

# Fast, Best-Effort Real-Time Scheduling Algorithms

Peng Li, *Student Member, IEEE*, and Binoy Ravindran, *Senior Member, IEEE*

**Abstract**—This paper presents two fast, best-effort real-time scheduling algorithms called MDASA and MLBESA. MDASA and MLBESA are novel in the way that they heuristically, yet accurately, mimic the behavior of the DASA and LBESA scheduling algorithms, but are faster with  $O(n)$  and  $O(n \lg(n))$  worst-case complexities, respectively. Experimental results show that the performance of MDASA and MLBESA, in general, is close to that of DASA and LBESA, respectively, for a broad range of realistic workloads. However, for a highly bursty workload, MLBESA is found to perform worse than LBESA. Furthermore, the task response times under MDASA and MLBESA are very close to the values under their counterpart scheduling algorithms. Thus, MDASA and MLBESA can substitute for DASA and LBESA algorithms, respectively, in adaptive resource allocation techniques for asynchronous real-time distributed systems where DASA and LBESA have previously been serious bottlenecks on computational costs.

**Index Terms**—Best-effort real-time scheduling, overload scheduling, response time analysis, asynchronous real-time systems, distributed real-time systems.

## 1 INTRODUCTION

ASYNCHRONOUS real-time distributed systems are emerging in many domains, including defense, telecommunication, and industrial automation, for the purpose of strategic mission management [22]. Such systems are fundamentally distinguished by the significant runtime uncertainties that are inherent in their application environment and system resource states [22], [24]. Consequently, it is difficult to postulate upper bounds on application workloads for such systems that will always be respected at runtime.

An example of such asynchronous real-time distributed systems is the US Navy's Anti-Air Warfare (AAW) system. AAW is an example air-defense system that is envisioned to be onboard the US Navy's future surface combatants. To better understand AAW, a software prototype that approximated AAW, called *DynBench*, was developed as part of DARPA's Quorum program [24]. Details of *DynBench* were presented in [38]. Such a system often operates in environments that are subject to significant uncertainties—parameters such as event arrival rates and execution time bounds cannot be accurately characterized at design time [15], [22].

Given the aforementioned uncertainties, the classical hard real-time satisfiability objective of “always meet all timing constraints” and hard real-time scheduling algorithms that achieve the satisfiability objective are not cost-effective or even impossible for asynchronous real-time distributed systems. Thus, in this paper, we consider timeliness requirements that are described by Jensen et al.'s benefit accrual model [23]. This model generalizes

the hard real-time “deadline” timing constraint to encompass nonbinary timing constraints. That is, completion of a task at any time yields a benefit that is specified by the task's *benefit function*. For example, a benefit function may specify a constant benefit before a task's deadline time and zero benefit after that, which is also known as a “firm deadline” in the real-time literature and is our model used in this paper.

To deal with the inherent nondeterminisms in asynchronous real-time distributed systems, we have developed adaptive resource allocation algorithms such as the RBA\* (Response-Time-Based Best-Effort Resource Allocation) algorithm [19]. The RBA\* algorithm seeks to maximize the aggregate benefit of trans-node real-time computations that have end-to-end timeliness requirements, specified using benefit functions. Furthermore, it considers “best-effort” process scheduling algorithms that have the scheduling objective of maximizing accrued benefit on each system node.

These scheduling algorithms are called “best-effort” in the sense that they seek to provide the best benefit to the application tasks, where the best benefit that can be accrued by an application task is application-specified using benefit functions. More precisely, the algorithms' scheduling objective is to maximize Accrued Benefit Ratio (ABR)—the ratio of total accrued benefit to the sum of all task benefits.

Two of the most prominent best-effort real-time scheduling algorithms include the Dependent Activity Scheduling Algorithm (DASA) [16] and Locke's Best Effort Scheduling Algorithm (LBESA) [29]. DASA and LBESA are equivalent to the Earliest Deadline First (EDF) algorithm during underloaded conditions [16], where EDF is optimal and guarantees that all deadlines are always satisfied. In the event of an overload situation, DASA and LBESA seek to maximize the aggregate task benefit.

DASA and LBESA are known to outperform other single processor real-time scheduling algorithms (during overload situations) in general. For example, in [16], Clark shows that

• The authors are with the Real-Time Systems Laboratory, The Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061.  
E-mail: {peli2, binoy}@ot.edu.

Manuscript received 28 Mar. 2002; revised 23 June 2003; accepted 10 Feb. 2004.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number 116177.

DASA generally outperforms LBESA during overload situations. In addition, in [19], Hegazy and Ravindran show that DASA outperforms the Robust Earliest Deadline First (RED) [13] and the Robust High Density (RHD) [12] scheduling algorithms. Furthermore, variants of DASA and LBESA have been implemented in the Alpha real-time operating system [21] and the MK7.3 kernel [41]. There are other implementations of overload real-time scheduling algorithms, such as the Maximum Urgency First (MUF) scheduler [40] implemented inside the CHIMERA II real-time operating system. However, to the authors' best knowledge, DASA and LBESA are the only two benefit accrual scheduling algorithms that have been implemented and deployed in mission-critical systems [15].

In spite of these remarkable results of DASA and LBESA, they have been serious performance bottlenecks for resource allocation algorithms such as RBA\*. The RBA\* algorithm is a proactive algorithm in the sense that it allocates resources to tasks for a future time window. To do that, RBA\* needs to determine task response times for the future time window, e.g., the RBA\* algorithm is invoked every several seconds. However, due to the complexity of analyzing process response times under DASA, RBA\* incurs a worst-case computational complexity that is a polynomial of the fourth degree.<sup>1</sup> In fact, in a middleware implementation of the RBA\* algorithm [27], we have observed that the overhead of invoking RBA\* can be up to 1 sec, which is quite high for deadlines in the magnitude of a few seconds.

Analyzing process response times under best-effort scheduling algorithms is difficult because best-effort schedulers such as DASA and LBESA make decisions which are functions of the remaining process execution times at an event. Thus, RBA\* estimates response times by determining the scheduling events that occur during a time interval (under which the response times are to be determined) and by applying the scheduling algorithm at each scheduling event to determine the scheduling decision. This is computationally expensive, as we incur the cost of the scheduler— $O(n^2)$  for DASA and LBESA given  $n$  tasks in the ready queue [16], [29]—whenever the scheduler is invoked at a scheduling event.

Thus, we are motivated to develop fast best-effort scheduling algorithms that can yield the same performance as that of DASA and LBESA, if not better. Furthermore, if such new scheduling algorithms behave similarly to DASA and LBESA, they can then be used by resource allocation algorithms such as RBA\* for determining response times under DASA and LBESA, respectively, in a faster way. This will reduce the computational complexity of RBA\*, resulting in good performance at low computational cost.

In this paper, we present two best-effort scheduling algorithms called **Modified DASA** (or MDASA) and **Modified LBESA** (or MLBESA) that exactly address the aforementioned problem. MDASA and MLBESA are novel in the way that they heuristically, yet accurately, mimic the

behavior of the DASA and LBESA algorithms, but are faster with  $O(n)$  and  $O(n \lg(n))$  worst-case complexities, respectively. MDASA and MLBESA reason about the behavior of DASA and LBESA by heuristically determining a feasible schedule of the process ready queue.

Our experimental results show that, in general, MDASA and MLBESA perform almost as well as DASA and LBESA, respectively. However, under highly bursty and heavily overloaded situations, DASA and LBESA may outperform MDASA and MLBESA. Furthermore, the process response times under MDASA and MLBESA are also found to be very close to the values under their counterpart scheduling algorithms. While MDASA has better performance than MLBESA and has better worst-case complexity, MLBESA guarantees the optimal schedule during underload situations.

The rest of the paper is organized as follows: We discuss the task model considered by MDASA and MLBESA in Section 2. Section 3 and Section 4 describe MDASA and MLBESA, respectively. We present the experimental evaluation of the algorithms in Section 5. In Section 6, we discuss the past and related works and contrast them with MDASA and MLBESA. Finally, the paper concludes with a summary of the work, its contributions, and identify future work in Section 7.

## 2 THE TASK MODEL

We consider a set of real-time tasks (or processes)<sup>2</sup> that are executed on a single processor. The set of tasks is denoted by the set  $T = \{T_1, T_2, \dots, T_n\}$ . Further, all tasks are assumed to be preemptable and, therefore, can be preempted at any time during their execution. Furthermore, tasks do not share any non-CPU resources or have precedence relations with other tasks and, therefore, are independent of each other.

Each task  $T_i$  is characterized by an arrival time  $A_i$  and an execution time  $C_i$ . For simplicity, we assume that the task execution time  $C_i$  is known by the scheduler at the time of arrival. This can be achieved by application profiling techniques, i.e., deriving the relationship between task workload and execution time, as done in [1], [36]. Similarly to [16], we assume downward “step” benefit functions for all tasks. Thus, completing a task any time before its deadline will result in uniform benefit; completing it after the deadline will result in zero benefit.

We denote the benefit of a task  $T_i$ —the height of  $T_i$ 's benefit function—as  $B_i$  and the absolute deadline of a task  $T_i$  as  $d_i$ , respectively.<sup>3</sup>

We use  $R_i(t)$  to denote the remaining execution time of a task  $T_i$  at time  $t$ . It is easy to see that  $R_i(t) \leq C_i, \forall t \geq 0$ . The slack of a task  $T_i$  at time  $t$  is therefore defined as  $S_i(t) = d_i - t - R_i(t), \forall t \geq 0$ .

For completeness, we also include definitions of several well-known quantities and terms here:

1. Given  $n$  trans-node tasks, a maximum of  $m$  processes per task,  $p$  processors, a minimum task period of  $k$ , and an adaptation window of length  $W$ , RBA\* incurs a worst-case computational complexity of  $O(p^2 m^4 n^4 \lceil W/k \rceil^4)$ .

2. We will use the terms *task* and *process* interchangeably in this paper unless otherwise stated.

3. The “benefit” of a task is also called “value” or “utility” in the literature [10].

1. For a generic real-time system with  $n$  tasks, where tasks can be activated dynamically, i.e., all tasks are aperiodic tasks, the processor load during the time interval of  $[t, d_i]$  is given by

$$\rho_i(t) = \frac{\sum_{d_k \leq d_i} c_k(t)}{(d_i - t)}$$

( $i \in [1, \dots, n]$ ), where  $c_k(t)$  refers to the remaining execution time of task  $T_k$  with deadline earlier than or equal to  $d_i$  [13]. Further, the total load in the interval of  $[t, d_n]$  is defined as  $\rho = \max_i(\rho_i)$  [13]. Furthermore, if  $\rho \leq 1$ , we say that the system is underloaded. Otherwise, we say the system is overloaded [13].

2. A schedule is an order of assigning tasks to the processor. That is, a schedule contains ordered executions of tasks.
3. We say a schedule is *feasible* if it can satisfy all timing constraints, such as deadlines of tasks within the schedule. Otherwise, the schedule is *infeasible*. A set of tasks is feasible if and only if there exists a feasible schedule on the set of tasks. If the task set is underloaded, EDF can produce a feasible schedule on the task set due to EDF's optimality [17].
4. The competitive factor measures the worst-case performance of an algorithm. This concept was introduced in [5] and is defined as follows: A scheduling algorithm  $\mathcal{A}$  has a competitive factor  $\varphi_{\mathcal{A}}$  if and only if it can guarantee a cumulative value (or benefit)  $\Gamma_{\mathcal{A}} \geq \varphi_{\mathcal{A}} \Gamma^*$ , where  $\Gamma^*$  is the cumulative value (or benefit) achieved by the optimal clairvoyant scheduler.

### 3 THE MDASA SCHEDULING ALGORITHM

The DASA algorithm makes scheduling decisions using the concept of *benefit densities*. The benefit density of a task is the benefit accrued per unit time by the execution of the task. Thus, the benefit density defines a measure of the "return of investment" for the task. We denote the benefit density of a task  $T_i$  at time  $t$  as  $BD_i(t)$ , which is given by  $BD_i(t) = B_i/R_i(t)$ .

The objective of DASA is to compute a schedule  $\Gamma$  that will maximize the aggregate task benefit. The aggregate task benefit is the cumulative sum of the benefit accrued by the execution of the tasks. Thus, a schedule that satisfies all deadlines of all tasks will clearly yield the maximum aggregate benefit (since task benefit functions are step-benefit functions).

During underload conditions, we know that EDF produces such a feasible schedule for all tasks [17]. Such an all-task-feasible schedule  $\Gamma$  will contain all tasks in  $T$ , if all tasks are present in the task ready queue. However, if not all task deadlines can be satisfied, then the schedule that will yield the maximum benefit will be an exact subset of the task set, i.e.,  $\Gamma \subset T$ , since some tasks will have to be excluded from the schedule. Therefore, DASA always seeks to compute a feasible schedule within which tasks can accrue as much benefit as possible.

Thus, at each scheduling event, DASA examines tasks in the task ready queue in decreasing order of their benefit densities. The algorithm then inserts each task into a tentative schedule at its deadline-position and checks the feasibility of the schedule. Tasks are maintained in increasing deadline-order in the tentative schedule. If inserting a task into the tentative schedule results in an infeasible schedule, then the algorithm removes the task from the schedule.

DASA repeats this process until all tasks in the ready queue have been examined. The algorithm then selects the earliest deadline task in the tentative schedule (which will be at the "head" of the schedule) as the next task to be executed. Note that, if all task deadlines can be satisfied, then DASA's output will be the same as that of EDF.

The rationale behind DASA's examination of tasks in decreasing order of their benefit densities is that, if tasks are examined in decreasing order of their benefit densities, then, at any instant in time, the task that is included in the schedule can accrue the maximal benefit per unit time among the set of nonexamined tasks. This will increase the chance of collecting as much benefit per unit time as possible, thereby increasing the likelihood of maximizing aggregate benefit.

Given  $n$  tasks in the task ready queue, DASA incurs a worst-case complexity of  $O(n^2)$  at each scheduling event. The computationally intensive steps of the algorithm include 1) sorting the tasks to order them in decreasing order of their benefit densities, which costs  $O(n \lg(n))$ , and 2) testing for the feasibility of the tentative schedule each time a task is inserted into the schedule.

The feasibility-testing costs  $O(n^2)$ . This is because, to determine the feasibility of a schedule, each task deadline must be examined in the increasing order of task deadlines. Each task deadline must then be compared against the cumulative sum of the remaining execution times of all tasks with lesser deadlines than the task deadline. This means that, to determine the feasibility of the schedule, DASA makes an  $O(n)$  pass over the task list. Since the algorithm tests for feasibility of the tentative schedule every time a task is inserted into the schedule, the resulting cost becomes  $O(n^2)$ .

Thus, the bottleneck of DASA's scheduling cost is the repeated feasibility test, which costs  $O(n)$  for a single test. We believe that one possible way to speed up the algorithm is to determine the feasible task subset without the repeated feasibility test.

#### 3.1 Observations and Heuristic

Our basic observation is that the tasks in the ready queue fall into one of the following three classes:

- **Class I:** tasks that will never appear in  $\Gamma$ ;
- **Class II:** tasks that will definitely be in  $\Gamma$ ; and
- **Class III:** tasks that may appear in  $\Gamma$ , i.e., those that do not belong to Class I or Class II.

The MDASA algorithm identifies tasks in Class I and Class II by determining if inserting a task  $T_i$  into  $\Gamma$  can maintain the feasibility of  $\Gamma$ . In general, most of the tasks in the ready queue can be identified as either a Class I or

Class II task and, hence, the possibility of making a different scheduling decision from that of DASA is small.

If a task  $T_i$  belongs to Class III, then it is impossible to determine whether the feasibility of  $\Gamma$  can be maintained without performing a feasibility test. In this case, MDASA *heuristically* calculates the probability that  $T_i$  will be included in  $\Gamma$ . We denote this probability as  $P_i$ .

Before describing MDASA, we first formulate the observations that identify the tasks in Class I and Class II as follows:

**Observation 1.** *A task  $T_i$  cannot be in the feasible schedule  $\Gamma$  if  $S_i(t) < 0$ .*

**Proof.** It follows that  $T_i$  will fail the feasibility test. Thus, only tasks with nonnegative slack could be in the feasible schedule.  $\square$

**Observation 2.** *Let  $T_f$  be the first task that has a nonnegative slack in decreasing benefit density order among all tasks. Then,  $T_f$  must be in the feasible schedule  $\Gamma$ .*

**Proof.** Note that the DASA algorithm examines the task ready queue in decreasing order of benefit densities. Thus, at the point of examining the task  $T_f$ , the tentative schedule  $\Gamma$  (which is feasible) is initialized to be empty and no task has been successfully inserted. Therefore, the expanded schedule  $\Gamma \cup T_f = \{T_f\}$  should be a feasible schedule as well.  $\square$

Now that  $T_f$  is in the feasible schedule  $\Gamma$ , any other task  $T_j$  that would cause task  $T_f$  to miss its deadline should be eliminated from  $\Gamma$ , as stated in Observation 3.

**Observation 3.** *If task  $T_j$  satisfies  $d_j < d_f$  and  $R_j(t) > S_f(t)$ , then inserting  $T_j$  into  $\Gamma$  will cause task  $T_f$  to miss its deadline.*

**Proof.** Since  $d_j < d_f$ ,  $T_j$  will be inserted before  $T_f$  in  $\Gamma$ . Hence, the start time of  $T_f$  will be postponed by time  $R_j(t)$ . This postponement implies that  $T_f$  will miss its deadline if  $R_j(t) > S_f(t)$ .  $\square$

**Observation 4.** *If a task  $T_k$  that has a nonnegative slack satisfies  $d_k > d_{max}$ , where  $d_{max}$  is the latest deadline in  $\Gamma$ , then  $T_k$  can be in  $\Gamma$  if and only if*

$$\frac{\sum_{T_i \in \Gamma} R_i(t) + R_k(t)}{d_k - t} \leq 1.0. \quad (1)$$

**Proof.** Recall that the feasible schedule  $\Gamma$  is ordered by increasing task absolute deadlines. Thus, task  $T_k$  will be inserted at the *tail* of  $\Gamma$ , which cannot affect the feasibility of any subset of  $\Gamma$ . Therefore, the feasibility of task set  $T_k \cup \Gamma$  only depends upon whether the cumulative processor time demand until  $d_k$  exceeds the available time ( $d_k - t$ ). This condition becomes (1).  $\square$

In practice, the scheduling algorithm may keep track of the total remaining execution time of all tasks currently in  $\Gamma$ , i.e.,  $\sum_{T_i \in \Gamma} R_i(t)$ , which results in constant complexity of using (1). Also notice that this feasibility condition is sufficient and necessary. Thus, the task  $T_k$  defined as above must belong to either Class I or Class II.

Once tasks that belong to Class I and Class II are identified, we can now proceed to the remaining tasks that constitute Class III.

Recall that, in DASA, tasks are inserted into the tentative schedule at their deadline positions [16]. That is, tasks in the tentative schedule are deadline-ordered. If the insertion of a new task  $T_i$  at its deadline position will cause any of the tasks that are already in the schedule  $\Gamma$  to miss its deadline,  $T_i$  should not be inserted into  $\Gamma$ . This procedure is called the “feasibility test” that ensures the feasibility of the tentative schedule. Therefore, for a task  $T_i$  in Class III, whether or not it should be included in  $\Gamma$  depends upon how significantly the insertion of  $T_i$  will affect the feasibility of  $\Gamma$ .

Furthermore, we observe that early deadline tasks are inserted near the beginning of the schedule by DASA. Our working hypothesis is that tasks with earlier deadlines are more likely to interfere with other tasks, causing them to miss their deadlines. Although we have not examined any other heuristics, the success of our resulting algorithms suggests this to be true (see Section 5).

We denote the number of tasks in the current  $\Gamma$  that have later deadlines than  $d_i$  as  $k$ , which can be approximated by the following equation:

$$k = |\Gamma| \times \frac{d_{max} - d_i}{d_{max} - d_{min}}, \quad (2)$$

where  $d_{max}$  and  $d_{min}$  are the latest and earliest deadlines among all tasks currently in  $\Gamma$ , respectively. Thus,  $k$  is the approximated number of tasks that will be affected by inserting  $T_i$ .

At this stage, it is not clear how  $P_i$  can be calculated based on the task parameters and  $k$ . Though there are a number of possible methods to calculate  $P_i$ , MDASA adopts a straightforward way by computing  $P_i$  as  $P_i = \frac{1}{k}$ . Our motivation for using this simple heuristic is to minimize the overhead of the scheduler. Another advantage is that  $P_i$  is automatically normalized, i.e.,  $P_i$  is guaranteed to be within the range of  $(0, 1]$  since  $k \geq 1$ .<sup>4</sup>

Note that the  $k$  value computed by (2) may not be a good approximation for task queues with nonuniformly distributed deadlines. For example, if the majority of the tasks have deadlines close to the earliest deadline and only a few tasks have very late deadlines, the calculation of  $k$  using (2) tends to overestimate the impact of the insertion of a new task. Consequently, the MDASA algorithm may overly reject tasks and thus may result in poor performance. We show the performance of MDASA under various experimental scenarios in Section 5.

### 3.2 Description of the MDASA Algorithm

Note that MDASA only schedules independent tasks. Thus, it is triggered at two scheduling events: 1) the arrival of a new task and 2) the termination of the currently executing task. Other scheduling events of the original DASA algorithm, namely, *request resource* and *release resource*, are not needed by MDASA.

To facilitate the computation of scheduling decisions, MDASA maintains the task ready queue in the order of task

4. If the calculated  $k$  is less than 1, it should be rounded to 1 because it is the number of tasks to be affected.

benefit densities. When a new task arrives, MDASA inserts the task into the ready queue at its benefit density position. If the processor is not idle, the algorithm extracts the currently executing task from the ready queue, updates its remaining execution time and its benefit density, and then reinserts the task into the ready queue at its new benefit density position. To select the next task to execute, MDASA invokes the algorithm *MDASA\_schedule* (Algorithm 1).

#### Algorithm 1 MDASA\_schedule

```

1: Input: Benefit density ordered task ready queue  $BQ$ ;
2: Initialize  $\Gamma \leftarrow \{\}$ ,  $C \leftarrow 0$ ,  $T_r \leftarrow \phi$ ; /*  $T_r$  is the
   selected task */
3: for each task  $T_i \in BQ$  in descending order of  $BD_i$  do
4:   if  $S_i(t) \geq 0$  then
5:     if  $\Gamma = \{\}$  then
6:       /*  $T_f$  is the first nonnegative slack task */
7:        $\Gamma \leftarrow \Gamma \cup T_i$ ,  $d_{max} \leftarrow d_i$ ,  $d_{min} \leftarrow d_i$ ,  $T_f \leftarrow T_i$ ;
8:     if  $d_i > d_{max}$  then
9:       if  $(C + R_i(t)) / (d_i - t) \leq 1.0$  then
10:         $\Gamma \leftarrow \Gamma \cup T_i$ ,  $d_{max} \leftarrow d_i$ ,  $C \leftarrow C + R_i(t)$ ;
11:     else
12:       if  $d_i < d_f$  then
13:         if  $R_i(t) < S_f(t)$  /*  $S_f(t)$  is the slack of task  $T_f$  /
           then
14:           ProbInclude( $\Gamma, T_i$ );
15:         else
16:           ProbInclude( $\Gamma, T_i$ );
17: return  $T_r$ ;

```

In the event of a task termination, MDASA removes the currently executing task from the ready queue. The algorithm then invokes *MDASA\_schedule* to select the next task to execute.

As shown in Algorithm 1, MDASA accepts a (benefit density-ordered) task ready queue as input and selects the next task to be executed. The algorithm examines the tasks in the ready queue in decreasing order of their benefit densities. By Observation 1, only tasks with nonnegative slacks can be in the feasible schedule. Thus, MDASA silently discards all tasks with negative slacks.

If a task  $T_i$  with a nonnegative slack has the latest deadline, (1) is used to check if the insertion of  $T_i$  can maintain the feasibility of  $\Gamma$ . Otherwise, its deadline is compared with the deadline of task  $T_f$ , i.e.,  $d_f$ .

Recall that  $T_f$  is the first nonnegative slack task in decreasing order of benefit density. By Observation 3, any task  $T_i$  that will cause task  $T_f$  to miss its deadline should be eliminated from  $\Gamma$ . Also recall that any task that does not belong to Class I or Class II falls into Class III. Thus, tasks in Class III are identified in the procedure of identifying Class I and Class II tasks in the task ready queue. These Class III tasks are further examined by the *ProbInclude* algorithm (Algorithm 2), which computes the probability  $P_i$  and determines if  $T_i$  should be included in  $\Gamma$ . Though a single scheduling decision in Algorithm 2 depends on the output of the *rand()* function and is thus random, this randomness is amortized if the algorithm is invoked repeatedly for a large number of times. Similar randomized scheduling policy has been employed in algorithms such as Lottery Scheduling [44].

TABLE 1  
Task Ready Queue Scheduled by MDASA at  $t = 10$

Tasks	$B_i$	$BD_i(t)$	$d_i$	$d_i - t$	$R_i(t)$	$S_i(t)$	Class	in $\Gamma$
$T_1$	360	60	15	5	6	-1	I	No
$T_2(T_f)$	250	50	20	10	5	5	II	Yes
$T_3$	80	40	14	4	2	2	III	Yes
$T_4$	180	30	18	8	6	2	I	No
$T_5$	200	20	22	12	10	2	I	No
$T_6$	20	10	30	20	2	18	II	Yes

#### Algorithm 2 ProbInclude

```

1: Input:  $\Gamma$  and associated state variables  $d_{min}$ ,  $C$ , and  $T_r$ ;
   task  $T_i$ ;
2:  $P_i \leftarrow (d_{max} - d_{min}) / (|\Gamma| \times (d_{max} - d_i))$ ;
3: if rand() <  $P_i$  /* rand() returns a random value
   within  $[0, 1]$  */ then
4:    $\Gamma \leftarrow \Gamma \cup T_i$ ;  $C \leftarrow C + R_i(t)$ ;
5:   if  $d_{min} > d_i$  then
6:      $d_{min} \leftarrow d_i$ ;  $T_r \leftarrow T_i$ ; /*  $T_r$  is the selected task */

```

Note that the latest and earliest task deadlines in  $\Gamma$  are computed in the procedure for computing the  $P_i$ s. Also, the task  $T_r$  is selected whenever a task  $T_i$  is determined in  $\Gamma$  and has the earliest deadline. Therefore, the MDASA algorithm, unlike the original DASA algorithm, only needs one pass and does not invoke the repeated feasibility test.

We now show an example of how the MDASA algorithm makes the scheduling decision. Assume that the MDASA algorithm is triggered at  $t = 10$ . The triggering event could be an arrival of a new task or the termination of the currently executing task. In either case, the resulting task ready queue (benefit density ordered) is shown in Table 1.

Task  $T_1$  is eliminated from  $\Gamma$  due to its negative slack time. Task  $T_2$  has positive slack time and is the first nonnegative slack task in descending order of benefit density. Thus, it becomes  $T_f$  and is included in  $\Gamma$ .

Task  $T_3$  has a positive slack and an earlier deadline than  $T_f$ . By the *MDASA\_Schedule* algorithm, its remaining execution time is compared against the slack time of  $T_2$ . Since  $R_3(10) < S_2(10)$ , task  $T_3$  belongs to Class III. Whether  $T_3$  should be included in  $\Gamma$  depends upon the output of *ProbInclude*( $\Gamma, T_3$ ). In this example, assume that  $T_3$  is included in  $\Gamma$ .

Task  $T_4$  has an earlier deadline than  $T_2$  and satisfies  $R_4(10) > S_2(10)$ . Therefore, the insertion of  $T_4$  will cause  $T_2$  to miss its deadline. Task  $T_5$  has the latest deadline, but fails the test by (1). On the contrary, task  $T_6$  has the latest deadline and enough slack, i.e.,  $(C + R_6(10)) / (d_6 - 10) \leq 1.0$ , where  $C = R_2(10) + R_3(10) = 7$ . The selected task  $T_r$  is the earliest deadline task in  $\Gamma$ , which is  $T_3$  in this example.

### 3.3 Worst-Case Complexity of MDASA

To analyze the worst-case complexity of MDASA, we consider a maximum of  $n$  tasks.

If the scheduler is triggered by a task arrival, the insertion of the new task involves searching for the correct task position in the schedule. This is followed by an insert operation. Using a binary search, the worst-case complexity of these operations is given by  $O(\lg(n)) + O(n) = O(n)$ .




Fig. 1. Example load calculation. (a) Load before removing  $T_1$ . (b) Load after removing  $T_1$ .

In the event of a nonidle processor, the state variables of the currently running task  $T_r$  need to be updated and  $T_r$  must be replaced in the ready queue to maintain the benefit density order. The task  $T_r$  can be kept track of by saving the index of  $T_r$  in a variable. Thus, the replacement operation of  $T_r$  could be implemented by searching the correct position for the updated  $BD_r$  and a “swap” operation. This costs  $O(\lg(n))$ .

The MDASA scheduler then invokes the *MDASA\_schedule* algorithm to make a scheduling decision. Algorithm *MDASA\_schedule* examines the task ready queue from the highest benefit density task to the lowest benefit density task. The worst-case complexity of *MDASA\_schedule* is  $O(n)$  as *MDASA\_schedule* makes only a single pass over the ready queue.

Therefore, the total worst-case complexity of MDASA at a scheduling event is given by  $O(n) + O(\lg(n)) + O(n) = O(n)$ .

#### 4 THE MLBESA ALGORITHM

LBESA [29] is another best-effort real-time scheduling algorithm. It is similar to DASA in that both algorithms schedule tasks using the notion of benefit densities and are equivalent to EDF during underload situations. However, the algorithms differ in the way they reject tasks during overload situations. In [16], Clark shows that DASA is generally better than LBESA in terms of aggregate accrued task benefit.

While DASA examines tasks in the ready queue in decreasing order of their benefit densities for determining feasibility, LBESA examines tasks in the increasing order of task deadlines. Like DASA, LBESA also inserts each task into a tentative schedule at its deadline-position and checks the feasibility of the schedule. Tasks are maintained in increasing deadline-order in the tentative schedule. If the insertion of a task into the tentative schedule results in an infeasible schedule, then, unlike DASA, LBESA removes the *least benefit density task* from the tentative schedule. LBESA continuously removes the least benefit density task from the tentative schedule until the tentative schedule becomes feasible. Once all tasks in the ready queue have been examined and a feasible tentative schedule is thus constructed, LBESA selects the earliest deadline task from the tentative schedule.

Again, the worst-case complexity of LBESA is  $O(n^2)$ , given  $n$  tasks in the ready queue. Furthermore, the major

component of the  $O(n^2)$  complexity is the repeated feasibility test performed by the algorithm before rejecting each task.

To speed up LBESA, our approach is to maintain the order of rejecting tasks and to *heuristically* determine the tasks to be rejected without invoking the feasibility test. We now discuss our observations, the resulting MLBESA heuristic, and its worst-case complexity in the subsections that follow.

##### 4.1 Observations and Heuristic

Note that LBESA continuously rejects tasks in increasing benefit density order if there is an overload. This means that if  $BD_i > BD_j$  and task  $T_i$  is rejected, then task  $T_j$  must also have been rejected.

For convenience, we denote the set of tasks that must be rejected to produce a feasible task subset as  $V$ . Therefore, to determine the tasks in  $V$ , we only need to determine the size of  $V$ .

Without loss of generality, we assume that the tasks are deadline-ordered in the set  $T = \{T_1, T_2, \dots, T_n\}$ . Furthermore, recall that a *feasible* task set satisfies  $\rho \leq 1$ , where  $\rho$  is the load of the task set [13]. Fig. 1a plots the load of a set of five example tasks specified in Table 2.

Our first observation on the load definition is formulated as the following:

**Observation 5.** Let  $\rho_m(t)$  be the load calculated during the time interval  $[t, d_m]$  such that it is the maximum load among all  $\rho_i(t)$  of the task set. Then, rejecting one or more tasks such that  $\rho_m(t) \leq 1.0$  is necessary, but not sufficient, to ensure the feasibility of the remaining tasks.

**Proof.** The necessity of this observation is obvious. We use the task set specified in Table 2 to clarify that it is not sufficient. Assume that task  $T_1$  is rejected due to its

TABLE 2  
A Set of Five Tasks

Tasks	$R_i(t)$	$d_i - t$	$\rho_i(t)$
$T_1$	4	5	0.8
$T_2$	5	6	1.5
$T_3$	3	10	1.2
$T_4$	17	20	1.45
$T_5$	1	25	1.2

TABLE 3  
Removing  $T_1$  from the Task Set

Tasks	$R_i(t)$	$d_i - t$	$\rho_i(t)$
$T_2$	5	6	0.83
$T_3$	3	10	0.8
$T_4$	17	20	1.25
$T_5$	1	25	1.04

minimum benefit density. Before  $T_1$  is rejected,  $\rho_m = \rho_2 = 1.5$ . After  $T_1$  is removed from the task set,  $\rho_2$  decreases to 0.83, as shown in Table 3. However,  $\rho_4$  and  $\rho_5$  are still larger than 1.0 and, hence, the remaining task set is not feasible. Fig. 1b plots the processor load after removing task  $T_1$ .  $\square$

Therefore, we can partition the task subset  $V$  into two parts:  $V = V_1 \cup V_2$ , where  $V_1$  is the task subset that must be rejected to decrease  $\rho_m$  to be less than 1.0, and  $V_2$  is the task subset that must be rejected to ensure the feasibility of the remaining tasks.

To determine the decrease in load that is acquired by rejecting one task, we use the following observation:

**Observation 6.** After rejecting task  $T_j$ , the processor load calculated during interval  $[t, d_i]$  is given by:

$$\rho'_i = \begin{cases} \rho_i - \frac{R_j(t)}{d_i - t} & \text{if } d_j \leq d_i \\ \rho_i & \text{otherwise,} \end{cases} \quad (3)$$

where  $\rho_i$  and  $\rho'_i$  are the calculated load before and after the task rejection, respectively.

**Proof.** If  $d_j \leq d_i$ , then the total processor time demand that needs to be completed before  $d_i$  must be subtracted from  $R_j(t)$  in the event of removing task  $T_j$ . It follows that

$$\rho'_i = \frac{\sum_{d_k \leq d_i} R_k(t) - R_j}{d_i - t} = \rho_i - \frac{R_j(t)}{d_i - t}.$$

In the case of  $d_j > d_i$ , the processor time demand of task  $T_j$  is not included in  $\rho_i$ . Thus, removing  $T_j$  cannot affect  $\rho_i$ .  $\square$

Therefore, a natural way to identify and reject tasks in the subset  $V_1$  is to continuously reject tasks, from the lowest benefit density task to the highest benefit density task, until  $\rho_m$  is less than 1.0. After each task is rejected,  $\rho_m$  is recomputed using (3).

Since rejecting tasks in  $V_1$  is necessary, but not sufficient, to ensure the feasibility of the remaining tasks, we *heuristically* determine and reject another task subset  $V_2 = V - V_1$ . Observe that a high  $\rho_m$  means that there is a severe overload situation for the task set. Thus, it is reasonable to use  $\rho_m$  to heuristically determine how many tasks should be rejected. These tasks are rejected to resolve the overload situation, as a higher  $\rho_m$  generally implies that more tasks should be rejected (though not necessarily).

We define the term *Reject Ratio* as the ratio of the number of tasks to be rejected from the subset  $(T - V_1)$ .

**Definition 1.** The reject ratio of subset  $(T - V_1)$  is defined as follows:

$$\eta = \frac{|V_2|}{|T - V_1|}. \quad (4)$$

Note that determining the exact number of tasks in  $V_2$  involves performing the feasibility test, which is not desired. Thus, we approximate the exact reject ratio with the help of  $\rho_m$ . This approximated reject ratio, indicated as  $\eta'$ , is computed as follows:

$$\eta' = \frac{\rho_m - 1}{\rho_m} = 1 - \frac{1}{\rho_m} \leq 1. \quad (5)$$

Now that the reject ratio is computed, we can determine the size of the subset  $|V_2|$  by Definition 1:

$$|V_2| = \eta' \times |T - V_1| = \eta' \times (|T| - |V_1|). \quad (6)$$

We choose this simple heuristic due to its implementation efficiency and its property of automatic normalization. It is worth noting that, under certain workload scenarios, this heuristic may not perform well. For example, if a task set is only slightly overloaded and the deadline order of a task queue is in the reverse order of its benefit density order, i.e., later deadline tasks have lower benefit densities, LBESA may still need to reject most of the tasks, whereas MLEBSA only rejects a small number of tasks due to a small  $\eta'$ . We examine the performance of the MLBESA algorithm in Section 5.

Once the size of  $V_1$  and  $V_2$  are determined, the algorithm simply rejects  $(|V_1| + |V_2|)$  tasks, from the lowest benefit density task to the highest benefit density task. The remaining tasks are conjectured to be feasible without performing any feasibility test.

Note that, if  $\rho_m$  is extremely large, the approximated reject ratio  $\eta'$  will be close to 1. In this case, all tasks should be rejected, which implies that no feasible schedule exists. Then, the algorithm selects the earliest deadline task from the remaining tasks, if there is any. In the example 5-task set,  $|V_1| = 1$  because rejecting task  $T_1$  is enough to reduce  $\rho_2$  to 0.83. Furthermore, the reject ratio is approximated as  $\eta' = 1 - \frac{1}{\rho_2} = 0.33$ . Therefore,  $|V_2| = \eta' \times (|T| - |V_1|) = 0.33 \times 4 = 1.33$ .

## 4.2 Description of the MLBESA Algorithm

The MLBESA scheduling algorithm is also triggered by task arrival and task termination events. However, unlike the MDASA algorithm, MLBESA orders the task ready queue in increasing deadlines to facilitate the initial feasibility test.

When a new task arrives, MLBESA inserts the task into the task ready queue at its deadline position and then invokes the Algorithm *MLBESA\_schedule* (Algorithm 3) for determining the next task to execute. If the algorithm is triggered by the termination of a task, MLBESA removes the currently running task from the task ready queue and then invokes the Algorithm *MLBESA\_schedule*.

### Algorithm 3 MLBESA\_schedule

- 1: **Input:** Absolute deadline ordered task ready queue  
 $DQ \leftarrow \{T_1, T_2, \dots, T_n\}$ ;
- 2: Initialize internal variables:  
 $\rho_m \leftarrow 0$ ;  $C \leftarrow 0$ ;  $T_r \leftarrow \phi$ ;  $|V_1| \leftarrow 0$ ,  $|V_2| \leftarrow 0$ ;
- 3:  $t \leftarrow \text{get\_current\_time}()$ ;
- 4: **for** each task  $T_i \in DQ$  in nondecreasing order of deadlines **do**

```

5:  $C \leftarrow C + R_i(t)$ ; /*  $C$  is the total processor time
   demand */
6:  $\rho_i(t) \leftarrow C / (d_i - t)$ ;
7: if  $\rho_i(t) > \rho_m$  then
8:    $\rho_m \leftarrow \rho_i(t)$ ;  $d_m \leftarrow d_i$ ;
9: if  $\rho_m \leq 1.0$  then
10:  return  $T_r \leftarrow T_1$ ;
11:else
12:  Sort  $DQ$  in another queue  $BQ$ , by benefit densities;
13:   $\eta' \leftarrow 1 - 1/\rho_m$ ;
14:  for each task  $T_i \in BQ$  in ascending order of benefit
   densities do
15:   Remove task  $T_i$  from  $BQ$ ,
    $BQ \leftarrow BQ - T_i$ ;  $|V_1| \leftarrow |V_1| + 1$ ;
16:   if  $d_i \leq d_m$  then
17:      $\rho_m \leftarrow \rho_m - R_i(t) / (d_m - t)$ ;
18:     if  $\rho_m \leq 1.0$  then
19:        $|V_2| \leftarrow \eta' \times (|DQ| - |V_1|)$ ;
20:       break;
21:   for  $i \leftarrow 1$  to  $|V_2|$  do
22:     Remove the lowest benefit density task from  $BQ$ ;
23:   if  $BQ \neq \emptyset$  then
24:     Select  $T_r \in BQ$ , which has the earliest deadline among
     all tasks within  $BQ$ ;
25:  return  $T_r$ ;

```

The *MLBESA\_schedule* algorithm first checks the feasibility of the task set by examining the tasks (in the ready queue) in their deadline-order. The algorithm identifies the maximum load factor  $\rho_m$  and its associated deadline  $d_m$ . If no overload is detected, i.e.,  $\rho_m \leq 1$ , then the algorithm simply applies the EDF algorithm to select the next task to be executed, which is task  $T_1$  at the “head” of the deadline-ordered queue.

In the event of an overload, the algorithm first computes the reject ratio  $\eta'$ . Tasks are then continuously rejected in ascending order of their benefit densities until  $\rho_m \leq 1$ . The number of tasks rejected so far is the size of the subset  $V_1$ . Once  $|V_1|$  tasks are rejected, *MLBESA\_schedule* determines the size of the subset  $|V_2|$  using (6). Note that the tasks in  $V_1 \cup V_2$  are temporarily removed from the task ready queue. Such tasks may be determined to be feasible at future scheduling events. After  $|V_1| + |V_2|$  tasks have been rejected, the algorithm selects the earliest deadline task from the remaining task set, if any exists.

### 4.3 Worst-Case Complexity of MLBESA

To analyze the complexity of MLBESA, we consider a deadline-ordered ready queue, denoted as  $DQ$ . We store  $DQ$  as a linear array. Thus, inserting a new task involves searching for the correct deadline position of the task in  $DQ$ , followed by the insert operation. Given  $n$  tasks in  $DQ$ , this has a worst-case complexity of  $O(\lg(n)) + O(n) = O(n)$ .

MLBESA first checks the feasibility of the task set by making a single pass over  $DQ$ , which costs  $O(n)$ . If no overload is detected, the algorithm returns the earliest deadline task in  $DQ$ , i.e., the task at the head of  $DQ$ . In this case, the complexity of MLBESA is  $O(n)$ .

However, in the event of an overload, the task queue needs to be reordered according to benefit densities. Tasks

in subset  $V_1 \cup V_2$  are then continuously rejected by the algorithm to produce a feasible subset. In the worst case, all tasks in the ready queue could be rejected. Thus, the complexity of MLBESA during overload is the sum of the cost of sorting the task ready queue (according to benefit densities) and rejecting the tasks. This total cost becomes  $O(n \lg(n)) + O(n) = O(n \lg(n))$ .

Clearly, the worst-case complexity of MLBESA is dominated by the complexity during the overload situation. Therefore, the worst-case complexity of MLBESA is  $O(n \lg(n))$ .

It is interesting to note that the algorithm could also order the task ready queue according to benefit densities, as in the case of the MDASA algorithm. This benefit density ordered-ready queue will help to reduce the complexity of rejecting the tasks during an overload situation, but will incur the same total worst-case complexity. Furthermore, such a benefit density-ordered task ready queue must be reordered according to deadlines. Our approach of ordering the task ready queue by deadlines may avoid the expensive sorting operation if there is no overload. Therefore, the deadline order scheme has a better average-case complexity.

## 5 PERFORMANCE EVALUATION

In experimentally evaluating MDASA and MLBESA, our main goal is to determine how well the algorithms perform with respect to their counterpart DASA and LBESA scheduling algorithms, respectively. The performances of MDASA and MLBESA are also compared against the  $D^{over}$  algorithm [25] and the BE-v algorithm proposed in [33]. A detailed discussion of related algorithms can be found in Section 6. Furthermore, we are also interested in determining the task response times under MDASA and MLBESA. We conducted simulation studies to determine these performance metrics using synthetic workloads.

### 5.1 Performance of MDASA and MLBESA

To evaluate the performance of the scheduling algorithms, we considered tasks with randomly distributed parameters. Each experimental setting was characterized by parameters including task execution time  $C_i$ , slack  $S_i$ , timeliness benefit  $B_i$ , and task interarrival time  $I_i$ .

Table 4 and Table 5 summarize the baseline experimental settings for exponential distributions and normal distributions, respectively. Note that the task slack is specified as the ratio of its slack time to its corresponding execution time. Once  $C_i$  and  $S_i$  of a task  $T_i$  are determined, the relative deadline of the task is given by  $D_i = C_i + S_i$ . Furthermore, the task interarrival time depends on the average load  $\rho_A$ <sup>5</sup>. Thus, only the ratio  $I_i/C_i$  is specified.

Each simulation experiment lasted 8,000 seconds and generated a stream of tasks, based on the specified task parameters. We conducted experiments for  $\rho_A = 0.1 \sim 2.0$ . The number of tasks generated during each experiment varied from approximately 1,000 to 43,000, due to the change in the average load. To eliminate random effects,

5. This average processor load definition has been used in the literature such as [12], [16], which may or may not be directly related to the processor load defined in Section 2.



TABLE 4  
Task Parameters (Exponential Distributions)

Para	Distribution	Mean	Std. Dev.
$C_i$	exponential	0.5 sec	0.5 sec
$\frac{S_i}{C_i}$	exponential	0.25	0.25
$B_i$	exponential	10	10
$\frac{I_i}{C_i}$	exponential	$\frac{1.0}{\rho}$	$\frac{1.0}{\rho}$

each experiment was independently repeated five times for the same task stream.

We measured the *Deadline Satisfaction Ratio (DSR)* and *Accrued Benefit Ratio (ABR)* produced by MDASA and MLBESA. DSR is defined as the ratio of the number of satisfied deadlines to the total number of task releases. Fig. 2a and Fig. 2b show the DSR and ABR of MDASA and those of DASA under two baseline experimental settings. As we can see, MDASA performs almost exactly the same as DASA for both experimental settings. The difference in the DSR and ABR of the two algorithms is found to be generally less than 5 percent. This result implies that the MDASA algorithm almost always makes the same scheduling decisions as those of DASA.

To evaluate the performance of the scheduling algorithms for highly bursty workload, we also considered the general Pareto distribution for task parameters. The general Pareto distribution is interesting because its variance could be infinity given a finite mean value—the maximal nondeterminism one can have.<sup>6</sup> Furthermore, the general Pareto distribution is the base for synthesizing self-similar communication traffic that can be found in real-time applications such as compressed video transmission (see [26] for an introduction to self-similar traffic).

In our baseline Pareto distribution experiments, task execution times, slack times, interarrival times, and task benefit values all follow the general Pareto distribution. For task execution time, the average task execution time is 0.5 sec and the  $\alpha$  value is 1.9. Note that, once the average value and the  $\alpha$  value are determined,  $\beta$  of the distribution can be calculated as  $Mean \times (\alpha - 1)/\alpha$ . Similarly, the average slack time is five times that of the task execution time—a long slack time that allows tasks to remain feasible for a long time—and the average task benefit is 10. Furthermore, the  $\alpha$  values for the distributions of task interarrival times and task benefit are all set to 1.1.

We show DSRs and ABRs of MDASA and DASA schedulers under the general Pareto distribution in Fig. 2a and Fig. 2b, respectively. Compared with the performance under exponential and normal distributions, both DASA and MDASA satisfy a smaller percentage of task deadlines. Furthermore, DASA slightly outperforms MDASA in terms of DSR. This performance gap is echoed in Fig. 2b, where DASA also performs better than MDASA. In addition, we observe that both DASA and MDASA accrue more benefit under the general Pareto distribution than under normal or exponential. This is because the

6. A general Pareto distribution is governed by a shape parameter  $\alpha$  and a scale parameter  $\beta$  that determine the minimal value of the random variable. If  $\alpha$  is between 1 and 2, the variance of the distribution is infinity and the mean value is  $\beta\alpha/(\alpha - 1)$  [43].

TABLE 5  
Task Parameters (Normal Distributions)

Para	Distribution	Mean	Std. Dev.
$C_i$	normal	0.5 sec	0.5 sec
$\frac{S_i}{C_i}$	normal	0.25	0.25
$B_i$	normal	10	10
$\frac{I_i}{C_i}$	normal	$\frac{1.0}{\rho}$	$\frac{1.0}{\rho}$

general Pareto distribution may produce very large task benefit values for a few tasks. Thus, satisfying deadlines of these high-benefit tasks can dramatically increase the algorithm performance in terms of ABR.

Since MDASA uses the *rand()* function to heuristically compute the schedule, it is possible that random factors can affect the algorithm performance given a set of of input tasks. Thus, for a given input task stream, we repeat the experiments with five independent random seeds that are used by the MDASA algorithm. These five experiments are grouped together as they all have the same input. Standard deviations of DSRs and ABRs of the MDASA scheduler are then calculated to evaluate the effects of random factors.

We show the average standard performance deviations of MDASA for the five groups of experiments in Fig. 3 (using the baseline Pareto distributions). As shown in the figure, the performance of MDASA is reasonably stable, i.e., the performance standard deviations are less than 2 percent in our experiments. Therefore, we conclude that the performance of MDASA is close to that of DASA for a broad range of workload scenarios.

Furthermore, Fig. 4 shows the ratios of tasks that fall into each class of the MDASA algorithm under the baseline general Pareto distribution. We observe that the ratio of Class I tasks, i.e., at any given scheduling event, the number of Class I tasks to the total number of tasks in the ready queue remains approximately 0.1 when load is greater than 0.6.

On the contrary, the ratio of Class II tasks monotonically decreases while the ratio of Class III tasks increases. This is because a heavy load tends to generate long ready queues and more tasks cannot be identified as Class I or Class II. Furthermore, the increasing ratio of Class III tasks explains why MDASA performs worse than DASA during heavy load. We also observe that, even when the system is heavily overloaded, i.e., load is 2.0, and the workload pattern is highly bursty, i.e., task parameters follow the general Pareto distribution, approximately half of the tasks still fall into either Class I or Class II.

We conducted the same experiments for MLBESA and LBESA. The DSRs and ABRs of MLBESA versus LBESA are shown in Fig. 5a and Fig. 5b, respectively. We observe that, under normal and exponential distributions, MLBESA is only slightly worse than LBESA. However, under the general Pareto distribution, LBESA may accrue as much as 10 percent more benefit than MLBESA—a visible performance gap. We conjecture that this is because, under a highly bursty workload pattern, MLBESA may reject too many tasks that have high benefit but long execution time as well. Thus, these rejected tasks have nonnegligible impacts on the aggregate benefit ratio.

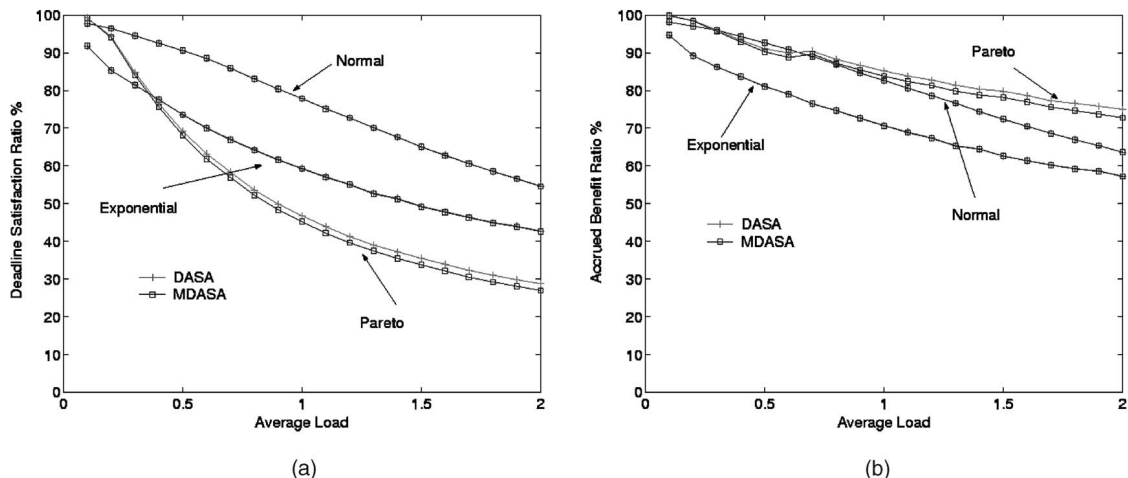


Fig. 2. MDASA versus DASA (baseline settings). (a) DSRs of MDASA and DASA. (b) ABRs of MDASA and DASA.

Furthermore, in Fig. 6, we show the average ratio of tasks that belong to task set  $V_1$  and the average reject ratio computed by MLBESA. As can be seen from the figure, both ratios increase when load increases. It is also worth noting that no more than 35 percent percentage of all tasks (at load 2.0) belong to task set  $V_1$  that are discarded to possibly produce feasible schedules. Recall that the MLBESA algorithm heuristically determines if a task in the remaining task set  $(T - V_1)$  is in the feasible schedule. Thus, a relatively small  $|V_1|/T$  ratio (compared with MDASA's Class I and Class II ratios in Fig. 4) implies small chances for MLBESA to make different scheduling decisions than LBESA. This result explains the performance gap between LBESA and MLBESA in Fig. 5b.

As a comparison, we implemented the  $D^{over}$  algorithm [25] and the BE-v algorithm [33] in the simulation experiments.

The  $D^{over}$  algorithm has been proven to be optimal in the sense that it achieves the upper bound of the competitive factor for any online scheduling algorithm. Note that the competitive factor of an online algorithm measures its *worst-case* performance. However, the optimal competitive factor of  $D^{over}$  does not imply the best *average* performance

because  $D^{over}$  may make pessimistic scheduling decisions so that it can guarantee the optimal competitive factor.

The BE-v algorithm was inspired by the  $D^{over}$  algorithm. Furthermore, Mosse et al. show that BE-v outperforms the original LBESA algorithm on short teaser tasks.

We first show the DSRs and ABRs of the six algorithms under the baseline general Pareto distributions. Our reason for choosing the general Pareto distribution is because a highly busy workload helps to manifest the performance difference between different algorithms, as we learned from previous experiments for MDASA and MLBESA. However, the results for other distributions, such as normal distribution and exponential distribution, all have similar trends and, thus, are omitted here.

As shown in Fig. 7, the DSRs of the algorithms do not have a significant difference. However, DASA, LBESA, BE-v, and MDASA perform significantly better than MLBESA and  $D^{over}$  in terms of ABR (see Fig. 8). This result is consistent with that shown in Fig. 5b, i.e., LBESA performs better than MLBESA. Furthermore, the fact that  $D^{over}$  performs poorly has been described [12].

Finally, we study task response times under various algorithms. For the purpose of comparison, we use a metric called "Response Time Ratio" that is defined as the ratio of

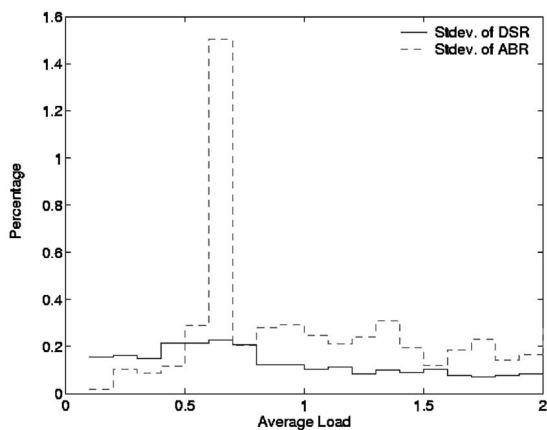


Fig. 3. Standard deviations of MDASA performance.

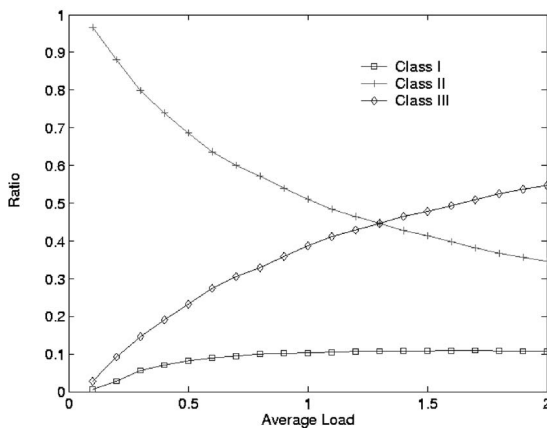


Fig. 4. Percentages of tasks in MDASA classes.

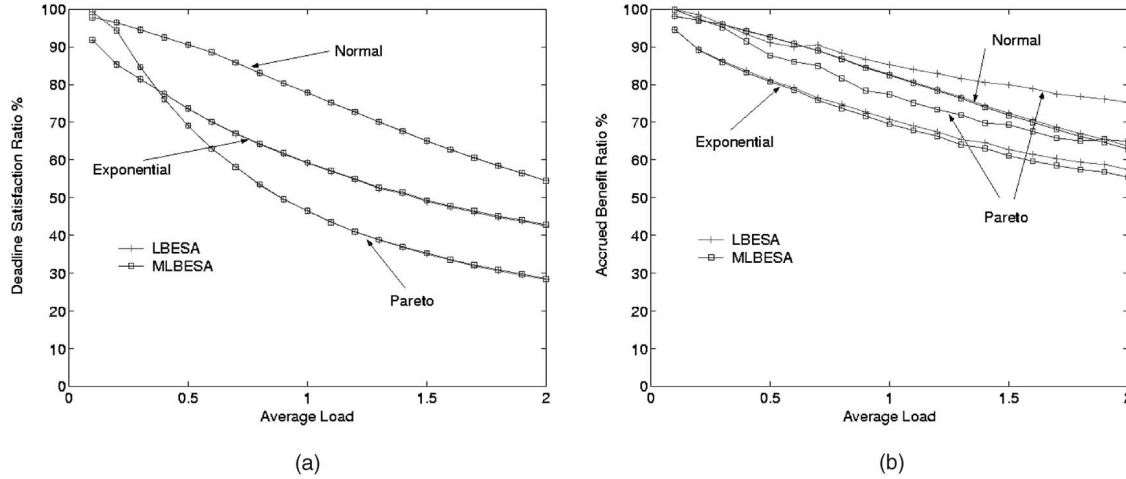


Fig. 5. MLBESA versus LBESA (baseline settings). (a) DSRs of MLBESA and LBESA. (b) ABRs of MLBESA and LBESA.

the actual task response time to its execution time. This definition does not depend on the execution times of particular tasks and thus enables a fair comparison of large number of tasks.

In Fig. 9, we show the average response time ratios of the six algorithms under the baseline general Pareto distributions. As shown in the figure, most of the response time ratios are close, especially for light to medium load conditions, i.e., load no more than 0.8. However, when the load increases, BE-v and  $D^{over}$  generally produce longer response times than other algorithms.

## 5.2 Task Response Times under MDASA and MLBESA

The experimental results in Section 5.1 show that MDASA and MLBESA, in general, have close performance to their counterpart algorithms. We conjecture that MDASA and MLBESA make almost the same scheduling decisions as those made by DASA and LBESA at most scheduling events, respectively. Consequently, most of the task response times under MDASA and MLBESA should be close to that under DASA and LBESA, respectively.

If task response times under MDASA and MLBESA are almost the same as those under DASA and LBESA, respectively, then MDASA and MLBESA can be used as approximate response time analysis algorithms for estimating response times under DASA and LBESA. This would be a useful property of the algorithms as determining exact response times under best-effort schedulers such as DASA and LBESA is impossible without constructing entire task schedules [19]. Such schedule construction is computationally expensive due to the high cost of DASA and LBESA.

Thus, we could use MDASA and MLBESA to estimate response times under DASA and LBESA, respectively, in a much faster way as MDASA and MLBESA have low computational costs.<sup>7</sup>

To measure how close the task response times are under different scheduling algorithms, we define the relative error of task response times as follows:

$$RelError = \begin{cases} \frac{Resp - P\_resp}{Resp} & Resp < \infty \text{ and } P\_resp < \infty \\ 100\% & Resp = \infty \text{ and } P\_resp < \infty \\ -100\% & Resp < \infty \text{ and } P\_resp = \infty \\ 0 & Resp = \infty \text{ and } P\_resp = \infty \end{cases}, \quad (7)$$

where  $Resp$  and  $P\_resp$  are the response times of a task under DASA (or LBESA) and MDASA (or MLBESA),

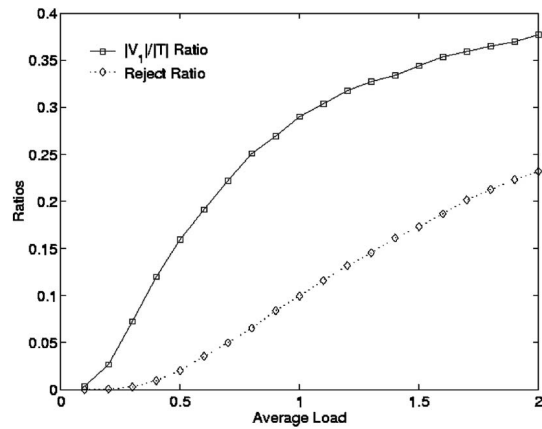


Fig. 6. Percentages of tasks in MLBESA categories.

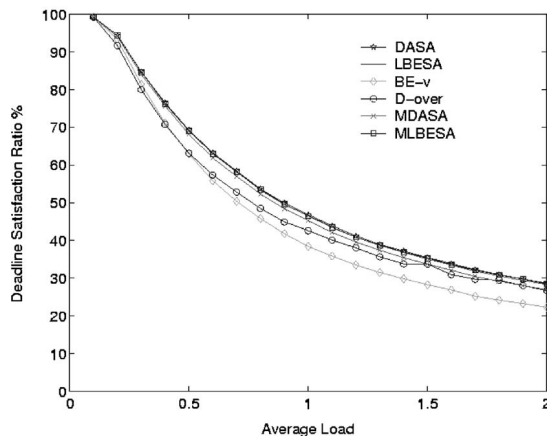


Fig. 7. Comparison of DSRs (Pareto distribution).

7. In fact, this has been one of our primary motivations for developing MDASA and MLBESA in the first place (as described in Section 1).




Fig. 8. Comparison of ABRs (Pareto distribution).

respectively. Note that, in the case of a deadline miss, the task response time becomes infinity. This is because an infeasible task is never scheduled by schedulers such as DASA and LBESA and thus will be aborted. Since task response is defined as the time interval between the finishing time of a task and its arrival time, the response time of an aborted task is calculated as infinity. This case is also handled by the above error definition.

We first show the mean and standard deviation of the relative errors under MDASA in Fig. 10. The experimental data are collected with the baseline exponential distribution of task parameters shown in Table 4. Each error bar in the plot corresponds to the *Mean* and *StdDev* at one load point, which is centered at the *Mean* and is  $2 \times StdDev$  long.

From Fig. 10, we see that the *Means* are always very close to zero. We also observe that the relative errors are more spread out as the load increases, which is implied by the larger *StdDev*. Thus, there are larger, and possibly more, errors during heavy load.

To illustrate the distribution of the relative errors, we show a sample histogram of MDASA's response time relative errors at  $\rho_A = 2.0$  (for the baseline exponential distribution experiment) in Fig. 11. We observe that more than 95 percent of the relative errors are very close to 0.

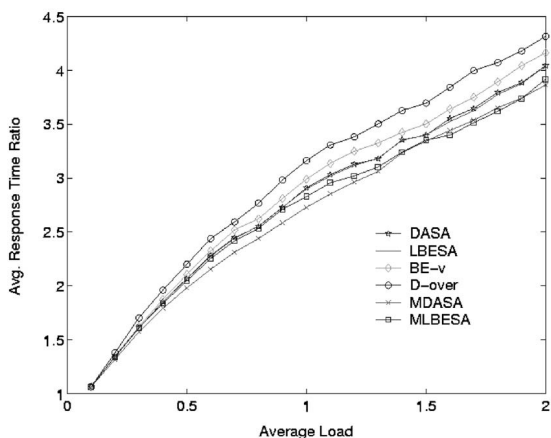


Fig. 9. Comparison of response ratios (Pareto distribution).




Fig. 10. Response time errors under MDASA.

Furthermore, we see that there are few very large errors, 100 percent or -100 percent for example.

Thus, the relatively large *StdDev* could be due to the few large errors. The error bar plot is, hence, misleading. For capturing the error distribution more accurately, we use an extended *outlier* measurement. We define a 5 percent *Hit Ratio* as the ratio of the relative errors within  $\pm 5\%$  to the total number of samples. A 10 percent *Hit Ratio* and a 20 percent *Hit Ratio* have similar definitions.

Fig. 12a shows the 5 percent, 10 percent, and 20 percent hit ratios for the same baseline exponential experiment under MDASA at  $\rho_A = 2.0$ . Being consistent with the histogram, the hit ratios are all approximately 99 percent. This means that almost all the task response times under MDASA are very close to the values under DASA. The hit ratio plots under MLBESA are shown in Fig. 12b.

### 5.3 Impact of Task Parameters on Response Times

In this section, we study the impact of task parameters on task response times. Since one of our objectives is to use the task response times obtained under the MDASA and MLBESA algorithms as an approximation for the response times that can be obtained under their deterministic counterparts, we are interested in determining how relative errors in response times are affected by task parameters.




Fig. 11. Histogram of MDASA response time errors.




Fig. 12. Hit ratios under (a) MDASA and (b) MLBESA.




Fig. 13. Impact of slack times on MDASA hit ratios. (a) Impact of mean slack. (b) Impact of StdDev slack.

We studied the impact of task parameters in two ways: 1) by determining the effect on *Hit Ratio*, defined in Section 5.2 and 2) by determining the correlation between task parameters and response time errors. In general, a distribution is characterized by its central tendency and dispersion, which are usually specified as *Mean* and *StdDev*. We chose the normal distribution as the basic distribution in these experiments because the *Mean* and *StdDev* of a normal distribution are independent parameters. Furthermore, the normal distribution may be approximated by the Central Limit Theorem [32] in many real applications. The general Pareto distribution is also used in the experiments because it allows studying the effect of infinite variance.

In each experiment, the *Mean* or *StdDev* of one of the four task parameters  $C_i$ ,  $S_i$ ,  $B_i$ , and  $I_i$ , is varied while the other parameters are kept the same, as shown in Table 5.

### 5.3.1 Impact of Timing Parameters

Our first experiment compared the 10 percent hit ratio of MDASA for 0.5 seconds of mean slack and for 0.25 seconds of mean slack. The results are shown in Fig. 13a. The impact of the *StdDev* of the slack for MDASA is shown in Fig. 13b.

We observe that either a larger *Mean* or a larger *StdDev* of task slack slightly degrades the performance of the algorithm. A larger *StdDev* implies that the tasks are more “bursty.” Thus, we conjecture that the length of the task ready queue may vary in a larger range. Recall that MDASA *heuristically* determines the probability that a task is in the feasible schedule based on the queue length. Therefore, MDASA has a greater chance of making a different scheduling decision from that of DASA if there is a long ready queue. Furthermore, we use the “truncated” normal distribution, i.e., only positive values are allowed, to generate the task slack since we expect a task to be feasible at the time of its arrival. Thus, increasing the *Mean* could also potentially generate a more bursty task stream. However, the performance degradation of the algorithm is found to be very small in both cases.

We observed that the impact of task execution times and interarrival times follows a similar pattern as that of the slack. That is, doubling *Mean* or doubling *StdDev* of the task execution time or interarrival time may result in a slightly worse performance, e.g., 1 percent degradation in terms of 10 percent hit ratio. We also observed similar results for MLBESA. For brevity, we omit these plots here.




Fig. 14. Average queue length under MDASA.

To reinforce the results presented in Fig. 13a and Fig. 13b, we also conducted experiments for studying the effect of task slack times on algorithm performance. These experiments are conducted with the baseline general Pareto distribution. Furthermore, we changed the mean slack times to two times (slack2), five times (slack5), and nine times (slack9) as that of task execution times.

Fig. 14 shows the average queue length produced by MDASA. Note that infeasible tasks are dropped from the ready queue to avoid infinite ready queue. These plots clearly show that, besides workload governed by task execution times and interarrival times, task slack times have significant impact on the length of ready queues. Furthermore, in Fig. 15, we show DSRs achieved by MDASA for the baseline Pareto distribution with various slack times. As can be seen from the figure, a large slack time implies a less stringent timing constraint and, thus, yields better performance.

5.3.2 Impact of Benefit

Unlike the task timing parameters such as  $C_i$ ,  $S_i$ , and  $I_i$ , the real-time benefit of a task indicates the importance level of a task relative to other tasks in the system. Thus, we conjecture that the performance of the algorithm is benefit scale-invariant.




Fig. 15. Impact of slack times on MDASA performance.


Fig. 16a and Fig. 16b show the 10 percent hit ratio curves with different Mean and StdDev of real-time benefit for the MDASA algorithm. As we expect, the performance of the algorithm does not change significantly.

5.3.3 Correlation between Task Parameters and Response Time Errors

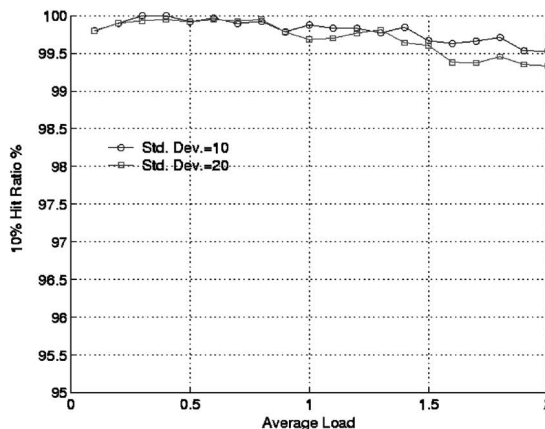
We are also interested in determining whether there is any significant correlation between task parameters and response time errors. For example, does a task with a longer execution time has a larger response time error than others? We consider the Pearson and Spearman correlations [32] for measuring this aspect.

Table 6 and Table 7 show the correlation coefficients between the task parameters and the response time relative error at  $\rho_A = 0.2$  and  $\rho_A = 2.0$ , respectively.

Note that the task interarrival time depends upon the load and the execution time of the task stream. Thus, it is not shown as a separate parameter in the tables. The Pearson correlation measures the linear relationship between two quantities, which may be a too strict requirement in our case. But, the rank-order correlation measured by Spearman correlation does not show any significance either.



(a)



(b)

Fig. 16. Impact of benefit on MDASA performance. (a) Impact of Mean benefit. (b) Impact of StdDev benefit.

TABLE 6  
Correlations at  $\rho_A = 0.2$

Correlation	$S_i$	$C_i$	$B_i$
Pearson	0.0041	0.0518	0.0036
Spearman	0.0137	0.0371	0.0075

Therefore, we conclude that the parameters of a task do not have a significant correlation with the response time difference between the MDASA and DASA algorithms. The correlation coefficients under MLBESA also lead us to the same conclusion.

## 6 RELATED WORK

Value or benefit-based overload scheduling has been studied in different contexts [2], [6], [10], [11], [14], [20], [25], [39], [3], [7], [8], [9], [13], [33], [34], [37].

Baruah et al. established the  $1/(1 + \sqrt{k})^2$  upper bound on the competitive factor of any online scheduling algorithm, given the *importance ratio* of  $k$  [6]. Note that  $k$  is defined as the maximum task value density divided by the minimum task value density among the task set. This upper bound is achieved by the  $D^{over}$  algorithm presented in [25]. However, the competitive factor of an online algorithm only measures the *worst-case* performance of the algorithm, which may not be directly related to the *average* performance for a large number of random tasks. Furthermore, our experimental results shown in Section 5.1 suggest that  $D^{over}$  may perform worse than MDASA and MLBESA for random task sets, which is also described in [12].

Baruah et al. also proposed two metrics for measuring the performance of a scheduling algorithm during overload situations. This includes the “Effective Processor Utilization (EPU)” and “Completion Count (CC)” [9]. It is shown that no online algorithm is competitive with respect to CC during overload situations [9]. In the case of “uniform value,” where the value of a task is equal to its execution time, optimal scheduling with known minimum slack factor and multiple processors is explored in [8] and [7], respectively.

Besides optimal algorithms such as  $D^{over}$ , heuristic algorithms have also been proposed. In [39], the authors consider the problem of computing schedules in a dynamic environment by ensuring the feasibility of the system. Furthermore, several algorithms similar to the LBESA algorithm have been designed [2], [13], [33]. These algorithms are similar to LBESA in that they reject tasks by ascending order of task benefit density or variant metrics of benefit density, to resolve any overload situation.

The RED (Robust Earliest Deadline) algorithm proposed by Buttazzo and Stankovic [13] combines many features, including graceful performance degradation during overload, deadline tolerance, and resource reclaiming. Since we assume that the execution time is the exact processor time demand of a task, no resource reclaiming mechanism is relevant to our task model. Similarly, no deadline tolerance is allowed for step benefit functions. Note that the RED algorithm itself does not specify any particular policy to

TABLE 7  
Correlations at  $\rho_A = 2.0$

Correlation	$S_i$	$C_i$	$B_i$
Pearson	0.0025	0.0101	0.0335
Spearman	0.02	0.0208	0.0314

reject tasks when overload occurs. In the case of rejecting tasks in the ascending order of task benefit densities, the RED scheduler behaves like an LBESA scheduler. Thus, we do not directly compare the performance of MDASA and MLBESA with the RED algorithm.

Mosse et al. derive several variants of the LBESA algorithm in [33]. They also show that the BE-v variant algorithm may outperform LBESA when the value accrued by rejecting short teasers is greatest. Our simulation experiments show that MDASA, MLBESA, and BE-v have very close performance. Furthermore, a comparative study shows that robust algorithms, such as RED, perform better than plain algorithms, such as plain EDF, and guarantee-based algorithms [11]. Therefore, we focus our performance comparison on only a class of robust algorithms, where the BE-v algorithm serves as an example.

Aldarmi and Burns propose the concept of “timeliness-function” in [3]. Unlike the step benefit function, which drops at the task deadline, the timeliness-function decreases when the slack time of a task reaches zero. In [2], the authors show that scheduling the task with the highest Dynamic Timeliness-Density (DTD) is more effective than scheduling the highest benefit density task. However, the “timeliness” of a task has already been examined in the feasibility test of DASA and LBESA, which aborts any task with negative slack. Thus, DTD scheduling bears a lot of similarities to the DASA algorithm. Burns et al. also study the problem of assigning value to tasks in a systematic and rational way [10], [34].

Apart from the task model with one segment of execution per task, the concept of “imprecise computations” has also been proposed as an effective technique to handle overloads [28]. The imprecise computation model assumes that a task consists of a mandatory part that must finish execution before its deadline and an optional part that can be executed to improve the quality of the computation’s output results.

In [31], Mejia-Alvarez et al. consider this imprecision computation model and use combinatorial algorithms to find near-optimal solutions. The work on feedback control theory scheduling further extends the imprecise computation model and assumes the presence of  $N$  versions of the same task ( $N \geq 2$ ) [30]. The MDASA and MLBESA algorithms do not assume the existence of multiple versions of the same task.

In [35], the author presents an overload management scheme that can satisfy deadlines of at least  $m$  instances of a periodic task within  $k$  consecutive releases, which is called the  $(m, k)$  *firm guarantee*. However, no benefit is defined for the application tasks, which fundamentally differs from our model.

In [14], [20], the authors consider the problem of scheduling a set of nonpreemptive tasks to maximize the

accrued benefit, which complements the work of scheduling independent tasks. Richardson and Sarkar consider scheduling tasks *before* overload but *after* a fault occurs [37], where a fault could be the reason of the overload situation. Our task model does not assume this knowledge.

Furthermore, the majority of the algorithms with comparable performance to DASA and LBESA repeatedly determine the task set feasibility and, hence, are expensive. The proposed MDASA and MLBESA algorithms, on the other hand, heuristically determine the feasible task subset and are much faster.

Although probabilistic analysis has been used in real-time scheduling, such as in the efforts [4], [18], [42], these algorithms mainly deal with scheduling real-time processes that have varying computational demand.

The original LBESA algorithm itself is the closest work to the MDASA and MLBESA algorithms. LBESA statistically "guesses" if there is an overload, based on the variance of the total process slack time. If there is an overload, LBESA removes the least benefit density process from the feasible process set that it computes, where the benefit density of a process is the process benefit divided by its remaining execution time. This procedure is repeated until the algorithm "guesses" and concludes that there is no overload. MDASA and MLBESA algorithms avoid this iterative decision procedure of LBESA by heuristically determining the processes in the *final* feasible schedule.

## 7 CONCLUSIONS

In this paper, we present two fast, best-effort, real-time scheduling algorithms, called MDASA and MLBESA. MDASA and MLBESA are novel in the way that they heuristically, yet accurately, mimic the behavior of the DASA and LBESA algorithms, respectively, but are faster with  $O(n)$  and  $O(n \lg(n))$  worst-case complexities. MDASA and MLBESA reason about the behavior of DASA and LBESA by heuristically determining a feasible schedule of the task ready queue.

Our experimental results show that MDASA and MLBESA perform almost as well as DASA and LBESA, respectively, unless the workload is highly bursty and the system is heavily overloaded. Furthermore, the task response times under MDASA and MLBESA are found to be very close to the values under their counterpart scheduling algorithms. While MDASA performs better than MLBESA and has a better worst-case complexity, MLBESA guarantees the optimal schedule during underload situations. Our prototype implementation of MDASA/MLBESA in a middleware [27] also suggests their effectiveness, e.g., response time analysis using MDASA/MLBESA can be 100 ~ 200 msec faster than that using DASA/LBESA.

## ACKNOWLEDGMENTS

This work was supported by the US Office of Naval Research under Grant N00014-99-1-0158 and N00014-00-1-0549.

## REFERENCES

[1] T. Abdelzaher, "An Automated Profiling Subsystem for QoS-Aware Services," *Proc. IEEE EEE Real-Time Technology and Applications Symp. (RTAS)*, pp. 208-217, 2000.

[2] S. Aldarmi and A. Burns, "Dynamic Value-Density for Scheduling Real-Time Systems," *Proc. Euromicro Conf. Real-Time Systems*, pp. 270-277, June 1999.

[3] S.A. Aldarmi and A. Burns, "Time-Cognizant Value Functions for Dynamic Real-Time Scheduling," Technical Report YCS-306, Dept. of Computer Science, Univ. of York, U.K., 1998.

[4] A. Atlas and A. Bestavros, "Statistical Rate Monotonic Scheduling," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 123-132, 1998.

[5] S. Baruah, G. Koren, D. Mao, B. Mishra, A. Raghunathan, and L. Rosier, "On the Competitiveness of On-Line Task Real-Time Task Scheduling," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 106-115, 1991.

[6] S. Baruah, G. Koren, B. Mishra, A. Raghunathan, L. Rosier, and D. Shasha, "On-Line Scheduling in the Presence of Overload," *Proc. IEEE Ann. Symp. Foundations of Computer Science*, pp. 101-110, 1991.

[7] S.K. Baruah, "Overload Tolerance for Single-Processor Workloads," *Proc. IEEE EEE Real-Time Technology and Applications Symp. (RTAS)*, pp. 2-11, 1998.

[8] S.K. Baruah and J. Haritsa, "Robust: A Hardware Solution to Real-Time Overload," *Proc. ACM SIGMETRICS Conf. Measurement and Modeling of Computer Systems*, vol. 21, no. 1, pp. 207-216, 1993.

[9] S.K. Baruah, J.R. Haritsa, and N. Sharma, "On-Line Scheduling to Maximize Task Completions," *J. Combinatorial Math. and Combinatorial Computing*, vol. 39, pp. 65-78, 2001.

[10] A. Burns et al., "The Meaning and Role of Value in Scheduling Flexible Real-Time Systems," *J. Systems Architecture*, vol. 46, pp. 305-325, 2000.

[11] G. Buttazzo, M. Spuri, and F. Sensini, "Value vs. Deadline Scheduling in Overload Conditions," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 90-95, 1995.

[12] G.C. Buttazzo, *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Academic Publishers, 1997.

[13] G.C. Buttazzo and J. Stankovic, "Adding Robustness in Dynamic Preemptive Scheduling," *Responsive Computer Systems*, D. Fussel and M. Malek, eds., pp. 67-88, Kluwer Academic, Oct. 1995.

[14] K. Chen and P. Muhlethaler, "A Scheduling Algorithm for Tasks Described by Time Value Function," *J. Real-Time Systems*, vol. 10, no. 3, pp. 293-312, 1996.

[15] R. Clark et al., "An Adaptive, Distributed Airborne Tracking System," *Proc. IEEE Int'l Workshop Parallel and Distributed Real-Time Systems*, pp. 353-362, 1999.

[16] R.K. Clark, "Scheduling Dependent Real-Time Activities," PhD dissertation, Carnegie Mellon Univ., 1990.

[17] M.L. Dertouzos, "Control Robotics: The Procedural Control of Physical Processes," *Proc. IFIP Congress*, pp. 807-813, 1974.

[18] M.K. Gardner, "Probabilistic Analysis and Scheduling of Critical Soft Real-Time Systems," PhD dissertation, Univ. of Illinois at Urbana-Champaign, 1999.

[19] T. Hegazy and B. Ravindran, "Using Application Benefit for Proactive Resource Allocation in Asynchronous Real-Time Distributed System," *IEEE Trans. Computers*, vol. 51, no. 8, pp. 945-962, Aug. 2002.

[20] S.-I. Hwung, C.-M. Chen, and A.K. Agrawala, "Scheduling an Overloaded Real-Time System," *Proc. IEEE Int'l Phoenix Conf. Computers and Comm.*, pp. 22-28, June 1996.

[21] E.D. Jensen and J.D. Northcutt, "Alpha: A Non-Proprietary Operating System for Large, Complex, Distributed Real-Time Systems," *Proc. IEEE Workshop Experimental Distributed Systems*, pp. 35-41, 1990.

[22] E.D. Jensen and B. Ravindran, "Guest Editor's Introduction to Special Section on Asynchronous Real-Time Distributed Systems," *IEEE Trans. Computers*, vol. 51, no. 8, pp. 881-882, Aug. 2002.

[23] E.D. Jensen, C.D. Locke, and H. Tokuda, "A Time-Driven Scheduling Model for Real-Time Systems," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 112-122, Dec. 1985.

[24] G. Koob, "Quorum," *Proc. DARPA ITO General PI Meeting*, pp. A-59-A-87, Oct. 1996.

[25] G. Koren and D. Shasha, "D-over: An Optimal On-Line Scheduling Algorithm for Overloaded Real-Time Systems," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 290-299, 1992.

[26] W.E. Leland, M. Taqqu, W. Willinger, and D.V. Wilson, "On the Self-Similar Nature of Ethernet Traffic," *IEEE/ACM Trans. Networking*, vol. 2, no. 1, pp. 1-15, Feb. 1994.

[27] P. Li, B. Ravindran, J. Wang, and G. Konowicz, "Choir: A Real-Time Middleware Architecture Supporting Benefit-Based Proactive Resource Allocation," *Proc. IEEE Int'l Symp. Object-Oriented Real-Time Distributed Computing (ISORC)*, pp. 292-299, 2003.



- [28] J.W.S. Liu, W.-K. Shih, K.-J. Lin, R. Bettati, and J.-Y. Chung, "Imprecise Computations," *Proc. IEEE*, vol. 82, no. 1, pp. 83-94, Jan. 1994.
- [29] C.D. Locke, "Best-Effort Decision Making for Real-Time Scheduling," PhD dissertation, Carnegie Mellon Univ., 1986.
- [30] C. Lu, J. Stakovic, G. Tao, and S. Son, "Feedback Control Real-Time Scheduling: Framework, Modeling and Algorithms," *J. Real-Time Systems*, vol. 23, nos. 1/2, pp. 85-126, July 2002.
- [31] P. Mejia-Alvarez, R. Melhem, and D. Mosse, "An Incremental Approach to Scheduling during Overloads in Real-Time Systems," *Proc. IEEE Real-Time Systems Symp. (RTSS)*, pp. 283-293, 2000.
- [32] W. Mendenhall and T. Sincich, *Statistics for the Engineering and Computer Sciences*, third ed. 1991.
- [33] D. Mosse, M.E. Pollack, and Y. Ronen, "Value-Density Algorithm to Handle Transient Overloads in Scheduling," *Proc. Euromicro Conf. Real-Time Systems*, pp. 278-286, 1999.
- [34] D. Prasad and A. Burns, "A Value-Based Scheduling Approach for Real-Time Autonomous Vehicle Control," *Robotica*, vol. 31, pp. 273-279, 2000.
- [35] P. Ramanathan, "Overload Management in Real-Time Control Applications Using the (m, k) Firm Guarantee," *IEEE Trans. Parallel and Distributed Systems*, vol. 10, no. 6, pp. 549-559, June 1999.
- [36] B. Ravindran, "Engineering Dynamic Real-Time Distributed Systems: Architecture, System Description Language, and Middleware," *IEEE Trans. Software Eng.*, vol. 28, no. 2, pp. 30-57, Jan. 2002.
- [37] P. Richardson and S. Sarkar, "Adaptive Scheduling: Overload Scheduling for Mission Critical Systems," *Proc. IEEE EEE Real-Time Technology and Applications Symp. (RTAS)*, pp. 14-23, 1999.
- [38] B. Shirazi, L.R. Welch, B. Ravindran, C. Cavanaugh, and E. Huh, "Dynbench: A Benchmark Suite for Dynamic Real-Time Systems," *J. Parallel and Distributed Computing Practices*, vol. 3, no. 1, pp. 89-107, Mar. 2000.
- [39] J.A. Stankovic and K. Ramamritham, "The Spring Kernel: A New Paradigm for Real-Time Systems," *IEEE Software*, vol. 8, no. 3, pp. 62-72, May 1991.
- [40] D. Stewart and P. Khosla, "Real-Time Scheduling of Sensor-Based Control Systems," *IFAC/IFIP Workshop Real-Time Programming*, pp. 144-150, May 1991.
- [41] The Open Group, *MK7.3a Release Notes*. Cambridge, Mass.: The Open Group Research Inst., 1998.
- [42] T.-S. Tia et al., "Probabilistic Performance Guarantee for Real-Time Tasks with Varying Computation Times," *Proc. IEEE EEE Real-Time Technology and Applications Symp. (RTAS)*, pp. 164-173, 1995.
- [43] D. von Seggern, *CRC Standard Curves and Surfaces*. CRC Press, 1993.
- [44] C.A. Waldspurger and W.E. Wehl, "Lottery Scheduling: Flexible Proportional-Share Resource Management," *Operating Systems Design and Implementation*, pp. 1-11, 1994.



**Peng Li** is a PhD candidate in the Department of Electrical and Computer Engineering at the Virginia Polytechnic Institute and State University. He received the BS degree in automatic control (1997) and ME degree in communication and information systems (2000) from the University of Electronic Science and Technology of China (UESTC). He was the recipient of the "Excellent Student Award" during all of the years of his study at UESTC (1993 to 2000) for his distinguished academic performance. His current research interests include real-time distributed systems, embedded systems, operating systems, and real-time communication. He is also a recipient of a Graduate Student Support Grant from the Ninth IEEE Real-Time and Embedded Technology and Applications Symposium. He is a student member of the IEEE and IEEE Computer Society.



**Binoy Ravindran** received the Master's degree from the New Jersey Institute of Technology (NJIT) in 1994 and the PhD degree from the University of Texas at Arlington (UT Arlington) in 1998, both in computer science. He is an assistant professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech. He received two "Student Achievement Award" from NJIT for outstanding research by a graduate student (1994, 1995) and a "Doctoral Dissertation Award" from UT Arlington for outstanding doctoral research (1998). His current research focus include real-time distributed systems having application-level, end-to-end quality of services (e.g., timeliness, survivability, security). Recent sponsors of his research include The MITRE Corp., US Office of Naval Research, and NASA. During 1999, he was an invited speaker at INRIA, France, and at the ARTES (Swedish National Strategic Research Initiative in Real-Time Systems) Real-Time Week, Sweden. He served as the program chair of the 1999 IEEE International Workshop on Parallel and Distributed Real-Time Systems. He was a co-guest editor for the *IEEE Transactions on Computers* special issue on asynchronous real-time distributed systems. He is a senior member of the IEEE and a member of the IEEE Computer Society.

► For more information on this or any computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).