, t)$ ($1 \leq i \leq m \& k \geq 0$), $LSD_l(t) \leq LSD_j(t)$ ($i \geq l \geq j$) holds when $ScheduleRT(i, k, C_i, t)$ terminates at time t .*

Proof. This can be proven by contradiction. Suppose that $LSD_l(t) > LSD_j(t)$ ($i \geq l \geq j$) when $ScheduleRT(i, k, C_i, t)$ terminates at t . If $LSD_l(t) < t$, then $CalcHPPreempt(i, k, t_1, t_2)$ does not terminate according to line 14 because $d_{l,t_1} < LSD_l(t) = t$. Thus, $LSD_l(t) \geq LSD_l(t) = t$. Let $LSD_l(t) = t_2$ in $CalcHPPreempt(l, k_l, t_1, t_2)$, which must be invoked before $ScheduleRT(i, k, C_i, t)$ terminates at t . As we assume that $LSD_j(t) < LSD_l(t) = t_2$, similarly, $CalcHPPreempt(l, k_l, t_1, t_2)$ has not reached the point to terminate according to line 14. This contradicts the assumption. \square

The next example illustrates how the *DS-FP* algorithm works with the transaction set in Example 3.1.

Example 3.2. Table 3 presents the comparison of (release time, deadline) pairs assigned by *ML* and *DS-FP* (Algorithm 3.1) for the jobs of τ_1, τ_2 , and τ_3 in Example 3.1. Note that only release times and deadlines before time 40 are depicted in the table. Note also that τ_1 has the same release times and deadlines for all jobs under *ML* and *DS-FP*.

TABLE 3
Release Time and Deadline Comparison

Job	τ_1	τ_2		τ_3	
	<i>ML/DS-FP</i>	<i>ML</i>	<i>DS-FP</i>	<i>ML</i>	<i>DS-FP</i>
0	(0,1)	(0, 3)	(0, 3)	(0, 6)	(0, 6)
1	(4,5)	(7, 10)	(7, 10)	(14, 20)	(18, 20)
2	(8,9)	(14, 17)	(14, 17)	(28, 34)	(35, 38)
3	(12,13)	(21, 24)	(22, 24)
4	(16,17)	(28, 31)	(30, 32)		
5	(20,21)	(35, 38)	(38, 40)		
6	(24,25)		
7	(28,29)				
8	(32,33)				
9	(36,37)				

However, $J_{2,3}, J_{2,4}, J_{2,5}, J_{3,1}$, and $J_{3,2}$ have different release times and deadlines under *ML* and *DS-FP*. Algorithm 3.1 starts at the *system initialization* time. It calculates deadlines for $J_{1,0}, J_{2,0}$, and $J_{3,0}$. Upon the completion of $J_{3,0}$ at time 6, $d_{3,1}$ is set to $r_{3,0} + \mathcal{V}_3 = 20$. Then, Algorithm 3.1 invokes $ScheduleRT(3, 1, 2, 20)$ at line 29, which derives $r_{3,1}$. At this moment, Algorithm 3.1 has already calculated the complete schedule up to $d_{3,0}$ (time 6), but the schedule in the interval $(6, 20]$ has only been partially derived. Specifically, only the schedule information of $J_{1,0}, J_{1,1}, J_{1,2}, J_{1,3}, J_{2,0}$, and $J_{2,1}$ has been derived for τ_1 and τ_2 . Algorithm 3.2 ($ScheduleRT$) obtains $r_{3,1} = 20 - 2 = 18$ at line 3 and then invokes $CalcHPPreempt(3, 1, 18, 20)$. Algorithm 3.3 ($CalcHPPreempt$) finds that $LSD_2 = 10 < t_2 = 20$ and, then, it jumps to the *for* loop starting at line 13 to build the complete schedule of τ_1 and τ_2 in the interval $(6, 20]$, where the release times and deadlines for $J_{1,4}, J_{1,5}, J_{2,2}, J_{1,6}$, and $J_{2,3}$ are derived. Thus, higher priority transactions τ_1 and τ_2 have a complete schedule before time 20. Note that $r_{1,6}$ and $d_{1,6}$ for $J_{1,6}$ are derived when we calculate $r_{2,3}$ and $d_{2,3}$ such that the complete schedule up to time $d_{2,3}$ is built for transactions with priorities higher than τ_2 . As $r_{2,2}$ is set to 14 by earlier calculation, $d_{2,3}$ is set to 24. It derives $r_{2,3}$ backward from $d_{2,3}$ and sets it to 22 because $\Theta_2(22, 24) = 0$. Similarly, $d_{3,1}$ and $r_{3,1}$ are set to 20 and 18, respectively.

3.3 Comparison of *DS-FP* and *ML*

Note that *ML* is based on the *periodic* task model, while *DS-FP* adopts the *aperiodic* task model. The relative deadline of a transaction in *DS-FP* is not fixed. Theoretically, the separation of two consecutive jobs of τ_i in *DS-FP* $r_{i,j} - r_{i,j-1}$ satisfies the following condition:

$$\mathcal{V}_i - C_i \geq r_{i,j} - r_{i,j-1} \geq \mathcal{V}_i - WCRT_i \quad (j \geq 1), \quad (9)$$

where $WCRT_i$ is the worst-case response time of the jobs of τ_i in *DS-FP*. Note that the maximal separation of $J_{i,j}$ and $J_{i,j-1}$ ($j \geq 1$), $\max_j \{r_{i,j} - r_{i,j-1}\}$, cannot exceed $\mathcal{V}_i - C_i$, which can be obtained when there are no higher priority preemptions in the execution of jobs $J_{i,j}$ s (for example, the highest priority transaction τ_1 always has separation $\mathcal{V}_1 - C_1$ for $J_{1,j}$ and $J_{1,j-1}$). Thus, the processor utilization for *DS-FP* should be greater than $\sum_{i=1}^m \frac{C_i}{\mathcal{V}_i - C_i}$, which is the CPU

workload resulting from the maximal separation $\mathcal{V}_i - C_i$ of each transaction.

If $f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}$, where $f_{i,0}^{ml}$ is the first job's response time (that is, the worst-case response time) of τ_i 's job in *ML*, *ML* can be regarded as a special case of *DS-FP* in which the sampling (or release) time $r_{i,j+1}^{ml}$ and deadline $d_{i,j+1}^{ml}$ ($j \geq 0$) can be specified as follows:

$$d_{i,j+1}^{ml} = r_{i,j}^{ml} + \mathcal{V}_i, \quad (10)$$

$$r_{i,j+1}^{ml} = d_{i,j+1}^{ml} - (\Theta_i(r_{i,0}^{ml}, f_{i,0}^{ml}) + C_i). \quad (11)$$

It is clear that $\Theta_i(r_{i,0}^{ml}, f_{i,0}^{ml}) + C_i = f_{i,0}^{ml}$ when $r_{i,0}^{ml} = 0$ ($1 \leq i \leq m$) in *ML*.

Theorem 3.1. *Given a synchronous update transaction set \mathcal{T} with known C_i and \mathcal{V}_i ($1 \leq i \leq m$), if $(\forall i) f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}$ in *ML*, then*

$$WCRT_i \leq f_{i,0}^{ml},$$

where $WCRT_i$ and $f_{i,0}^{ml}$ denote the worst-case response times of τ_i in *DS-FP* and *ML*, respectively.

Proof. This can be proven by contradiction. Suppose that τ_k is the highest priority transaction such that $WCRT_k > f_{k,0}^{ml}$ holds in *DS-FP*. Also, it is assumed that the response time of $J_{k,n}$ ($n \geq 0$), $R_{k,n}$, is the worst for τ_k in *DS-FP*. Note that schedules of τ_1 in *ML* and *DS-FP* are the same, as, in both cases, τ_1 jobs have the same relative deadline C_1 and separation/period $\mathcal{V}_1 - C_1$. Therefore, $1 < k \leq m$ holds.

As $WCRT_k > f_{k,0}^{ml}$, there must be a transaction τ_l such that 1) τ_l has a priority higher than τ_k ($1 \leq l < k$), 2) at least two consecutive jobs of τ_l , $J_{l,j-1}$ and $J_{l,j}$ overlap with $J_{k,n}$, and 3) the separation of $J_{l,j-1}$ and $J_{l,j}$ satisfies the following condition:

$$r_{l,j} - r_{l,j-1} < \mathcal{V}_l - f_{l,0}^{ml} \quad (j > 0), \quad (12)$$

where $\mathcal{V}_l - f_{l,0}^{ml}$ is the period (that is, separation) of the jobs of τ_l in *ML*.

Claim 1 is true because $k > 1$. It is straightforward that if each higher priority transaction of τ_k only has one job overlapping with $J_{k,n}$, then $R_{k,n} \leq f_{k,0}^{ml}$. This implies that Claim 2 is true. Finally, for $(\forall l < k)$ and $J_{l,j-1}$ and $J_{l,j}$ overlapping with $J_{k,n}$, if

$$r_{l,j} - r_{l,j-1} \geq \mathcal{V}_l - f_{l,0}^{ml} \quad (j > 0),$$

then $R_{k,n} > f_{k,0}^{ml}$ cannot be true because the amount of preemptions from higher priority transactions received by $J_{k,n}$ in *DS-FP* is not more than that received by $J_{k,0}$ in *ML*. Thus, Claim 3 is also true.

We know that the release time $r_{l,j}$ in *DS-FP* is derived as follows:

$$r_{l,j} = d_{l,j} - R_{l,j}, \quad (13)$$

where $R_{l,j}$ is the calculated response time of job $J_{l,j}$, that is, $\Theta_l(r_{l,j}, d_{l,j}) + C_l$. Following (12) and (13),

$$\begin{aligned} d_{l,j} - R_{l,j} &= r_{l,j} \text{ \{by (13)\}} \\ &< r_{l,j-1} + \mathcal{V}_l - f_{l,0}^{ml} \text{ \{by (12)\}} \\ &= d_{l,j} - f_{l,0}^{ml} \text{ \{by (3)\}}. \end{aligned}$$

Finally,

$$R_{l,j} > f_{l,0}^{ml}. \quad (14)$$

Equation (14) contradicts the assumption that τ_k is the highest priority transaction such that $WCRT_k > f_{k,0}^{ml}$ holds. Therefore, the theorem is proven. \square

The following theorem gives a sufficient condition for the schedulability of *DS-FP*.

Theorem 3.2. *Given a synchronous update transaction set \mathcal{T} with known C_i and \mathcal{V}_i ($1 \leq i \leq m$), if $(\forall i) f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}$ in *ML*, then \mathcal{T} is schedulable with *DS-FP*.*

Proof. If $f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}$, then the worst-case response times of τ_i ($1 \leq i \leq m$) in *DS-FP* $WCRT_i$ satisfy the following condition (by Theorem 3.1):

$$WCRT_i \leq f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}.$$

That is, $WCRT_i$ is not more than $\frac{\mathcal{V}_i}{2}$. Because the following equations hold in *DS-FP*, according to (5) and (6):

$$r_{i,j} = d_{i,j} - R_{i,j}, \quad (15)$$

$$d_{i,j+1} = r_{i,j} + \mathcal{V}_i, \quad (16)$$

$$d_{i,j+1} = r_{i,j+1} + R_{i,j+1}. \quad (17)$$

Replacing $r_{i,j}$ and $d_{i,j+1}$ in (16) with (15) and (17), respectively, it follows that

$$r_{i,j+1} + R_{i,j+1} = d_{i,j} - R_{i,j} + \mathcal{V}_i.$$

That is,

$$r_{i,j+1} - d_{i,j} + R_{i,j+1} + R_{i,j} = \mathcal{V}_i. \quad (18)$$

Because

$$R_{i,j+1} + R_{i,j} \leq 2 \cdot WCRT_i \leq \mathcal{V}_i,$$

it follows from (18) that $r_{i,j+1} - d_{i,j} \geq 0$ holds. This ensures that it is schedulable to schedule two jobs of τ_i in one validity interval \mathcal{V}_i under *DS-FP*. Thus, \mathcal{T} is schedulable with *DS-FP*. \square

The following corollary states the correctness of *DS-FP*.

Corollary 3.1. *Given a synchronous update transaction set \mathcal{T} with known C_i and \mathcal{V}_i ($1 \leq i \leq m$), if $(\forall i) f_{i,0}^{ml} \leq \frac{\mathcal{V}_i}{2}$ in *ML*, then *DS-FP* correctly guarantees the temporal validity of real-time data.*

Proof. As deadline assignment in *DS-FP* follows (3), the largest distance of two consecutive jobs $d_{i,j+1} - r_{i,j}$ ($j \geq 0$) does not exceed \mathcal{V}_i . The *validity constraint* can be satisfied if all jobs meet their deadlines, which is guaranteed by Theorem 3.2. \square

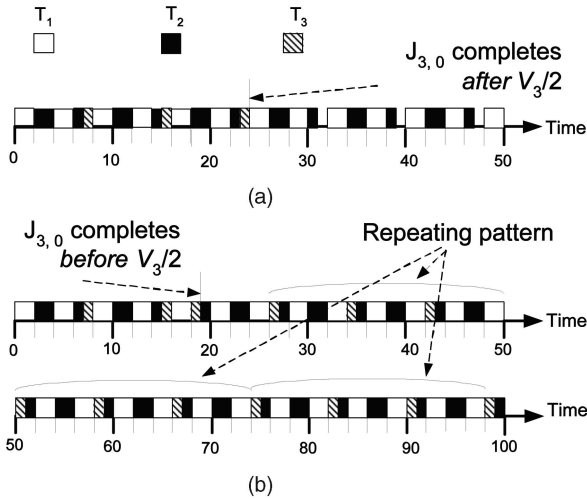


Fig. 5. A transaction set schedulable by *DS-FP* but not by *ML*. (a) *ML* is unschedulable. (b) *DS-FP* is schedulable.

If \mathcal{T} can be scheduled by *ML*, then by *ML* definition, $(\forall i) f_{i,0}^{ml} \leq \frac{V_i}{2}$. Thus, Corollary 3.2, which states a sufficient schedulability condition for *DS-FP*, directly follows from Theorem 3.2.

Corollary 3.2. *Given a synchronous update transaction set \mathcal{T} with known C_i and V_i ($1 \leq i \leq m$), if \mathcal{T} can be scheduled by *ML*, then it can also be scheduled by *DS-FP*.*

However, the converse statement of Corollary 3.2 is not true. That is, if \mathcal{T} can be scheduled by *DS-FP*, then it is not necessarily true that \mathcal{T} can also be scheduled by *ML*. This is demonstrated in the following example.

Example 3.3. Consider a set of three transactions $\{\tau_1, \tau_2, \tau_3\}$ with computation times 2, 3, and 3 and validity intervals 6, 15, and 47, respectively. Fig. 5a depicts a schedule of the transactions under *ML*. The first job of τ_3 , $J_{3,0}$, completes at time 24, which is greater than $\frac{V_3}{2}$ (that is, 23.5). Thus, the set of transactions is not schedulable by *ML*.

Fig. 5b depicts a schedule of the transactions under *DS-FP*. The same transaction set is schedulable by *DS-FP*, because the schedule pattern between times 26 and 50 repeats itself forever.

In summary, if a set of synchronous update transactions can be scheduled by *ML* to satisfy the *validity constraint*, then it can also be scheduled by *DS-FP*. However, the converse statement is not true, which implies that *DS-FP* outperforms *ML* in terms of schedulability. Thus, the following corollary directly follows from both Corollary 3.2 and Example 3.3.

Corollary 3.3. *DS-FP outperforms ML in terms of schedulability for satisfying the validity constraint.*

Discussion of jitters. Our results can be easily extended to the case where jitter between the sampling time and the release time of a job is nonzero if the maximum jitter of a transaction is known. In the presence of nonzero jitters, we can transform a transaction τ'_i (with validity length V'_i and maximum jitter δ_i) to a transaction τ_i (with validity length

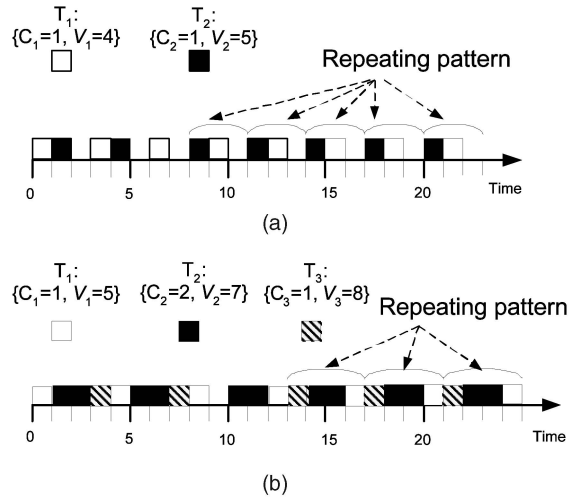


Fig. 6. *DS-FP* schedules with *fixed patterns*. (a) A *DS-FP* schedule for two transactions. (b) A *DS-FP* schedule for three transactions.

$V_i = V'_i - \delta_i$ and zero jitter). Such a transformation guarantees that if τ'_i can meet its validity constraint, then τ'_i can also meet its validity constraint.

3.4 Theoretical Estimation of Processor Utilization for *DS-FP*

This section presents means of estimating the average CPU utilization. Note that *DS-FP* does not usually schedule transactions periodically. Thus, it is hard to derive its exact CPU utilization unless there is a fixed pattern that repeats itself in a *DS-FP* schedule. In what follows, we shall investigate two cases in order: 1) a *DS-FP* schedule that has a detected pattern repeating itself from a certain point in time and 2) a *DS-FP* schedule that has no detected pattern.

3.4.1 *DS-FP* with a Detected Pattern

We introduce a *fixed pattern* in a *DS-FP* schedule with a simple example, which is shown in the following:

Example 3.4. Consider a *DS-FP* schedule for two transactions τ_1 and τ_2 in Fig. 6a. Note that transaction parameters (C_i s and V_i s) are given in the figure. We observe that there is a *fixed pattern* repeating itself in the schedule every three time units, starting from time 8. If time goes to infinity, we can estimate that the average CPU utilization of the *DS-FP* schedule is about 66.7 percent. Similarly, given the three transactions in Fig. 6b, we observe a *fixed pattern* repeating itself in the schedule every four time units, starting from time 13. Again, we can easily estimate that its CPU utilization is close to 100 percent.

Needless to say, the average CPU utilization for a *DS-FP* schedule can be approximated based on a *fixed pattern* if such a pattern exists in the schedule. However, it is not true that we can always easily detect a fixed pattern in every *DS-FP* schedule. It becomes harder to detect a *fixed pattern* in a *DS-FP* schedule if the size of the transaction set is larger. This is because the complexity of pattern detection grows exponentially with the size of the transaction set. Indeed, it remains open whether there is always a fixed pattern in

every *DS-FP* schedule or not. As many *DS-FP* schedules may not be detected to have such *fixed patterns*, it becomes more important to estimate the average CPU utilization for such *DS-FP* schedules.

3.4.2 *DS-FP without a Detected Pattern*

We now present an approximation of the average processor utilization of *DS-FP* from a statistical perspective *in the absence of detected patterns in DS-FP schedules*. Note that *our approximation only works provided that \mathcal{T} can be scheduled by ML*. This implies that the approximation is applicable to transaction sets where all deadlines are not greater than their corresponding periods in *ML*. Our approximation is quite close to the average CPU utilization obtained in our experiments. The CPU utilization approximation depends on the approximate values of the average deadline \bar{D} and period \bar{P} of transactions, which is described as follows.

Given a set of transactions $\mathcal{T} = \{\tau_i\}_{i=1}^m$, let \bar{U}_{DS} denote the average processor utilization in *DS-FP* and let \bar{P}_j be the average period for τ_j . The average relative deadline of τ_i , namely, \bar{D}_i , is approximated as follows:

$$\bar{D}_i = C_i + \sum_{j=1}^{i-1} \left[\left(\frac{\bar{D}_j}{\bar{P}_j} \right) \times C_j \right] \quad (1 \leq i \leq m). \quad (19)$$

Let $P_{i,j}$ and $D_{i,j+1}$ ($1 \leq i \leq m \wedge j \geq 0$) denote $r_{i,j+1} - r_{i,j}$ and $d_{i,j+1} - r_{i,j+1}$ in (4), respectively. It follows that

$$P_{i,j} + D_{i,j+1} = \mathcal{V}_i. \quad (20)$$

Thus, the following equation holds, given an arbitrarily large n ($n \rightarrow \infty$), where n is the number of jobs in averaging:

$$\bar{P}_i + \bar{D}_i = \mathcal{V}_i. \quad (21)$$

Following (19) and (21), \bar{D}_i and \bar{P}_i ($1 \leq i \leq m$) can be calculated (from the highest priority transaction τ_1 to the lowest priority transaction τ_m), respectively, as follows:

$$\bar{D}_i = \frac{C_i}{1 - \sum_{j=1}^{i-1} \frac{C_j}{\bar{P}_j}} \quad (1 \leq i \leq m), \quad (22)$$

$$\bar{P}_i = \mathcal{V}_i - \bar{D}_i \quad (1 \leq i \leq m). \quad (23)$$

Finally, \bar{U}_{DS} , which is the average utilization of the transaction set \mathcal{T} under *DS-FP*, can be approximated as

$$\bar{U}_{DS} = \sum_{i=1}^m \frac{C_i}{\bar{P}_i} = \sum_{i=1}^m \left(\frac{C_i}{\mathcal{V}_i - \frac{C_i}{1 - \sum_{j=1}^{i-1} \frac{C_j}{\bar{P}_j}}} \right). \quad (24)$$

The following example illustrates how the average utilization is estimated.

Example 3.5. Given the transaction set in Table 2, we calculate the average relative deadline and period of τ_i ($i = 1, 2, 3$) as follows:

$$\bar{D}_1 = C_1 = 1, \quad \bar{P}_1 = \mathcal{V}_1 - \bar{D}_1 = 4,$$

$$\bar{D}_2 = \frac{C_2}{1 - \frac{C_1}{\bar{P}_1}} = 2.7, \quad \bar{P}_2 = \mathcal{V}_2 - \bar{D}_2 = 7.3,$$

$$\bar{D}_3 = \frac{C_3}{1 - \left(\frac{C_1}{\bar{P}_1} + \frac{C_2}{\bar{P}_2} \right)} = 4.2, \quad \bar{P}_3 = \mathcal{V}_3 - \bar{D}_3 = 15.8.$$

The average processor utilization is $\bar{U}_{DS} = \sum_{i=1}^m \frac{C_i}{\bar{P}_i} = 0.65$. Given the transaction set in Table 2, it can be verified that the processor utilization for the first 200 time units is 63 percent, which is very close to our theoretical estimation and lower than the processor utilization from *ML* (68 percent).

Discussion of fixed patterns. A *fixed pattern* in a *DS-FP* schedule may be exponentially long (with respect to the number of transactions). Thus, it can be very expensive to detect. Assume that the minimal number of jobs per transaction in this pattern is n . If n is large, then (24) can be used to estimate the average CPU utilization of the fixed pattern, which, in turn, is the utilization estimation of the schedule.

In summary, the average CPU utilization of a *DS-FP* schedule can be approximated based on a fixed pattern if such a pattern exists in the schedule. Otherwise, the CPU utilization can be estimated by (24) if the transaction set is schedulable according to Corollary 3.2.

3.5 The Nonoptimality of *DS-FP*

We have proven in Section 3.4 that *DS-FP* is close to optimal in terms of minimizing the CPU workload from a statistical perspective. Intuitively, *DS-FP* should be very close to an optimal algorithm because it always defers the execution of a job as late as possible, hence reducing the workload as much as possible. We have also proven that *DS-FP* can schedule any transaction set that is schedulable by *ML* in Section 3.3. Now, it is interesting to know if *DS-FP* is an optimal algorithm in terms of schedulability. That is, *given any transaction set, if it is schedulable by a fixed-priority scheduler, can it be scheduled by DS-FP?* Unfortunately, the answer to the aforementioned question is negative, which can be demonstrated with the following example.

Example 3.6. Consider a set of three transactions $\{\tau_1, \tau_2, \tau_3\}$ with computation times 4, 4, and 3 and validity intervals 12, 22, and 36, respectively. This set is not schedulable by *DS-FP* as it fails at time 36, as shown in Fig. 7a. In this case, the second job of τ_3 cannot be completed by the end of its first validity interval. However, if $J_{1,2}$ is scheduled two time units earlier, this transaction set can be successfully scheduled because there is a *fixed pattern* repeating itself every 32 time units starting from time 27, as depicted in Fig. 7b. Note that such a schedule is also a fixed-priority schedule because no lower priority jobs may interrupt a higher priority job once the higher one is released. By doing so, the release time of $J_{2,1}$ is postponed to time 18, as shown in Fig. 7b (from time 14 in Fig. 7a). Hence, the deadline of $J_{2,2}$ is also postponed.

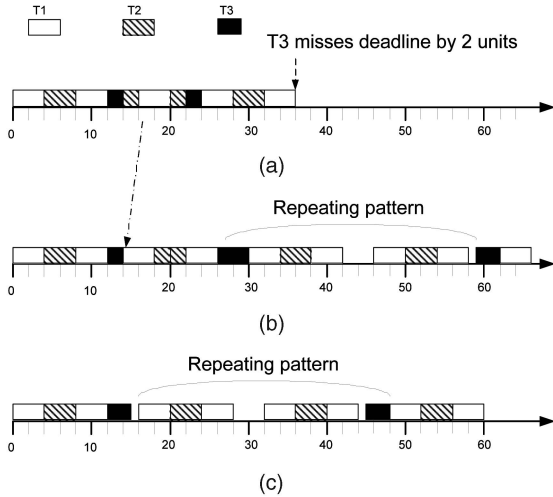


Fig. 7. *DS-FP* is not optimal. (a) An unsuccessful *DS-FP* schedule. (b) A successful schedule. (c) An asynchronous *DS-FP* schedule.

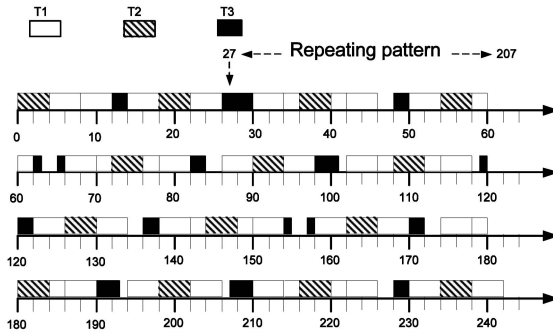


Fig. 8. *SVF* is not optimal for *DS-FP*.

DS-FP requires that every transaction τ_i should finish its first two jobs in $[0, \mathcal{V}_i)$. If the requirement is relaxed so that the first two jobs are allowed to finish in $[r_{i,0}, r_{i,0} + \mathcal{V}_i)$, where $r_{i,0}$ denotes the time at which $J_{i,0}$ actually starts, then *DS-FP* can schedule the set in Example 3.6. In this case, the first jobs of transactions start *asynchronously*. An asynchronous schedule for the same transaction set in Example 3.6 is depicted in Fig. 7c, in which there is a *fixed pattern* between times 16 and 48 repeating itself forever. In general, whether the asynchronous *DS-FP* algorithm is optimal in terms of schedulability remains an open question.

Another interesting observation is that the transaction set in Example 3.6 is schedulable by *DS-FP* if a priority order different from *SVF* is used. For example, if we swap the priorities of τ_1 and τ_2 , *DS-FP* can schedule the set, as depicted in Fig. 8. In this case, there is a *fixed pattern* between times 27 and 207 repeating itself forever.

In summary, *DS-FP* is not optimal for a set of synchronous update transactions in terms of schedulability, but it remains open if it is optimal for asynchronous transactions or transaction priorities assigned differently from *SVF*.

4 PERFORMANCE EVALUATION

This section presents the important results from our simulation studies. Section 4.1 describes our simulation

TABLE 4
Experimental Parameters and Settings

Parameters	Meaning	Value
N_{CPU}	No. of CPU	1
N_T	No. of real-time data objects	[10, 300]
\mathcal{V}_i (ms)	Validity interval of data X_i	[4000, 8000]
C_i (ms)	CPU time for updating X_i	[5, 15]
Length	No. of data to update	1

model and parameters. Section 4.2 compares *DS-FP* with the *ML* algorithm. *ML* is known to outperform Half-Half [25], which is not compared here. The experiments demonstrate that our proposed approaches reduce CPU utilization while guaranteeing data validity constraints.

4.1 Simulation Model and Parameters

Our experiments compare the update transaction workloads produced by *DS-FP* and *ML*. It is demonstrated that *DS-FP* produces a lower CPU workload than *ML*. Also, the experiments demonstrate that the increase in the average sampling period from *DS-FP* is the main reason for its lower workload. The primary performance metric used in the experiments is the CPU workload.

A summary of the parameters and default settings used in the experiments is presented in Table 4. The baseline values for the parameters follow those used in [25], which are originally from air traffic control applications. We consider a single CPU main memory-based RTDBS. The number of real-time data objects varies from 10 to 300 and the validity interval of each real-time data object is uniformly distributed between 4,000 and 8,000 ms. Each transaction updates one real-time data object and the CPU time for each transaction is uniformly distributed between 5 and 15 ms. In the experiments, 95 percent confidence intervals have been obtained, whose widths are less than ± 5 percent of the point estimate for the performance metrics.

4.2 Experimental Results

In our experiments, the CPU workloads of update transactions produced by *ML* and *DS-FP* are quantitatively compared. Update transactions are generated randomly according to the parameter settings in Table 4.

The resulting CPU workloads generated from *ML* and *DS-FP* are depicted in Fig. 9. From the results, we observe that the CPU workload produced by *DS-FP* is consistently lower than that of *ML*. In fact, the difference widens as the number of update transactions increases. The difference reaches 18 percent when the number of transactions is 300. It is also observed that the CPU utilization of *DS-FP* measured in our experiments (*DS-FP*(Expt.)) nearly matches the CPU workload estimation \bar{U}_{DS} (24), shown as *DS-FP*(Est.) in the figure. Moreover, the *DS-FP* CPU workload is only slightly higher than $\sum_{i=1}^m \frac{C_i}{\mathcal{V}_i - C_i}$, which is the CPU workload resulting from the maximal separation $\mathcal{V}_i - C_i$ ($1 \leq i \leq m$) of each transaction (see Section 3.3). In fact, the difference is insignificant in Fig. 9. The improvement in the CPU workload of *DS-FP* is due to the fact that *DS-FP* adaptively samples real-time data objects at a lower

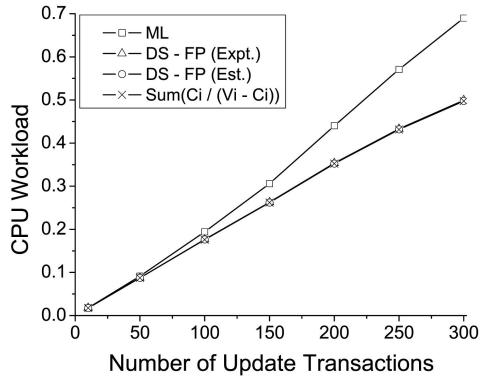


Fig. 9. CPU workload comparison.

rate. This is verified by the average sampling periods of update transactions obtained from experiments.

Fig. 10 shows the average sampling period for each transaction in *DS-FP* when the number of update transactions is 300. Given a set of update transactions, the period of transaction τ_i in *ML* P_i^{ml} is a constant and it can be calculated offline [25], while the separation of the sampling times of two consecutive jobs from the same transaction in *DS-FP* is dynamic and it is obtained online in the experiments. The mean value of the separations, that is, the average sampling period, \bar{P}_i^{ds} for transaction τ_i is calculated as follows, where n is the number of jobs generated by τ_i in the experiments:

$$\bar{P}_i^{ds} = \frac{1}{n-1} \sum_{j=1}^{n-1} (r_{i,j} - r_{i,j-1}). \quad (25)$$

In Fig. 10, it is observed that \bar{P}_i^{ds} is consistently larger than P_i^{ml} , while the difference $(\bar{P}_i^{ds} - P_i^{ml})$ increases with the decrease in the transaction's priority. *DS-FP* reduces the average sampling rate more for lower priority transactions, thus reducing the workload of CPU. Fig. 10 also shows that the trend of $(\frac{\bar{P}_i^{ds}}{P_i^{ml}})$ increases similarly to that of $(\bar{P}_i^{ds} - P_i^{ml})$, although it fluctuates.

Fig. 11 depicts how much the CPU workload estimation *DS-FP*(Est.) differs from the actual CPU utilization obtained from the experiments *DS-FP*(Expt.) in finer granularity. The x-axis depicts the size of update transactions and the y-axis

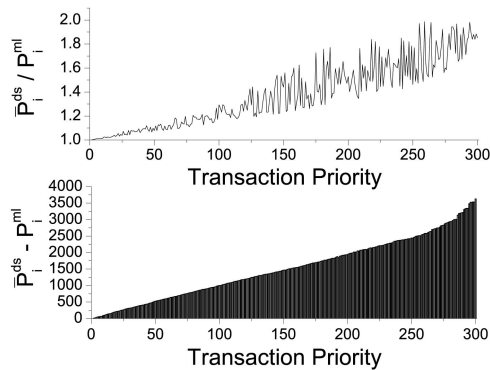


Fig. 10. Average sampling period comparison.

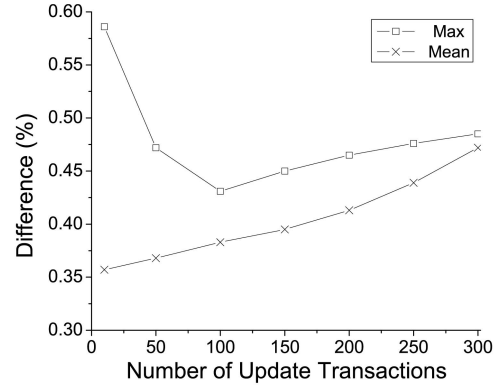


Fig. 11. CPU workload estimation error.

depicts the relative difference between *DS-FP*(Est.) and *DS-FP*(Expt.), which is defined as

$$\frac{|DS-FP(Expt.) - DS-FP(Est.)|}{DS-FP(Expt.)} \times 100\%.$$

Both the maximum and mean relative differences are depicted in the figure. In our experiments, it is observed that *DS-FP*(Expt.) is consistently higher than *DS-FP*(Est.). As observed in the figure, our CPU workload estimation nearly matches the measured CPU utilization in our experiments as the maximum relative difference never exceeds 0.6 percent.

We also conducted a set of experiments by varying $\frac{V_i}{C_i}$ and fixing $\sum_{i=1}^m \frac{C_i}{V_i}$ of the update transaction set at 45 percent. The results are depicted in Fig. 12, which compares *ML*, *DS-FP*, $\sum_{i=1}^m \frac{C_i}{V_i}$, and $\sum_{i=1}^m \frac{C_i}{V_i - C_i}$. Similarly to Fig. 9, the actual utilization for *DS-FP* is very close to the utilization estimation \bar{U}_{DS} (shown as *DS-FP*(Est.)). Note that $\sum_{i=1}^m \frac{C_i}{V_i - C_i}$ is the CPU workload resulting from the possible maximum separation $V_i - C_i$ satisfying the *validity constraint* for each transaction τ_i . It is a CPU lower bound ignoring transaction interference. It is observed in Fig. 12 that a CPU workload of *DS-FP* is very close to that of $\sum_{i=1}^m \frac{C_i}{V_i - C_i}$. The larger $\frac{V_i}{C_i}$ is, the closer *DS-FP* and $\sum_{i=1}^m \frac{C_i}{V_i - C_i}$ are. This is because the probability of transaction interference decreases for *DS-FP* when $\frac{V_i}{C_i}$ increases.

We have conducted more experiments to study the performance of *DS-FP* with different experimental settings.

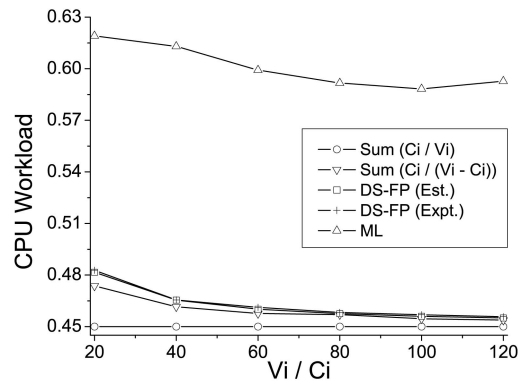


Fig. 12. CPU workloads with fixed $\sum_{i=1}^m \frac{C_i}{V_i}$.

The results are reported in [29] and they are omitted here due to space limitations.

In summary, when a set of update transactions is scheduled by *DS-FP* to maintain the temporal validity of real-time data objects, it produces a schedule with a much lower CPU workload than *ML* does. Thus, more CPU capacity is available for improving the performance of other workloads (for example, triggered transactions [27]) in the system.

5 RELATED WORK

There has been a lot of work on RTDBSs in which validity intervals are associated with real-time data [21], [22], [1], [12], [13], [14], [5], [26], [11], [30], [8], [7], [25], [10]. In [6], a safety-critical automotive application, that is, adaptive cruise control, is studied. It deals with critical data and involves deadline-bound computations on data gathered from the automobiles' environment. These applications have stringent requirements on the freshness of data objects and completion times of the tasks. In [8], a vehicular application with embedded engine control systems is presented and an on-demand scheduling algorithm is proposed for enforcing base and derived data freshness. In [7], an algorithm (ODTB) is proposed for updating data items that can skip unnecessary updates, allowing for better utilization of the CPU in the vehicular application. Such systems introduce the need to maintain data temporal consistency in addition to logical consistency.

In the model introduced in [22], a real-time system consists of periodic tasks which are either read-only, write-only, or update (read/write) transactions. Data objects are temporally inconsistent when their ages or dispersions are greater than the absolute or relative thresholds allowed by the application. Two-phase locking and optimistic concurrency control algorithms, as well as rate-monotonic and earliest deadline first scheduling algorithms are studied in [22]. In [12], [13], real-time data semantics are investigated, and a class of real-time data access protocols, called Similarity Stack Protocols (SSP), is proposed. The correctness of SSP is based on the concept of similarity, which allows different but sufficiently timely data to be used in a computation without adversely affecting the outcome.

In [14], similarity-based principles are coupled with the *Half-Half* approach to adjust the real-time transaction load by skipping the execution of task instances. The concept of *data deadline* is proposed in [26]. It also proposes data-deadline-based scheduling, forced-wait, and similarity-based scheduling techniques to maintain the temporal validity of real-time data and to meet transaction deadlines in RTDBSs. In [10], Jha et al. study whether, given an update transaction schedule, a periodic query would read mutually consistent data. They propose design approaches to decide the period and relative deadline of a query so that it satisfies mutual consistency. They then suggest ways of reducing the complexity of the solution approach by using harmonic periods in general.

Our work is related to the *ML* scheme in [2], [25], [30]. *ML* guarantees a bound on the sampling time of a periodic transaction job and the finishing time of its next job, but, as we showed, the deadline and the period of a periodic

transaction are derived from the worst-case response time of the transaction. This is different from the aperiodic task model-based *DS-FP* algorithm in which the deadline of a transaction job is derived adaptively and the separation of two consecutive jobs is not a constant. *DS-FP* further reduces the CPU workload resulting from update transactions by adaptively adjusting the separation of two consecutive jobs while satisfying the *validity constraint*. *DS-FP* is also different from the *distance constrained scheduling*, which guarantees a bound of the finishing times of two consecutive instances of a task [9]. The *EDL* algorithm proposed in [4] processes tasks as late as possible based on the Earliest Deadline scheduling algorithm [17]. However, *EDL* assumes that all deadlines of tasks are given, whereas *DS-FP* derives deadlines dynamically. Finally, our *DS-FP* algorithm is applicable to the scheduling of *age constraint* tasks in real-time systems [24].

6 CONCLUSIONS AND FUTURE WORK

This paper proposes a novel algorithm, namely, deferrable scheduling for fixed priority transactions (*DS-FP*). Distinct from past studies of maintaining the freshness (or temporal validity) of real-time data in which the periodic task model is adopted, *DS-FP* adopts the *aperiodic* task model. The deadlines of jobs and the separation of two consecutive jobs of an update transaction are adjusted judiciously so that the farthest distance of the sampling time of a job is achieved and the completion time of its next job is bounded by the validity length of the updated real-time data. This paper presents a sufficient condition for its schedulability. It also proposes a theoretical estimation of the processor utilization of *DS-FP*, which is verified in our experimental studies. It is also demonstrated in our experiments that *DS-FP* greatly reduces update workload compared to *ML*, while guaranteeing the validity constraint.

However, there are still many open questions to be answered for *DS-FP*. For example, it is not clear what a sufficient and necessary condition for the schedulability of *DS-FP* is, if time 0 is a critical instant for a synchronous transaction set scheduled by *DS-FP*, and if there is a least upper bound of CPU utilization for *DS-FP*. Moreover, the concept of *deferrable scheduling* is only used to schedule update transactions with fixed priority in this paper. It is possible for the same concept to be used in the scheduling of update transactions with dynamic priority, for example, in the Earliest Deadline scheduling [17], [4] of update transactions.

ACKNOWLEDGMENTS

Preliminary versions of this paper appeared in [27], [28]. The authors would like to thank Professor Krithi Ramamritham and Professor Aloysius K. Mok for fruitful discussions on this work. This work was partially done while Song Han was at City University of Hong Kong.

REFERENCES

- [1] N.C. Audsley, A. Burns, M.F. Richardson, and A.J. Wellings, "Data Consistency in Hard Real-Time Systems," *Informatica*, vol. 19, no. 2, 1995.

- [2] A. Burns and R. Davis, "Choosing Task Periods to Minimize System Utilization in Time-Triggered Systems," *Information Processing Letters*, vol. 58, pp. 223-229, 1996.
- [3] D. Chen and A.K. Mok, "Scheduling Similarity-Constrained Real-Time Tasks," *Proc. Int'l Conf. Embedded Systems and Applications and Int'l Conf. VLSI*, pp. 215-221, 2004.
- [4] H. Chetto and M. Chetto, "Some Results of the Earliest Deadline Scheduling Algorithm," *IEEE Trans. Software Eng.*, vol. 15, no. 10, pp. 1261-1269, Oct. 1989.
- [5] R. Gerber, S. Hong, and M. Saksena, "Guaranteeing End-to-End Timing Constraints by Calibrating Intermediate Processes," *Proc. 15th IEEE Real-Time Systems Symp.*, 1994.
- [6] G. Goud, N. Sharma, K. Ramamritham, and S. Malewar, "Efficient Real-Time Support for Automotive Applications: A Case Study," *Proc. 12th IEEE Int'l Conf. Embedded and Real-Time Computing Systems and Applications*, 2006.
- [7] T. Gustafsson and J. Hansson, "Data Management in Real-Time Systems: A Case of On-Demand Updates in Vehicle Control Systems," *Proc. 10th IEEE Real-Time and Embedded Technology and Applications Symp.*, pp. 182-191, 2004.
- [8] T. Gustafsson and J. Hansson, "Dynamic On-Demand Updating of Data in Real-Time Database Systems," *Proc. 19th ACM Symp. Applied Computing*, 2004.
- [9] C.C. Han, K.J. Lin, and J.W.-S. Liu, "Scheduling Jobs with Temporal Distance Constraints," *SIAM J. Computing*, vol. 24, no. 5, pp. 1104-1121, Oct. 1995.
- [10] A.K. Jha, M. Xiong, and K. Ramamritham, "Mutual Consistency in Real-Time Databases," *Proc. 28th IEEE Real-Time Systems Symp.*, 2006.
- [11] K.D. Kang, S. Son, J.A. Stankovic, and T. Abdelzaker, "A QoS-Sensitive Approach for Timeliness and Freshness Guarantees in Real-Time Databases," *Proc. 14th Euromicro Conf. Real-Time Systems*, 2002.
- [12] T. Kuo and A.K. Mok, "Real-Time Data Semantics and Similarity-Based Concurrency Control," *Proc. 13th IEEE Real-Time Systems Symp.*, 1992.
- [13] T. Kuo and A.K. Mok, "SSP: A Semantics-Based Protocol for Real-Time Data Access," *Proc. 14th IEEE Real-Time Systems Symp.*, 1993.
- [14] S. Ho, T. Kuo, and A.K. Mok, "Similarity-Based Load Adjustment for Real-Time Data-Intensive Applications," *Proc. 18th IEEE Real-Time Systems Symp.*, 1997.
- [15] J.P. Lehoczky, "Fixed Priority Scheduling of Periodic Task Sets with Arbitrary Deadlines," *Proc. 11th IEEE Real-Time Systems Symp.*, 1990.
- [16] J. Leung and J. Whitehead, "On the Complexity of Fixed-Priority Scheduling of Periodic Real-Time Tasks," *Performance Evaluation*, vol. 2, pp. 237-250, 1982.
- [17] C.L. Liu and J. Layland, "Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment," *J. ACM*, vol. 20, no. 1, 1973.
- [18] D. Locke, "Real-Time Databases: Real-World Requirements," *Real-Time Database Systems: Issues and Applications*, A. Bestavros, K.-J. Lin, and S.H. Son, eds., pp. 83-91, Kluwer Academic, 1997.
- [19] K. Ramamritham, "Real-Time Databases," *Distributed and Parallel Databases*, vol. 1, pp. 199-226, 1993.
- [20] K. Ramamritham, "Where Do Time Constraints Come From and Where Do They Go?" *Int'l J. Database Management*, vol. 7, no. 2, pp. 4-10, Spring 1996.
- [21] X. Song and J.W.S. Liu, "How Well Can Data Temporal Consistency Be Maintained?" *Proc. IEEE Symp. Computer-Aided Control Systems Design*, 1992.
- [22] X. Song and J.W.S. Liu, "Maintaining Temporal Consistency: Pessimistic versus Optimistic Concurrency Control," *IEEE Trans. Knowledge and Data Eng.*, vol. 7, no. 5, pp. 786-796, Oct. 1995.
- [23] J.A. Stankovic, S. Son, and J. Hansson, "Misconceptions about Real-Time Databases," *Computer*, vol. 32, no. 6, pp. 29-36, June 1999.
- [24] S. Vestal, "Real-Time Sampled Signal Flows through Asynchronous Distributed Systems," *Proc. 11th IEEE Real-Time and Embedded Technology and Applications Symp.*, pp. 170-179, 2005.
- [25] M. Xiong and K. Ramamritham, "Deriving Deadlines and Periods for Real-Time Update Transactions," *Proc. 20th IEEE Real-Time Systems Symp.*, 1999.
- [26] M. Xiong, K. Ramamritham, J.A. Stankovic, D. Towsley, and R.M. Sivasankaran, "Scheduling Transactions with Temporal Constraints: Exploiting Data Semantics," *IEEE Trans. Knowledge and Data Eng.*, vol. 14, no. 5, pp. 1155-1166, Sept./Oct. 2002.

- [27] M. Xiong, S. Han, and K.Y. Lam, "A Deferrable Scheduling Algorithm for Real-Time Transactions Maintaining Data Freshness," *Proc. 26th IEEE Real-Time Systems Symp.*, 2005.
- [28] M. Xiong, S. Han, and D. Chen, "Deferrable Scheduling for Temporal Consistency: Schedulability Analysis and Overhead Reduction," *Proc. 12th IEEE Int'l Conf. Embedded and Real-Time Computing Systems and Applications*, 2006.
- [29] M. Xiong, S. Han, D. Chen, and K.Y. Lam, "Deferrable Scheduling for Maintaining Real-Time Data Freshness: Algorithms, Analysis, and Results," Technical Report TR-07-44, Dept. of Computer Sciences, Univ. of Texas at Austin, Sept. 2007.
- [30] M. Xiong, B. Liang, K.Y. Lam, and Y. Guo, "Quality of Service Guarantee for Temporal Consistency of Real-Time Transactions," *IEEE Trans. Knowledge and Data Eng.*, vol. 18, no. 8, pp. 1097-1110, Aug. 2006.



Ming Xiong received the PhD degree in computer science from the University of Massachusetts, Amherst, in 2000. He is currently a member of technical staff at Bell Laboratories. His research interests include database systems, real-time systems, and mobile computing. He is a member of the IEEE.



Song Han received the BS degree in computer science from Nanjing University, Nanjing, People's Republic of China, in 2003 and the MPhil degree in computer science from City University of Hong Kong in 2006. He is currently working toward the PhD degree in the Department of Computer Sciences at the University of Texas at Austin. His research interests include real-time systems, database systems, wireless networks, and data mining. He is a student member of the IEEE.



Kam-Yiu Lam received the BSc (Hons; with distinction) degree in computer studies and the PhD degree from City University of Hong Kong in 1990 and 1994, respectively. He is currently an associate professor in the Department of Computer Science at City University of Hong Kong. His research interests include real-time database systems, real-time active database systems, mobile computing, and distributed multimedia systems.



Deji Chen received the PhD degree in computer science from the University of Texas at Austin in 1999. He is currently a senior principal software engineer at Emerson Process Management. His research interests include real-time systems and wireless process control. He is a member of the IEEE and the IEEE Computer Society.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.