Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 08-003

Exploring the Throughput-Fairness Tradeoff of Deadline Scheduling
in Heterogeneous Computing Environments

Vasumathi Sundaram, Abhishek Chandra, and Jon Weissman

January 31, 2008

Exploring the Throughput-Fairness Tradeoff of Deadline Scheduling in Heterogeneous Computing Environments

Vasumathi Sundaram, Abhishek Chandra, and Jon Weissman
 Department of Computer Science and Engineering,
 University of Minnesota,
 Minneapolis, MN 55455, USA
 {vassun,chandra,jon}@cs.umn.edu

Abstract—The scalability and computing power of large-scale computational platforms that harness processing cycles from distributed nodes has made them attractive for hosting compute-intensive time-critical applications. Many of these applications are composed of computational tasks that require specific deadlines to be met for successful completion. In scheduling such tasks, replication becomes necessary due to the heterogeneity and dynamism inherent in these computational platforms. In this paper, we show that combining redundant scheduling with deadline-based scheduling in these systems leads to a fundamental tradeoff between throughput and fairness. We propose a new scheduling algorithm called *Limited Resource Earliest Deadline (LRED)* that couples redundant scheduling with deadline-driven scheduling in a flexible way by using a simple tunable parameter to exploit this tradeoff. Our evaluation of LRED using trace-driven and synthetic simulations shows that LRED provides a powerful mechanism to achieve desired throughput or fairness under high loads and low timeliness environments, where these tradeoffs are most critical.

I. INTRODUCTION

Large-scale computational platforms such as Grid infrastructures [1] and cycle sharing systems [2], [3] have grown in popularity for running computationally intensive applications in areas spanning Bio-informatics [4], High Energy Physics [5], Climate prediction [6], etc. These systems provide scalability and enormous computational power by harnessing idle processing cycles from computing hosts distributed around the Internet. Their low deployment and operational cost in addition to their scalability has made these infrastructures attractive for hosting large-scale time-critical applications, for example, Biomedical applications such as medical image processing [7], [8], and Real-time MRI analysis [9].

Many of these applications are composed of computational tasks that require specific deadlines to be met for successful completion. These deadlines could result either from the time-constrained or real-time nature of the applications, or due to internal task dependencies, requiring some of the critical tasks to be finished in a timely manner. However, scheduling such time-constrained tasks in cycle sharing systems is challenging because of several reasons. First, the nodes in such a system are highly heterogeneous, with different CPU speeds, network connectivity, and load conditions. In addition, the capacity of individual nodes is also highly variable due to varying loads,

fluctuating network bandwidth and churn. Such heterogeneity and dynamism makes it extremely difficult to select the right nodes to execute tasks in a timely manner. As a result, redundant scheduling [3], [4], [10], where a task is assigned to multiple nodes to improve its chance of successful completion, is often employed in such systems.

However, the use of redundant scheduling creates a fundamental dilemma in choosing the right order of task scheduling. Giving preference to low deadline tasks, as is done by traditional deadline-based scheduling algorithms such as Earliest Deadline First (EDF) [11], results in consuming more resources, since such tasks have more stringent deadlines and need more resources for their timely completion. On the other hand, ordering the tasks in decreasing order of their deadlines (we refer to this ordering as Latest Deadline First or LDF), while potentially providing better resource utilization, is likely to starve tighter deadline (and hence potentially more important) tasks. This dilemma can be understood as a tradeoff between throughput and fairness in the system: should a scheduler focus on successfully completing more tasks, or should it partition the available resources more equitably among tasks with different deadlines? The goal depends on the specific type of application that is being run on the system.

There are several scenarios in which achieving a higher throughput is more important than ensuring good fairness or vice versa. For instance, image processing applications [9] often have a large number of independent tasks processing different parts of an image database. Similarly, a transportation planning program [12] may employ independent traffic analysis tasks working on different parts of a large traffic matrix. In such cases, the goal is to maximize the total number of tasks executed by the system within a time period. On the other hand, in workflow-based applications, such as those in the domains of bioinformatics [13] and astronomy [14], a task with large number of dependencies will tend to have a shorter deadline, and would need to be finished first even at the expense of overall throughput, while improving fairness (as we will show in this paper). Service-oriented computing is another scenario where tasks belonging to different requests or services will need to be executed fairly, irrespective of their individual deadlines.

In this paper, we propose a new scheduling algorithm called

Limited Resource Earliest Deadline (LRED) that is specifically designed to address this throughput-fairness tradeoff in heterogeneous, dynamic computational environments. *LRED couples redundant scheduling with deadline-driven scheduling* in a flexible way to exploit this tradeoff. Intuitively, LRED works by limiting the number of resources consumed per task (thus improving throughput), while scheduling the selected tasks in earliest deadline order (thus improving fairness). An important feature of LRED is that it can achieve a desired throughput-fairness tradeoff using a simple tunable parameter.

The design of the LRED algorithm has resulted in the following key research contributions:

- We define a statistical notion of *timeliness* for a computational node which can incorporate both inter-node heterogeneity as well as intra-node dynamism, and provide a simple technique to estimate this timeliness based on the node's past execution history.
- LRED uses these timeliness values to couple redundant scheduling with deadline-driven scheduling in a seamless manner.
- LRED can achieve the desired throughput-fairness tradeoff in the system by using a tunable parameter to control the scheduling order of the tasks. LRED is a generalization of EDF and LDF, so that, by tuning this parameter, LRED reduces to EDF in one extreme, and to (a close variant of) LDF in another extreme.

We use trace-driven and synthetic simulations to evaluate the performance of LRED in a heterogeneous environment under different system conditions such as load and varying overall timeliness level of the system. Our results show that the load and the timeliness environment has a significant impact on the throughput-fairness tradeoff of task scheduling. We find that LRED provides a powerful mechanism to achieve desired throughput or fairness, particularly under high loads and low timeliness environments, where these tradeoffs are most critical.

II. SYSTEM MODEL

In this section, we present our system model and define some of the concepts that will be used throughout the rest of the paper.

A. Task model

Our task model consists of a task pool with a set of homogeneous¹ tasks in terms of their computational requirements. These tasks are continuously created and submitted to the pool by an application. Each task T_i is associated with a *deadline* D_i that is defined as the time by which the task T_i must be completed². The task is deemed as completed successfully, only if a result is computed within time D_i from the instant the task arrived in the system, otherwise, it is assumed to be failed. Different tasks can have different deadlines, either based on

¹Many of the applications we consider produce homogeneous tasks. We intend to explore heterogeneous tasks in the future.

²By task completion time, we mean the time at which the task result is returned back to the scheduler. We will use response time and task completion time interchangeably in the rest of the paper, unless otherwise noted.

Fig. 1. Computing the timeliness of two workers for a given deadline D .

their time of arrival, or because of difference in importance or priorities (e.g., a task with large number of task dependencies may be assigned a shorter deadline).

B. Computational model

The computing environment consists of a set of heterogeneous worker nodes that provide their computational resources for executing the tasks. We assume a pull-based task scheduling model (commonly used in large-scale computational systems such as BOINC [3]), where each worker node requests work from a central scheduler and is assigned a task from the existing task pool. It then executes the task and return the results back to the scheduler. Different worker nodes can take different amounts of time to complete the same task due to differences in their computational capabilities (e.g., CPU speeds), network bandwidth, and load. Further, each worker may provide different response times for the same task during different periods, due to dynamic conditions such as varying loads and fluctuating network latency. Thus, we assume *heterogeneity across the worker nodes*, as well as *dynamism over the same node over time*, in terms of the completion time for a task.

C. Timeliness model

Based on the above worker and task model, we present a timeliness model that incorporates the inter-node heterogeneity as well as intra-node dynamism of the worker nodes for different task deadlines. We associate a response time distribution with each worker, which models the probability with which a worker is able to finish a task within a given amount of time. Such a distribution can be constructed based on a worker's past execution history. Using this distribution, we can estimate the likelihood that a worker will be able to meet a deadline.

Definition 1: Timeliness: The timeliness $\tau_i(D)$ of a worker W_i for a task with deadline D is defined as the probability that the worker will be able to finish the task within time D .

$$\tau_i(D) = CDF_i(D), \quad (1)$$

where CDF_i is the CDF of the worker's response time. Figure 1 illustrates this notion of timeliness.

Based on this definition, timeliness takes on a value between 0 and 1, where a timeliness value of 1 for a deadline D means that the worker node always returns within time D , while a value of 0 implies that the worker node will never return before D . Further, note that the timeliness of a worker depends on the deadline of a task, so that a worker will have a smaller timeliness value for more stringent deadlines. Finally, different workers will have different timeliness values for the same deadline based on their response time distributions. For example, in Figure 1, for the deadline D , the two workers W_i and W_j have different timeliness values $\tau_i(D)$ and $\tau_j(D)$ respectively.

D. Redundant Scheduling

Because of the worker heterogeneity and different ranges of task deadlines, it is possible that a task with a stringent deadline may have only a small probability of being successfully completed by any worker on its own. However, its success probability can be increased by redundantly allocating it to multiple workers. For instance, let us assume that a task T arrives with a deadline of 100 and there are two workers $W1$ and $W2$ with timeliness values of 0.8 and 0.6 respectively for this deadline. Now, if we want the task to be completed with a high probability, say 0.9, then the task cannot be scheduled to either of the workers individually and be expected to complete successfully with the desired probability. On the other hand, by assigning the task to both $W1$ and $W2$ and waiting for the first response, we increase the probability of successful task completion to 0.92, which meets our requirement.

We assume the use of such redundant scheduling in our system to achieve a desired *target success rate (TSR)* for each task. The target success rate can be defined as the desired probability with which each task must be completed within its deadline. TSR can be thought of as the overall task completion rate desired by an application. Note that a TSR of less than 1 may be acceptable for many applications in our target environment, largely because most of them use soft task deadlines with a possibility to re-execute a failed task in the worst case. Moreover, it may be infeasible to achieve a TSR of 1 in an uncertain environment that we are considering.

The redundancy level required to satisfy a given TSR for a task is then determined by the task deadline and the timeliness of the available workers. In particular, it can be determined using the following notion of group timeliness of a group of workers for a task:

Definition 2: Group Timeliness: The group timeliness $\tau_G(D)$ for a group G of workers $\{W_1, W_2, \dots, W_n\}$ for a task with deadline D is defined as the probability of successful completion of the task within time D by at least one of the workers in G .

$$\tau_G(D) = 1 - \prod_{i=1}^n (1 - \tau_i(D)) \quad (2)$$

Thus, to meet the TSR for a task with deadline D , a redundant scheduling algorithm would assign the task to a group of workers whose group timeliness value exceeds the TSR.

III. COMBINING DEADLINE-DRIVEN AND REDUNDANT SCHEDULING

A. Understanding the Throughput-Fairness Tradeoff

Having defined our system model and the notion of timeliness, we next examine the implications of using redundant scheduling for deadline-based tasks with respect to two key metrics of interest: *throughput* and *fairness*. In the context of deadline-based scheduling, we define throughput as the total number of tasks that are completed within their deadlines. Fairness can be defined as a measure of the share of the worker resources utilized for tasks with different deadlines. In other words, fairness can be thought as capturing the difference in the proportion of tasks completed for different deadlines. The smaller this difference is, the more fair is the scheduling algorithm. We measure the fairness of a schedule using Jain's fairness index [15]³. This simple index possesses many desirable properties of a fairness index compared to other indices such as Min-max ratio, Variance and Coefficient of Variation. It is continuous so that any change in the proportion of tasks completed changes the fairness index. It has normalized values bounded between 0 and 1 so that it can be easily used to compare different algorithms. In addition, it is independent of the number of task deadlines used.

Let us start by examining possible ways of combining deadline-driven scheduling with redundant scheduling. *Earliest Deadline First (EDF)* [11] is a classical scheduling algorithm used for deadline-driven scheduling. EDF always selects the task with the shortest deadline for execution. EDF has been shown to provide an optimal schedule for a uniprocessor environment as long as the tasks do not require more computing power than is available in the system. In other words, EDF can provide the highest throughput in such an environment. The following example illustrates how EDF will perform in our system model, in the presence of redundant scheduling.

Example 1: Consider a set of tasks $T1$, $T2$ and $T3$ in the task pool with successively higher deadlines, and a set of workers $W1$, $W2$, and $W3$. Let us assume that the deadlines of the tasks $T2$ and $T3$ are such that they can be successfully completed with a high probability (based on a given TSR) by any one of the workers. However, the deadline of task $T1$ is so stringent that it needs to be assigned to all three workers to have a high likelihood of timely completion. Since EDF selects the task with minimum deadline, it will schedule $T1$ to $W1$, $W2$, and $W3$, and ignore the other two tasks, resulting in a throughput of 1. Let us see what happens if we use a different scheduler, namely, *Latest Deadline First (LDF)*, that schedules tasks in decreasing order of their deadlines. LDF will assign tasks $T2$ and $T3$ to $W2$ and $W3$ respectively (assuming ties are broken arbitrarily), while the task $T1$ will be ignored, resulting in a throughput of 2.

The above example illustrates the problem of naively coupling redundant scheduling with deadline-based scheduling. Redundant scheduling results in non-uniform resource requirements for different tasks. Since tasks with more stringent

³For ease of exposition, we defer the quantitative definition of Jain's fairness index to Section V, and use fairness here in a more intuitive sense.

Fig. 2. Comparison of LDF with EDF showing the throughput-fairness tradeoff

deadlines consume more resources, giving priority to such tasks will not always result in higher throughput for the system. The above example raises the question whether a scheduler like LDF is more desirable in heterogeneous and non-dedicated environment, as it is likely to achieve higher throughput. We examine this question using the following example.

Example 2: Figure 2 shows an example scenario where at time 0, there are 4 tasks in the task pool: tasks T1 and T2 with deadline D1, and tasks T3 and T4 with deadline D2, where $D2=2 \cdot D1$. Further assume there are two workers W1 and W2 in the system at this time, each of which can finish a task with deadline D2 (with probability TSR), but both of them need to be grouped together to finish a task with deadline D1. With this setup, at time 0, EDF will use both workers to schedule task T1, while LDF will assign one worker each to tasks T3 and T4.

Now, suppose both workers successfully finish their tasks and come back at time D1. At this point, the deadlines of T1 and T2 would have expired, while the effective deadlines of T3 and T4 would have reduced to D1 due to passage of time. In this case, EDF will use both workers to schedule task T3 (since T2 has already missed its deadline), while LDF will have no remaining task to schedule (as T1 and T2 have already missed their deadlines).

Thus, over the two scheduling instances, both EDF and LDF achieve the same throughput of 2. However, while EDF schedules one task for each deadline, LDF schedules both higher deadline tasks, while starving the lower deadline tasks. We can see that EDF provides more fairness for this case which can also be verified using Jain's fairness index [15] which gives a fairness value of 1 for EDF and 0.5 for LDF (higher is fairer).

This example shows that while LDF may have equal or better throughput than EDF in general, it suffers from higher unfairness because of its bias towards longer deadline tasks, at the expense of starving short deadline tasks. Note that even though EDF is biased towards shorter deadline tasks, it is also likely to schedule higher deadline tasks, as their deadlines become more stringent with the passage of time. On the other hand, by preferring more lax deadline tasks initially, LDF further decreases the possibility of eventually executing shorter deadline tasks in the future.

To summarize, the above examples lead to the following key insights:

- **Insight 1:** Scheduling shorter deadline tasks consumes a larger number of resources than higher deadline tasks, resulting in lower throughput.
- **Insight 2:** Scheduling a shorter deadline task before a higher deadline task is likely to achieve higher fairness.

In other words, there is a clear tradeoff between throughput and fairness, when coupling redundant scheduling with deadline-driven scheduling. In particular, EDF and LDF represent the opposite ends of the spectrum in this tradeoff, with one achieving higher fairness while the other achieves higher throughput.

B. Formalization of Key Insights

We now present theoretical results to formalize the insights presented above. For all the results presented in the section, we assume the system model described in Section II, and assume a combination of redundant and deadline-based scheduling, so that a schedule is required to meet each task deadline with a given target success rate (TSR). In addition, we make the following assumptions:

- 1) All the worker nodes in the environment are available at each scheduling instant.
- 2) The timeliness of any worker node does not change for the entire duration of scheduling.
- 3) The set of tasks to be scheduled is fixed.

We omit the explicit mention of these assumptions in the results presented below unless required. Further, proofs for all lemmas and theorems are given in the Appendix.

1) Throughput-related Results:

Lemma 1: If T_i and T_j are two tasks such that $D_i \leq D_j$, and if T_i can be assigned to a worker group G_i , then $\exists G_j \subseteq G_i$ to which T_j can be assigned.

Corollary 1: If T_i and T_j are two tasks such that $D_i \leq D_j$, and if T_i and T_j respectively need at most n and m workers to satisfy the TSR, then $m \leq n$.

Corollary 2: Let $\mathcal{T} = \{T_1, \dots, T_i, \dots, T_k\}$ be the set of k tasks in a schedule that contains task T_i , and let $\mathcal{G} = \{G_1, \dots, G_i, \dots, G_k\}$ be the corresponding set of worker groups assigned to the tasks. Let $\mathcal{T}' = (\mathcal{T} - \{T_i\}) \cup \{T_j\}$, where, $T_j \notin \mathcal{T}$ is a task such that $D_i \leq D_j$, and $\mathcal{G}'$ be the corresponding set of worker groups assigned to the tasks in $\mathcal{T}'$. Then,

$$|\bigcup_{G_l \in \mathcal{G}'} G_l| \leq |\bigcup_{G_l \in \mathcal{G}} G_l|.$$

Lemma 1 and the above corollaries basically correspond to the Insight 1 above that says that shorter deadline tasks consume a larger number of resources than higher deadline tasks.

Theorem 1: At a given scheduling instant, if $\mathcal{S}^{\text{EDF}}$ is a task schedule generated by EDF scheduler and $\mathcal{S}^{\text{LDF}}$ is a task schedule generated by LDF, then $|\mathcal{S}^{\text{EDF}}| \leq |\mathcal{S}^{\text{LDF}}|$, where, $|\mathcal{S}|$ denotes the number of tasks scheduled in $\mathcal{S}$.

This theorem states that the throughput achieved by an LDF scheduler will be higher than that of an EDF scheduler. Intuitively, this result follows from the Insight 1 above.

2) Fairness-related Results:

Theorem 2: If the tasks are divided into two non-empty sets $\mathcal{T}_1$ and $\mathcal{T}_2$, s.t. any task in $\mathcal{T}_1$ has a lower or equal deadline to any task in $\mathcal{T}_2$, and if r_1^{EDF} , r_2^{EDF} and r_1^{LDF} , r_2^{LDF} respectively are the ratio of tasks completed in those bins by EDF and LDF, where $r_1^{EDF}, r_2^{LDF} > 0$, then, $FI_{EDF} \geq FI_{LDF}$, if and only if

$$\frac{r_2^{EDF}}{r_1^{EDF}} \geq \frac{r_1^{LDF}}{r_2^{LDF}}. \quad (3)$$

This theorem basically provides the necessary and sufficient condition for the fairness of EDF to be higher than that of LDF. This condition is illustrated in Figure 3, where Figure 3 (a)

Fig. 3. Necessary condition for Theorem 2

shows a hypothetical LDF schedule with $N_2 = 2 \cdot N_1$, where N_1 is the number of tasks completed in bin B1 and N_2 is the number of tasks completed in bin B2. We assume that each bin has the same total number of tasks. Figures 3(b), (c) and (d) illustrate corresponding EDF schedules, with x additional tasks in the bin $\mathcal{T}_1$ and y fewer tasks in the bin $\mathcal{T}_2$. Figures 3(b) and (c) show cases where the condition (Equation 3) holds, while (d) shows a case violating this condition. From cases (b) and (c), we can observe that the relative imbalance among the two bins is less than or equal to that for case (a), so that the fairness of EDF will be at least as large as that of LDF. However, for case (d), we can observe that the relative imbalance between the bins goes below that for case (a), which would result in a lower fairness for EDF compared to LDF.

We now show that the likelihood of $FI_{EDF} \geq FI_{LDF}$ is high based on the likelihood of the condition 3 holding true. To determine this likelihood, we first establish bounds on the possible behavior of LDF and EDF in terms of the number of tasks each can schedule from sets of different deadline tasks. We use the following lemmas to establish the limits for the values taken by the fractions in Equation (3).

Lemma 2: For two non-empty sets of tasks $\mathcal{T}_1$ and $\mathcal{T}_2$ such that the deadline of any task in $\mathcal{T}_1$ is less than or equal to the deadline of any task in $\mathcal{T}_2$, if LDF completes a ratio of tasks r_1^{LDF} in $\mathcal{T}_1$ and r_2^{LDF} in $\mathcal{T}_2$, and if EDF completes a ratio of tasks r_1^{EDF} in $\mathcal{T}_1$ and r_2^{EDF} in $\mathcal{T}_2$ then,

$$r_1^{EDF} \geq r_1^{LDF}$$

and

$$r_2^{EDF} \leq r_2^{LDF}.$$

Lemma 3: In LDF scheduling, given two non-empty sets of unscheduled tasks $\mathcal{T}_1$ and $\mathcal{T}_2$ such that the deadline of any task in $\mathcal{T}_1$ is less than or equal to the deadline of any task in $\mathcal{T}_2$, if r_1 is the ratio of tasks completed from $\mathcal{T}_1$ and r_2 is the ratio of tasks completed from $\mathcal{T}_2$, then,

$$(i) \ r_1 \leq r_2,$$

and

$$(ii) \ r_1 > 0 \Rightarrow r_2 = 1.$$

Lemma 4: Let $L = \frac{r_1^{LDF}}{r_2^{LDF}}$ and $E = \frac{r_2^{EDF}}{r_1^{EDF}}$. L lies in the interval $[0, 1]$ and E lies in the interval $[0, \frac{1}{L}]$ for a given $L = z$.

Intuitively, Lemma 3 states that LDF will *always* schedule higher deadline tasks before lower deadline tasks. From this lemma, the value of $\frac{r_1^{LDF}}{r_2^{LDF}}$ lies in the range $[0, 1]$, with the minimum value of 0 occurring when no task in B_1 is completed, and the maximum value of 1 occurring *only when all the tasks* in both bins are completed.

For a given value of r_1^{LDF} and r_2^{LDF} , the minimum value E can take is 0 and the maximum value it can assume is $\frac{1}{L}$ when $r_2^{EDF} = r_2^{LDF}$ and $r_1^{EDF} = r_1^{LDF}$ from Lemma 2. With the idealized assumption that both L and E are uniformly distributed in their respective intervals, we now formally state how often the condition (3) holds in such an environment.

Theorem 3: Let $L = \frac{r_1^{LDF}}{r_2^{LDF}}$ and $E = \frac{r_2^{EDF}}{r_1^{EDF}}$. If we assume the values of L to be uniformly distributed in the interval $[0, 1]$ and the values of E to be uniformly distributed in the interval $[0, \frac{1}{L}]$ for a given $L = z$, then $P(FI_{EDF} \geq FI_{LDF}) = \frac{2}{3}$.

This theorem states that EDF is more likely to achieve a higher fairness than LDF and it follows intuitively from Insight 2 above. Note that we have assumed independent and uniform distributions for L and E for tractability, however, it is likely that the actual distributions will be non-uniform—in particular, from Lemma 3, L will be highly skewed towards 0, which will correspond to a large range of values for E , and will actually increase the likelihood of EDF's fairness being higher than that of LDF. We present empirical results in Section V-F to support this theoretical result by showing that this result holds over a large number of scenarios. In fact, in our empirical results, we were unable to find any instance of LDF's fairness being higher than that of EDF.

IV. LIMITED RESOURCE EARLIEST DEADLINE SCHEDULING

In this section, we present *Limited Resource Earliest Deadline Scheduling (LRED)*: a general deadline-driven scheduler that explicitly incorporates redundant scheduling, and thus provides a flexible way to exploit the throughput-fairness tradeoff in a heterogeneous, dynamic computation system.

In the previous section, we saw that in a heterogeneous, dynamic resource environment, there is a clear tradeoff between throughput and fairness, when coupling redundant scheduling with deadline-driven scheduling. We also saw that EDF and LDF represent the opposite ends of the spectrum in this tradeoff, with one ensuring higher fairness while the other achieves higher throughput. LRED uses these insights to exploit the throughput-fairness tradeoff. We first present the

high-level intuition behind the algorithm, followed by the key concepts used by this algorithm, and then describe the algorithm's working in detail.

- N : Total no. of tasks in the task pool
 L : Total no. of workers available at the scheduler
- S_k : Set of tasks with ascending deadline values each of which can be completed with probability TSR by the workers $\{W_1, \dots, W_k\}$ for $k = 1, 2, \dots, \max$. S_k could be empty.
 S_∞ : Set of tasks that cannot be completed with probability TSR by any number of workers in $\{W_1, \dots, W_L\}$

Fig. 4. Partitioning of the task list by LRED

Fig. 5. Scheduling by LDF, EDF and LRED

A. Intuitive Description of LRED

Intuitively, LRED works by limiting the number of resources consumed per task (thus improving throughput), while scheduling the selected tasks in earliest deadline order (thus improving fairness). To achieve this goal, LRED sorts the task pool in increasing order of deadlines, so that shorter deadline tasks require more resources compared to higher deadline tasks. Then, LRED schedules the tasks in earliest deadline first order *starting* from the first task that needs a (specified) limited number of resources. This limit is specified as a parameter to the algorithm, and allows LRED to control the throughput-fairness tradeoff. For instance, specifying a limit of 3 would result in LRED begin scheduling from the shortest deadline task that needs only 3 nodes, resulting in a higher throughput but lower fairness than specifying a limit of 7. Next we describe the key concepts used by the algorithm and its working in more detail.

B. Key Concepts

Consider a set of N workers and L tasks in the system. Let us assume that the task list is sorted in increasing order of

task deadlines. Further, let us assume that the worker queue is sorted in decreasing order of the mean timeliness value τ of workers⁴. Then, we can define the following:

Definition 3: k-dependent task: A task is said to be k-dependent if it needs exactly the k most timely workers in the worker queue to complete successfully with a high probability (based on a desired target success rate)⁵.

Definition 4: k-dependent task set (S_k): The set of all k-dependent tasks in the task queue.

Figure 4 illustrates the concept of k-dependent tasks and sets. As seen from the figure, tasks T_m to T_p require the top n workers from the worker queue for their successful execution, and thus are n -dependent tasks, and belong to the set S_n . Similarly, tasks T_f to T_i belong to the set S_{max} , while tasks T_x to T_N belong to the set S_1 . Note that the set S_∞ represents the set of tasks that cannot be successfully completed with any number of workers from $\{W_1, \dots, W_L\}$, and are thus infeasible with the current worker pool.

In this way, the task list is divided into disjoint k-dependent sets for $k = 1, 2, \dots, \max$, where $\max$ is the maximum number of workers required by any task in the task list. Note that the size of one of these sets S_k could be zero, which means that there may be no tasks that can be completed with exactly k most timely workers in the worker queue. Also, it can be shown that each set S_k consists of consecutive tasks from the sorted task list, and all tasks in a set S_n have lower deadlines than the tasks in a set S_m , for $n > m$. These properties are based on Lemma 1 given in the Appendix.

C. Algorithm Description

We now describe how the LRED algorithm works in practice. We begin by describing how EDF and LDF would schedule tasks based on the concepts of k-dependent tasks and k-dependent sets presented above.

EDF schedules tasks in the increasing order of their deadlines, so that it will start with the tasks in S_∞ as shown in Figure 5. However, since none of the tasks in S_∞ can be successfully completed by the available workers, EDF will skip these tasks. Thus, EDF will effectively start scheduling from the tasks in S_{max} . Since it tends to consume larger number of nodes per task, it is likely to achieve lower throughput due to Insight 1, but because of its ordering of tasks, it achieves higher fairness due to Insight 2.

LDF, on the other hand, schedules the tasks in the decreasing order of their deadlines, so it will begin by scheduling the longest deadline task in S_1 , as shown in Figure 5. LDF will schedule tasks in S_2 only if S_1 is empty, and continue moving towards S_∞ as each successive set becomes empty. It, however, will not schedule any task in S_∞ since none of these tasks can be executed with the available workers. As we illustrated in Examples 1 and 2 in Section III, this ordering is

⁴Note that based on the shapes of the timeliness distributions, a worker with a higher mean timeliness value need not always have a higher timeliness for a given deadline value D , compared to a worker with a smaller mean timeliness value. We use this ordering mainly as a heuristic.

⁵From here on, we will assume successful completion to be dependent on a given TSR, and omit its mention unless required specifically.

likely to result in higher throughput due to Insight 1, but in lower fairness due to Insight 2.

LRED generalizes the scheduling orders of EDF and LDF by introducing a new set pointer $LRED_{start}$, which refers to the first k -dependent set S_k that the algorithm will schedule tasks from. By using S_k as the starting set pointer, we limit the number of workers to schedule a task to k , hence, we call it a *Limited Resource* algorithm. Once the tasks in the initial set S_k are exhausted, the algorithm moves on to the next k -dependent set with a smaller value of k . The values of the pointer $LRED_{start}$ can be seen to be S_1 and S_∞ respectively for LDF and EDF, as shown in Figure 5. The fairness of an LRED schedule can be increased from that of LDF if $LRED_{start}$ is set to S_2 instead of S_1 , but throughput will reduce accordingly. Similarly, the throughput of LRED can be increased but fairness reduced from that of EDF by setting $LRED_{start}$ to S_{max-1} instead of S_{max} (or S_∞). In this manner, the fairness and the throughput of the system can be adjusted by sliding the pointer $LRED_{start}$ along the task queue.

Besides the choice of the starting k -dependent set S_k , the other question is the ordering of the tasks within S_k . Using Insights 1 and 2, we can achieve higher fairness without sacrificing throughput by traversing S_k in the increasing order of deadlines. LRED also uses this ordering, referring to the *Earliest Deadline* part of the name. Note that because of this ordering, setting $LRED_{start}$ to S_1 will differ from a pure LDF algorithm (achieving higher fairness and throughput than LDF, as we will show in Section V).

To summarize, LRED initializes the task set pointer $LRED_{start}$ to start from a set S_n to obtain a particular level of throughput and fairness. The shortest deadline task from this set S_n is chosen and scheduled to the first group of n most timely workers from the worker list. S_n is traversed in the increasing order of task deadlines, and once S_n becomes empty, a task from the next non-empty set in the list $\{S_{n-1}, S_{n-2}, \dots, S_1\}$ is chosen for scheduling. To make the algorithm work-conserving, once all tasks in the sets $\{S_n, S_{n-1}, \dots, S_1\}$ are exhausted, the algorithm moves to tasks in S_{n+1} , and continues moving towards S_∞ as each successive set becomes empty.

Algorithm 1 LRED(n)

```

1:  $W \leftarrow$  Set of all available workers
2: Sort  $W$  in decreasing order of  $\tau$ 
3: Sort the task pool in increasing order of  $D$ 
4: while  $W$  is non-empty do
5:   Organize the task pool into the list  $\{S_1, S_2, \dots, S_{max}\}$  based
     on  $\tau$  of workers in  $W$ 
6:    $V \leftarrow$  Set of all tasks in the list  $\{S_n, \dots, S_2, S_1\}$ 
7:   if  $V$  is non-empty then
8:      $T \leftarrow$  First task from the first non-empty set  $S_k$  in  $V$ 
9:     Schedule  $T$  to  $k$  most timely workers
10:    Update  $W$  by removing the  $k$  assigned workers
11:   else if  $n < max$  then
12:     LRED( $n+1$ )
13:   else
14:     break
15:   end if
16: end while

```

Algorithm 1 shows the pseudocode for LRED. It takes a parameter n which corresponds to the set S_n to be used as the set pointer $LRED_{start}$. The basic algorithm works by scheduling the group of the n most timely workers among the available worker list to the shortest deadline task T in S_n . The value of n signifies which task among all the tasks in the task pool will be chosen to be scheduled first. When $n = 1$, the execution of LRED(1) corresponds to LDF (except for the ordering of tasks within S_1). When $n = max$ (or ∞), it schedules tasks from S_{max} until either S_{max} becomes empty or all the capable workers are exhausted before moving on to S_{max-1} . This corresponds to an execution of EDF. Also, to make the algorithm work-conserving, once it exhausts all tasks in the sets S_k , for $k = 1, \dots, n$, it recursively calls LRED($n+1$).

	Task pool	Deadline	LRED(2)	LRED(1)
S_2	T_1	D_1	$T_1 \rightarrow \{W_1, W_2\}$	
	T_2	D_1		
	T_3	D_1		
S_1	T_4	D_2	$T_4 \rightarrow \{W_3\}$	$T_4 \rightarrow \{W_1\}$
	T_5	D_2		$T_5 \rightarrow \{W_2\}$
	T_6	D_2		$T_6 \rightarrow \{W_3\}$

Fig. 6. An example schedule by LRED(n) for values of $n = 1, 2$. $L = 6$, $N = 3$

Figure 6 gives an illustration of the schedule created by the LRED(n) algorithm for values of $n = 1, 2$. The figure shows that a higher value of $n=2$ produces a lower throughput but completes more short deadline tasks, whereas with a smaller value of $n=1$, short deadline tasks starve while increasing the net throughput. We next provide detailed quantitative evaluation of this algorithm using a simulation study.

V. EVALUATION

We now evaluate the performance of the LRED algorithm through simulation. We begin by describing our simulation methodology, followed by a definition of the metrics used for evaluation and the results obtained from the simulations.

A. Simulation Methodology

The simulator consists of a central task scheduler and a set of workers that arrive at the scheduler requesting work. Each worker is associated with an underlying response time distribution, which is sampled to generate a response time for each task it is assigned. The scheduler maintains a database of timeliness behavior information for all the individual workers by observing the past history of response times of each worker. The timeliness distribution of each worker is then represented by a histogram of its observed response times for the tasks assigned to it in the past. A sample worker histogram is shown in Figure 7(a). This is used to calculate the timeliness of the worker with respect to a given task deadline D , while the group timeliness of a group of workers is computed as per Equation 2 presented in Section II-D. The length of the maintained history is bounded to accommodate recent changes in the worker's timeliness behavior into the estimated

timeliness. The histogram of each worker is updated by the scheduler whenever a worker returns a result for an assigned task. All the workers are initially assumed to be available at the scheduler, after which the workers return to the scheduler at time intervals based on task response times sampled from their underlying response time distributions.

To emulate the timeliness behavior of individual nodes, we have used a synthetic trace as well as a PlanetLab trace. The synthetic trace consists of the worker response time distributions emulated by normal distributions, with equal standard deviation and with the distribution means uniformly distributed over a specific deadline range. Figure 7(b) shows the distribution of mean worker timeliness values ranging from 80-400 time units. We have also used a trace from PlanetLab that contains the response times of 90 PlanetLab nodes running for 5 hours executing same size tasks, collected during a live execution of BLAST over BOINC [3]. BLAST is a bio-informatics application that performs genetic matching of an input DNA sequence against a gene sequence database. The PlanetLab trace and results based on this trace are presented in detail in Section V-H.

In our simulations, fixed-size tasks are generated with deadlines uniformly distributed in a given deadline range. We vary the task deadline range to emulate different overall timeliness levels of the computational environment for a given worker distribution. For instance, we use a task deadline range towards the lower end of the response time range of the workers, to represent a low timeliness environment (LowTE). Similarly a high (HighTE) or moderately timely environment (ModTE) is simulated by moving the task deadline range over the worker timeliness range accordingly. The inter-arrival times for new tasks arriving into the system follow an exponential distribution. The mean of this inter-arrival time distribution is varied to vary the load level in the system, so that a low mean value corresponds to a high volume of tasks arriving in the system.

For all our experiments, we use two additional parameters: *MaxWkrs* and *Waiting time(w)*. *MaxWkrs* is defined as the maximum number of workers in any group G a task can be replicated to run on. This parameter is used to avoid over-consumption of resources for executing a single task. *Waiting time* is defined as the maximum amount of time the scheduler waits between its scheduling decisions. If w is 0, it means that a scheduling decision is made every time a worker arrives at the scheduler requesting work. By setting w greater than 0, it is possible to have a larger number of workers available to the scheduler to enable better scheduling decisions. We discuss the significance and the impact of the waiting times in more detail later.

In all of our experiments, we have set the parameters *MaxWkrs* = 7 (i.e., the maximum group size that can be used to schedule a task) and *TSR* = 0.9. We executed the experiments for LRED(n) for values of n = 1, 4, 7 for each timeliness environment under different load conditions. We execute the EDF and LDF algorithm under the same conditions for comparison. In addition, we also simulated a Random algorithm (*Rand*), that selects tasks randomly in arbitrary order from the task pool, while keeping the worker list sorted.

B. Metrics

- *Throughput* : The total number of tasks that are completed within their deadlines.
- *Fairness* : We measure the fairness of a schedule using Jain's fairness index [15] as follows. Suppose the deadline range of the tasks is divided into m bins s.t. C_i is the number of tasks in bin i and X_i is the number of tasks successfully completed in bin i , then the fairness index FI is given by:

$$FI = \frac{\left[\sum_{i=1}^m x_i \right]^2}{m \sum_{i=1}^m x_i^2} \quad \text{where } x_i = \frac{X_i}{C_i} \quad (4)$$

The value of the fairness index ranges from 0 to 1, with a value of 1 representing absolute fairness in the system while 0 indicates absolute unfairness.

C. Throughput-Fairness Tradeoff

We start by presenting results for the synthetic worker timeliness distribution trace. Figures 8(a) and 8(b) show the tradeoff between fairness and throughput for a low timeliness environment (LowTE). In this environment, the task deadlines lie in the range 80-150, which makes only about 25% of the workers timely for these tasks. The load is kept high with a mean task arrival time of 5. The waiting time w is fixed at 2, which means that the scheduler will allocate tasks to workers every 2 time units. This waiting time gives a sufficient number of workers at each scheduling point.

Figures 8(a) and 8(b) plot the fairness index FI and throughput respectively for the different scheduling algorithms. As expected, the fairness of LRED increases as n increases, while throughput decreases as shown in the figures. EDF shows the highest fairness and lowest throughput. LDF has the lowest fairness, however, its throughput is lower than that of LRED(1), which demonstrates the benefit of scheduling tasks in the increasing order of deadlines within a k -dependent set (S_1 in this case). *Rand* shows slightly higher fairness and lower throughput than LDF because it happens to schedule a greater number of lower deadline tasks than LDF due to the randomness in choosing tasks. EDF does not show any dramatic improvement over LRED(7), because the majority of the tasks that could be finished with unlimited number of workers needed only a size of 7 at maximum.

To understand these results better, Figure 9(a) shows the ratio of tasks completed in each deadline bin by the different algorithms. The fairness level of an algorithm is indicated by how flat its curve is. As seen from the figure, LRED(7) could finish more low deadline tasks than that of LRED(1). In addition, if there are no timely workers to schedule the shortest deadline tasks, it finishes some tasks from higher deadline bins requiring smaller group sizes than 7. This gives LRED(7) a flatter curve and consequently higher fairness.

This result can be verified by Figure 9(b) which shows the ratio of completed tasks grouped by the k -dependent sets S_k to which they belonged at scheduling time. As seen

Fig. 7. Representing a single worker and all workers in the environment

Fig. 8. Comparing FI and Throughput of LDF and Rand with LRED for $n=1,4,7,\infty$ in a LowTE

Fig. 9. (a) LRED performing in each deadline bin. LRED(1) completes more tasks in higher bins while LRED(7) performs better in lower bins (b) LRED performing in each k-dependent set. LRED(1) spends more on S_1 while LRED(7) is more evenly spread across majority of the sets

from the figure, LRED(1) shows a heavy bias towards tasks from S_1 (high deadline tasks), while the tasks scheduled by LRED(7) are more equally distributed across the various sets. For instance, if we consider the sets S_1 and S_3 , LRED(1) finishes 67% of the tasks in S_1 as opposed to only 4% of the tasks in S_3 . On the other hand, LRED(7) finishes 30% of the tasks in S_1 and 18% in S_3 , showing a smaller deviation than LRED(1). This explains the high fairness for LRED(7) compared to LRED(1). Yet, the throughput is lower than LRED(1) since it effectively finishes less number of tasks from all the sets. Also, we see that none of the algorithms could do many tasks from $\{S_5, S_6, S_7\}$ because of the lack of timely workers in the LowTE.

D. Impact of Load

We next explore the effect of system load on the performance of LRED for different n values for LowTE. We varied the task arrival times and observed the resulting fairness and throughput for each of $n=1,4,7$, and for Rand in Figures 10(a) and 10(b). A high mean task arrival time represents a light load condition while a low mean value creates a heavy load condition.

LRED(1) has the lowest fairness and highest throughput when the system load is high, because tasks arrive at a faster rate and keep LRED(1) busy scheduling along the task list with lower k-dependent sets. So, most tasks use smaller group sizes and consequently the throughput is higher. However, LRED(1) has fewer chances to move to higher k-dependent sets due

Fig. 10. Impact of load on fairness and throughput of LRED in a LowTE

TE	n
LowTE	5.2
ModTE	3.1
HighTE	1.1

TABLE I
MAXIMUM n NEEDED FOR THE SHORTEST DEADLINE TASK SCHEDULED

to the higher task arrival rate and consequently exhibits less fairness compared to LRED(7) and LRED(4). On the other hand, as the mean task arrival rate increases, reducing load on the system, the fairness and throughput of LRED for all values of n converges, because now the number of waiting tasks is small enough that most tasks can be picked up by the workers irrespective of their deadlines. We also see that the total throughput decreases with decreasing load due to fewer tasks arriving in the system.

E. Impact of Timeliness Environment

We plot the impact of the underlying timeliness environment on the performance of LRED, Rand, LDF and EDF in Figures 11(a) and 11(b). The timeliness of the environment is increased by moving the task deadline ranges towards the higher range of the worker timeliness in the distribution Figure 7(b). LowTE, ModTE and HighTE correspond to the task deadline ranges of 80-150 (<25% workers are timely), 120-200 (approximately 40% workers are timely) and 200-350 (>75% are timely) respectively.

We make two main observations from these figures regarding the behavior of LRED. First, the figures show that as the overall timeliness level of the environment increases, the fairness as well as the throughput values converge for all values of n . Secondly, we also see that as the timeliness level gets higher, both fairness and throughput increase for the same value of n .

These results occur because as the timeliness level increases, smaller groups of workers are required to successfully complete most of the tasks. Evidence for this is seen from Table I which shows the maximum group size required for the shortest deadline task for each timeliness environment. As seen from the table, this value decreases from 5.2 to 1.1 as we go from LowTE to HighTE. This implies that while most tasks are spread between S_1 to S_6 for LowTE, most of the

tasks are concentrated in S_1 only for HighTE. As a result, for HighTE, there is less differentiation between the different values of n for LRED(n) resulting in similar throughput and fairness. Moreover, since most tasks require only 1 worker for successful execution in HighTE, the overall throughput and fairness also increase.

Rand and LDF, however, show poor fairness as well as throughput as compared to LRED even in HighTE, due to the poor choices they make for each task. They both might assign highly timely workers to very high deadline tasks leaving the short deadline tasks to be assigned to larger sized groups, reducing the overall throughput. It affects fairness due to the starvation of the short deadline tasks for the same reason.

Fig. 12. Equation (3) with varying timeliness level

F. Empirical Evidence for Theoretical Fairness Property

From Figure 11(a), we observe that the fairness of EDF is greater than that of LDF across all timeliness environments. Recall from Theorem 2 in Section III-B that this happens only if the condition in Equation (3) holds, which can be rewritten as:

$$\frac{N_2^{EDF}/N_1^{EDF}}{N_1^{LDF}/N_2^{LDF}} \geq 1,$$

assuming equal number of tasks in both bins. We refer to the LHS of the equation above as $Ratio(EDF, LDF)$. In Section III-B, we had given a probabilistic argument that this condition is highly likely to hold in general; we now present empirical results to back this claim. In particular, we show that

Fig. 11. Fairness-Throughput of LRED for different Timeliness environments with a fixed load

Equation (3) holds across a wide variation in the timeliness environment.

For these results, in addition to the LowTE, ModTE, and HighTE environments described above, we also simulated a number of intermediate timeliness environments by varying the task deadlines. For each timeliness environment, we divided the task pool into two equal sets $\mathcal{T}_1$ and $\mathcal{T}_2$ of equal size with deadlines of tasks in $\mathcal{T}_1$ less than the deadlines of tasks in $\mathcal{T}_2$. We then executed LDF and EDF and measured the throughput values of $\mathcal{T}_1$ and $\mathcal{T}_2$ as N_1 and N_2 respectively for each algorithm. Figure 12 plots the value of $Ratio(EDF, LDF)$ against increasing values of $\frac{N_1^{LDF}}{N_2^{LDF}}$. Note that each value on the x-axis corresponds to a different timeliness environment—in particular, the x-axis value increases with increasing timeliness levels, because LDF is able to finish more tasks in $\mathcal{T}_1$ for a better timeliness environment. From the figure, we can observe that the value of $Ratio(EDF, LDF)$ remains greater than 1 for all timeliness levels, including HighTE, where most tasks require only 1 worker (Table I). This result shows that even in extreme timeliness scenarios, the condition (Equation (3)) holds true, and hence, the fairness of EDF is highly likely to be higher than that of LDF in general.

G. Significance of waiting time

Since the scheduler needs to wait for workers to arrive before making scheduling decisions, we conducted experiments to see how long we can wait without affecting throughput. We used a high load setting in a LowTE and HighTE, and varied the waiting times. The results are shown in Figure 13.

Figure 13(a) shows the impact of waiting time on the throughput in a LowTE. As seen from the figure, the throughput initially increases as we increase the waiting time, but then it starts decreasing after a point. This is because, with no waiting, we may not have enough timely workers to form good groups, since most of the workers in this environment have poor timeliness. This is clearly shown by the low throughput at $w=0$. If we wait for some amount of time (in this case until $w=2$), we get a few more timely workers to form better groups so that more tasks can be finished. However, if we wait too long, the deadlines of the pooled tasks drop too much to be successfully completed in time.

Figure 13(b) shows that the drop in throughput happens at

a much lower waiting time value (with only a slight increase from $w=0$ to $w=1$). This is because, with HighTE, since the workers already have good timeliness without waiting, we can achieve a good throughput. Increasing the waiting time only results in reducing the deadlines of waiting tasks without much benefit in getting better workers.

Overall, waiting for some time may enable the arrival of better workers to result in better scheduling decisions, however, waiting too long could expire task deadlines. Therefore, these two effects have to be balanced based on the environment.

H. Results with PlanetLab Trace

In addition to the synthetic worker distributions, we conducted some experiments with a worker trace from PlanetLab. The trace contains the response times of 90 PlanetLab nodes running equal-sized tasks, recorded over a duration of 5 hours to capture sufficient timeliness data about all the workers. The worker distribution with their mean response times is shown in Figure 14(a). We used these traces to initialize the workers and ran experiments with LRED and LDF for LowTE, ModTE and HighTE with the task deadline ranges of 60-160 (approximately 20% of workers are highly timely), 130-230 (approximately 40% of workers are timely) and 200-280 (>70% are timely) respectively. The mean task arrival time was fixed at 5 representing high load and the waiting time w set to 2. The results are summarized in Figures 14(b) and 14(c). The observations are similar to that for synthetic worker distributions. LRED(1) leads in fairness while LRED(7) leads in throughput for LowTE and ModTE. However, all LRED(n) show the same performance for HighTE while LDF still lags behind.

I. Summary of Results

We summarize the major results below.

- Reducing the value of n with LRED provides better throughput because it uses fewer workers per task.
- Using LRED with higher values of n produces higher fairness, as the workers are now more spread out across different deadline tasks.
- When the load on the system is high, the throughput-fairness tradeoff is more visible than at low loads.

Fig. 13. Impact of waiting time on the throughput of LRED and Rand

Fig. 14. LRED for different TEs for PlanetLab traces

- The throughput-fairness tradeoff also becomes more prominent as the overall timeliness level of the environment decreases.

Some of the key conclusions we can draw from these results are as follows. When a system is heavily loaded in terms of task arrivals or the timeliness level is low (so that there are more stringent tasks or less timely workers in the system), there is a greater opportunity to exploit the tradeoff between throughput and fairness to suit the system requirements. LRED provides a way to tune this tradeoff by simply increasing or decreasing its group size parameter. When the system load reduces or the timeliness level increases, however, simple EDF can be used because a good throughput-fairness combination can be obtained irrespective of any limit on the group size. Note that the ability of LRED to achieve flexibility under high loads and low timeliness makes it powerful as these system conditions are more likely to impact an application's performance.

VI. RELATED WORK

Deadline-based scheduling: There is a large body of work on deadline-based scheduling [16]. Earliest-deadline-first (EDF) [11] is a commonly used deadline-based online scheduling algorithm that has been shown to be optimal for uniprocessor systems as long as the demand does not exceed the system capacity. Many ideas have been proposed to ensure the optimality of deadline scheduling under overloads [17], [18], while variations of EDF such as least slack first [19] have also been proposed over the years. In this paper, we have

looked at the problem of combining deadline-based scheduling with redundant scheduling in a heterogeneous environment, and our main focus is on providing probabilistic guarantees for soft deadline tasks.

Fairness in scheduling: Fairness has been emphasized by many scheduling algorithms in allocating processor bandwidth to applications in both uniprocessor as well as multiprocessor systems. Many systems use proportional share allocation based on Generalized Processing Sharing [20], which is an ideal algorithm in which each application is allocated bandwidth based on the weight assigned to the application. Many other algorithms [21], [22], [23], [24], [25] have been proposed that approximate GPS in different domains. Fairness has also been widely investigated in queuing systems [26], [27], [28], [29]. Adam Wierman et al. in [30], [31] provide a classification of scheduling policies with respect to fairness. Our focus is on exploring the tradeoff between fairness and throughput in the context of deadline-driven scheduling. Also, we focus on a heterogeneous system with dynamic resource availability.

In the context of scheduling tasks with constraints (e.g. deadlines), [32] introduced the concept of Pfairness. Under Pfair scheduling, tasks are scheduled according to a fixed-size allocation quantum so that deviation from an ideal allocation is strictly bounded. We focus on a heterogeneous multi-resource environment with probabilistic requirements.

Deadline scheduling in large-scale systems: One of the earlier works that proposed an on-line deadline scheduling algorithm for client-server Grid systems is [33]. Caron et al. [34] improved this work by associating a priority with each task. In [35], the authors propose a fair scheduling algorithm

for wireless networks that ensures packets are dropped fairly among users in case of missed deadlines. In this paper, our mechanism provides a way to control the fairness among the task deadlines.

Redundant scheduling: Redundancy has been used in several contexts such as Byzantine fault tolerant systems [36], [37]. Redundancy has also been widely employed by data storage systems to ensure high availability, performance and fault tolerance [38], [39]. Here, we use redundancy for achieving robust scheduling guarantees.

Reputation-based scheduling: There are a number of papers that have proposed the use of reputation to store the trust and reliability values of nodes for use in scheduling [40], [41], [10]. However, many of these systems do not consider task deadlines which creates a second dimension in the reputation computation, as shown by our use of node timeliness distribution, rather than each node having a fixed reputation across all tasks.

Scheduling in heterogeneous environments: Maheswaran et al. [42] study heuristics similar to EDF for scheduling independent tasks in heterogeneous computing systems, with the objective of maximizing throughput. Our work is on scheduling independent tasks with deadline constraints, with additional requirements about fairness.

VII. CONCLUSION

In this paper, we examined the problem of deadline-driven task scheduling in a heterogeneous and dynamic computational environment. We showed that combining deadline-based scheduling with redundant scheduling typically used in such systems leads to a fundamental tradeoff between throughput and fairness. In particular, we showed that earliest deadline first (EDF) results in lower throughput because the tasks with stringent deadlines consume more resources, while at the same time ensuring a higher fairness. However, while latest deadline first (LDF) provides better resource utilization, it is likely to starve stringent deadline tasks.

To exploit this tradeoff in such heterogeneous and dynamic environments, we proposed a new scheduling algorithm called Limited Resource Earliest Deadline (LRED) that couples redundant scheduling with deadline-driven scheduling in a flexible way by using a simple tunable parameter. This algorithm uses a statistical notion of timeliness for each computational node that captures both inter-node heterogeneity as well as intra-node dynamism in the system, and can be estimated based on the node's past execution history. Further, we showed that LRED is a generalization of EDF and LDF, so that, by tuning its parameter, LRED reduces to EDF in one extreme, and to (a close variant of) LDF in another extreme.

We used trace-driven and synthetic simulations to evaluate the performance of LRED. Our results show that load and the timeliness level of the underlying environment have a significant impact on the throughput-fairness tradeoff of task scheduling. We find that LRED provides a powerful mechanism to achieve desired throughput or fairness under high loads and low timeliness environments, where these tradeoffs are most critical.

In the future, we intend to implement this algorithm on a live system such as PlanetLab to observe the tradeoff when the nodes change in behavior abruptly during the course of the experiments. In addition, we plan to extend our algorithm to incorporate heterogeneous tasks as well.

APPENDIX A

Note: For all the theoretical results below to hold, we make the following assumptions.

- 1) We assume the system model described in Section II
- 2) We assume a combination of redundant and deadline-based scheduling, so that a schedule is required to meet each task deadline with a given target success rate (TSR) defined in Section II.
- 3) All the worker nodes in the environment are available at each scheduling instant.
- 4) The timeliness levels of the worker nodes do not change for the entire duration of scheduling.
- 5) The set of tasks to be scheduled is fixed.

Lemma 1: Let T_i and T_j be two tasks such that $D_i \leq D_j$ and $T_i \in S_n, T_j \in S_m$. If T_i can be assigned to worker group G_i , then $\exists G_j \subseteq G_i$ to which T_j can be assigned.

Proof: T_i can be assigned to $G_i \Rightarrow \tau_{G_i}(D_i) \geq \text{TSR}$. $T_i \in S_n \Rightarrow G_i = \{W_1, \dots, W_n\}$ where $\{W_1, \dots, W_n\}$ are the first n workers in the sorted worker queue.

Now, using Equation 2 in Section II-D, the timeliness of G_i for T_j is given by:

$$\tau_{G_i}(D_j) = 1 - \prod_{l=1}^n (1 - \tau_l(D_j)).$$

From Equation 2,

$$D_j \geq D_i \Rightarrow \tau_l(D_j) \geq \tau_l(D_i) \text{ for each worker } W_l, l = 1, 2, \dots, n$$

$$\Rightarrow (1 - \tau_l(D_j)) \leq (1 - \tau_l(D_i))$$

$$\Rightarrow \tau_{G_i}(D_j) \geq \tau_{G_i}(D_i)$$

This means that T_j can be assigned to the workers in the group G_i . Hence, $G_j \subseteq G_i$. ■

Corollary 1: If T_i and T_j are two tasks such that $D_i \leq D_j$, and if T_i and T_j respectively need at most n and m workers to satisfy the TSR, then $m \leq n$.

Corollary 2: Let $\mathcal{T} = \{T_1, \dots, T_i, \dots, T_k\}$ be the set of k tasks in a schedule that contains task T_i , and let $\mathcal{G} = \{G_1, \dots, G_i, \dots, G_k\}$ be the corresponding set of worker groups assigned to the tasks. Let $\mathcal{T}' = (\mathcal{T} - \{T_i\}) \cup \{T_j\}$, where, $T_j \notin \mathcal{T}$ is a task such that $D_i \leq D_j$, and $\mathcal{G}'$ be the corresponding set of worker groups assigned to the tasks in $\mathcal{T}'$. Then,

$$|\bigcup_{G_l \in \mathcal{G}'} G_l| \leq |\bigcup_{G_l \in \mathcal{G}} G_l|.$$

Theorem 1: At a given scheduling instant, if $\mathcal{S}^{\text{EDF}}$ is a task schedule generated by EDF scheduler and $\mathcal{S}^{\text{LDF}}$ is a task schedule generated by LDF, then $|\mathcal{S}^{\text{EDF}}| \leq |\mathcal{S}^{\text{LDF}}|$, where, $|\mathcal{S}|$ denotes the number of tasks scheduled in $\mathcal{S}$.

Proof: Let $\mathcal{S}^{\text{EDF}} = \{T_1, T_2, \dots, T_k\}$ be the set of tasks in the schedule generated by a EDF scheduler with $D_1 \leq D_2 \leq \dots \leq D_k$ and let $\mathcal{G} = \{G_1, G_2, \dots, G_k\}$ be the corresponding worker groups assigned.

Let $\mathcal{T} = \{T_N, T_{N-1}, \dots, T_{k+1}\}$ be the non-empty task pool with unscheduled tasks ordered such that $D_{k+1} \leq D_{k+2} \leq \dots \leq D_N$.

If $D_1 < D_N$, then replacing T_1 with T_N implies that $|\{G_N, G_2, \dots, G_k\}| \leq |\{G_1, G_2, \dots, G_k\}|$, from Corollary 2. Let us continuously replace each task $T_i \in \mathcal{S}^{\text{EDF}}$ with the tasks $T_{N-i+1} \in \mathcal{T}$, $i = 1, 2, \dots, k$, until $D_i \geq D_{N-i+1}$. If we call the new schedule $\mathcal{S}'$ and the new set of corresponding worker

groups $\mathcal{G}'$, then we can say that $|\mathcal{G}_r| = (|\mathcal{G}| - |\mathcal{G}'|) \geq 0$. Since the last task in $\mathcal{S}'$ has deadline greater than or equal to the first task from $\mathcal{T}$, $\mathcal{S}'$ is exactly the schedule $\mathcal{S}^{\text{LDF}}$ generated by LDF. To this schedule $\mathcal{S}'$, we can add tasks in LDF order from $\mathcal{T}$ until $|\mathcal{G}_r| \leq 0$. Hence, it follows that $|\mathcal{S}^{\text{LDF}}| \geq |\mathcal{S}^{\text{EDF}}|$. ■

Lemma 2: For two non-empty sets of tasks $\mathcal{T}_1$ and $\mathcal{T}_2$ such that the deadline of any task in $\mathcal{T}_1$ is less than or equal to the deadline of any task in $\mathcal{T}_2$, if LDF completes a ratio of tasks r_1^{LDF} in $\mathcal{T}_1$ and r_2^{LDF} in $\mathcal{T}_2$, and if EDF completes a ratio of tasks r_1^{EDF} in $\mathcal{T}_1$ and r_2^{EDF} in $\mathcal{T}_2$ then,

$$r_1^{\text{EDF}} \geq r_1^{\text{LDF}}$$

and

$$r_2^{\text{EDF}} \leq r_2^{\text{LDF}}.$$

Proof: Since LDF could complete a ratio r_1^{LDF} from $\mathcal{T}_1$, EDF will be able to complete atleast a ratio r_1^{LDF} from $\mathcal{T}_1$ due to the scheduling order it follows, which implies that $r_1^{\text{EDF}} \geq r_1^{\text{LDF}}$. If $r_2^{\text{EDF}} > r_2^{\text{LDF}}$, then $(r_1^{\text{EDF}} \cdot |\mathcal{T}_1| + r_2^{\text{EDF}} \cdot |\mathcal{T}_2|) > (r_1^{\text{LDF}} \cdot |\mathcal{T}_1| + r_2^{\text{LDF}} \cdot |\mathcal{T}_2|)$, which contradicts Theorem 1. Hence, $r_2^{\text{EDF}} \leq r_2^{\text{LDF}}$. ■

Theorem 2: If the tasks are divided into two non-empty sets $\mathcal{T}_1$ and $\mathcal{T}_2$, s.t. any task in $\mathcal{T}_1$ has a lower or equal deadline to any task in $\mathcal{T}_2$, and if $r_1^{\text{EDF}}, r_2^{\text{EDF}}$ and $r_1^{\text{LDF}}, r_2^{\text{LDF}}$ respectively are the ratio of tasks completed in those bins by EDF and LDF, where $r_1^{\text{EDF}}, r_2^{\text{LDF}} > 0$, then, $FI_{\text{EDF}} \geq FI_{\text{LDF}}$, if and only if

$$\frac{r_2^{\text{EDF}}}{r_1^{\text{EDF}}} \geq \frac{r_1^{\text{LDF}}}{r_2^{\text{LDF}}}. \quad (5)$$

Proof:

Using Lemma 2, we can assume $r_1^{\text{LDF}} = r_1, r_2^{\text{LDF}} = r_2$ and $r_1^{\text{EDF}} = r_1 + x, r_2^{\text{EDF}} = r_2 - y$ and $x, y \geq 0$.

From (4),

$$\begin{aligned} FI_{\text{LDF}} &= \frac{(r_1 + r_2)^2}{2(r_1^2 + r_2^2)} \\ FI_{\text{EDF}} &= \frac{[(r_1 + x) + (r_2 - y)]^2}{2[(r_1 + x)^2 + (r_2 - y)^2]} \end{aligned}$$

We need to prove,

$$FI_{\text{EDF}} \geq FI_{\text{LDF}}$$

$$\Leftrightarrow \frac{[(r_1 + x) + (r_2 - y)]^2}{2[(r_1 + x)^2 + (r_2 - y)^2]} \geq \frac{(r_1 + r_2)^2}{2(r_1^2 + r_2^2)} \quad (6)$$

Simplifying,

$$\begin{aligned} \Leftrightarrow \frac{[r_1 + r_2]^2}{[r_1 + x + r_2 - y]^2} &\geq \frac{r_1 r_2}{(r_1 + x)(r_2 - y)} \\ \Leftrightarrow \frac{[r_1 + r_2]^2}{r_1 r_2} &\geq \frac{[r_1 + x + r_2 - y]^2}{(r_1 + x)(r_2 - y)} \\ \Leftrightarrow \frac{r_1^2 + r_2^2 + 2r_1 r_2}{r_1 r_2} &\geq \frac{(r_1 + x)^2 + (r_2 - y)^2 + 2(r_1 + x)(r_2 - y)}{(r_1 + x)(r_2 - y)} \end{aligned}$$

$$\begin{aligned}
\Leftrightarrow \quad & \frac{r_1}{r_2} + \frac{r_2}{r_1} \geq \frac{r_1+x}{r_2-y} + \frac{r_2-y}{r_1+x} \\
\Leftrightarrow \quad & \frac{r_2}{r_1} - \frac{r_2-y}{r_1+x} \geq \frac{r_1+x}{r_2-y} - \frac{r_1}{r_2} \\
\Leftrightarrow \quad & \frac{1}{r_1(r_1+x)} \geq \frac{1}{(r_2-y)r_2} \\
\Leftrightarrow \quad & \frac{r_2-y}{r_1+x} \geq \frac{r_1}{r_2}
\end{aligned}$$

Hence, it is proved that $FI_{EDF} \geq FI_{LDF}$, if and only if the condition (5) is true. ■

Lemma 3: In LDF scheduling, given two non-empty sets of unscheduled tasks $\mathcal{T}_1$ and $\mathcal{T}_2$ such that the deadline of any task in $\mathcal{T}_1$ is less than or equal to the deadline of any task in $\mathcal{T}_2$, if r_1 is the ratio of tasks completed from $\mathcal{T}_1$ and r_2 is the ratio of tasks completed from $\mathcal{T}_2$, then,

$$(i) \quad r_1 \leq r_2,$$

and

$$(ii) \quad r_1 > 0 \Rightarrow r_2 = 1.$$

Proof: Let us first prove (ii) and then prove (i) using (ii). Suppose $r_1 > 0$ and $r_2 \neq 1$. This means that LDF could complete a task from $\mathcal{T}_1$, when there is at least one task in $\mathcal{T}_2$ that could not be completed. That is, it could find a worker group for a task in $\mathcal{T}_1$ but not for a task in $\mathcal{T}_2$, which contradicts Lemma 1. Hence $r_1 > 0 \Rightarrow r_2 = 1$.

Now let us prove that $r_1 \leq r_2$ using (ii). If $r_1 = 0$, then $r_1 \leq r_2$. If $r_1 > 0$, then $r_2 = 1$ from (ii). Hence, $r_1 \leq r_2$. ■

Lemma 4: Let $L = \frac{r_1^{LDF}}{r_1^{EDF}}$ and $E = \frac{r_2^{EDF}}{r_1^{EDF}}$. L lies in the interval $[0, 1]$ and E lies in the interval $[0, \frac{1}{z}]$ for a given $L = z$.

Proof: From Lemma 3, it follows that L lies in $[0, 1]$. From Lemma 2, it follows that E lies in $[0, \frac{1}{z}]$ for a given $L = z$. ■

Theorem 3: Let $L = \frac{r_1^{LDF}}{r_1^{EDF}}$ and $E = \frac{r_2^{EDF}}{r_1^{EDF}}$. If we assume the values of L to be uniformly distributed in the interval $[0, 1]$ and the values of E to be uniformly distributed in the interval $[0, \frac{1}{z}]$ for a given $L = z$, then $P(FI_{EDF} \geq FI_{LDF}) = \frac{2}{3}$.

Proof: Let us assume $r_1^{LDF} = r_1$, $r_2^{LDF} = r_2$ and $r_1^{EDF} = r_1 + x$, $r_2^{EDF} = r_2 - y$.

From Theorem 2, we know that $FI_{EDF} \geq FI_{LDF}$ only if $L \geq E$. If we prove that $P(L \geq E) = \frac{2}{3}$, then it is automatically proved that $P(FI_{EDF} \geq FI_{LDF}) = \frac{2}{3}$.

The probability mass function of L is given by

$$f_L(z) = \begin{cases} 1 & \text{if } z \in [0, 1] \\ 0 & \text{otherwise} \end{cases}$$

The conditional probability that $E \geq L$ given L is

$$P(E \geq L | L = z) = \int_z^{1/z} \frac{1}{\frac{1}{z}} dx = \int_z^{1/z} z dx = 1 - z^2$$

Now, the probability that $E \geq L$ is given by

$$\begin{aligned}
P(E \geq L) &= \int_0^1 P(E \geq L | L = z) f_L(z) dz \\
&= \int_0^1 (1 - z^2)(1) dz \\
&= \frac{2}{3}.
\end{aligned}$$

REFERENCES

- [1] I. Foster and C. Kesselman, *The Grid 2: Blueprint for a New Computing Infrastructure*. Morgan Kaufman, CA, USA, 2004.
- [2] D. Zhou and V. Lo, "Cluster Computing on the Fly: Resource Discovery in a Cycle Sharing Peer-to-Peer System," *CCGrid*, 2004.
- [3] D. Anderson, "BOINC: A System for Public-Resource Computing and Storage," *IEEE/ACM GRID*, 2004.
- [4] "Folding@home distributing computing project," <http://folding.stanford.edu>.
- [5] "PPDG: Particle Physics Data Grid," <http://www.ppdg.net>.
- [6] "Climate Prediction Network," <http://www.climateprediction.net/>.
- [7] P. Bonetto, M. Guarracino, and F. Inguglia, "Integrating Medical Imaging into a Grid Based Computing Infrastructure," *Computational Science and Its Applications - ICCSA*, vol. 3044, pp. 505–514, Apr 2004.
- [8] C. Germain, V. Breton, P. Clarysse, Y. Gaudreau, T. Glatard, E. Jeannot, Y. Legre, C. Loomis, J. Montagnat, J.-M. Moureaux, A. Osorio, X. Pennec, and R. Texier, "Grid-enabling medical image analysis," *CCGrid*, May 2005.
- [9] E. Bagarinao, L. F. G. Sarmenta, Y. Tanaka, K. Matsuo, and T. Nakai, "Real-Time Functional MRI Analysis Using Grid Computing," *High Performance Computing and Grid*, 2004.
- [10] K. Budati, J. Sonnek, A. Chandra, and J. Weissman, "RIDGE: Combining Reliability and Performance in Open Grid Platforms," *HPDC*, Jun 2007.
- [11] C. L. Liu and J. W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment," *Journal of the Association for Computing Machinery* 20, vol. 1, pp. 46–61, Jan 1973.
- [12] "Cube cluster," http://www.citilabs.com/cube_cluster.html.
- [13] T. Oinn, M. Addis, J. Ferris, D. Marvin, T. C. M. Greenwood, A. Wipat, and P. Li, "Taverna, A tool for the composition and enactment of bioinformatics workflows," *Bioinformatics journal*, 2004.
- [14] J. C. Jacob, R. Williams, J. Babu, S. G. Djorgovski, M. J. Graham, D. S. Katz, A. Mahabal, C. D. Miller, R. Nichol, D. E. V. Berk, and H. Walia, "Grist: Grid Data Mining for Astronomy," *Astronomical Data Analysis Software and Systems XIV*, Oct 2004.
- [15] R. Jain, D. Chiu, and W. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems," *DEC Research Report TR-301*, Sep 1984.
- [16] S.-C. Cheng, J.-A. Stankovic, and K. Ramamritham, "Scheduling Algorithms for Hard Real-Time Systems: A Brief Survey," *Tutorial: Hard Real Time Systems*, 1989.
- [17] S. Baruah, G. Koren, B. Mishra, A. Raghunathan, L. Rosier, and D. Shasha, "Online Scheduling in the Presence of Overload," *Foundations of Computer Science*, 1991.
- [18] T. Lam and K. To, "Performance guarantee for Online Deadline Scheduling in the Presence of Overload," *ACM SODA*, 2001.
- [19] J. Leung, "A New Algorithm for Scheduling Periodic Real-Time Tasks," *Algorithmica*, 1989.
- [20] A. K. Parekh and R. G. Gallager, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks-The Single Node Case," *ACM/IEEE Transactions on Networking*, 1991.
- [21] A. Demers, S. Keshav, and S. Shenkar, "Analysis and Simulation of a Fair Queuing Algorithm," *Journal of Internetworking Research and Experience*, 1990.
- [22] P. Goyal, X. Guo, and H. M. Vin, "A Hierarchical CPU Scheduler for Multimedia Operating Systems," *OSDI*, 1996.
- [23] I. Stoica, H. Abdel-Wahab, and K. Jeffay, "A Proportional Share Resource Allocation Algorithm for Real-Time Time-Shared Systems," *Real Time Systems Symposium*, 1996.
- [24] A. Chandra, M. Adler, P. Goyal, and P. Shenoy, "Surplus Fair Scheduling: A Proportional-Share CPU Scheduling Algorithm for Symmetric Multiprocessors," *USENIX OSDI*, 2000.
- [25] B. Caprita, W. Chan, J. Nieh, C. Stein, and H. Zheng, "Group Ratio Round-Robin: O(1) Proportional Share Scheduling for Uniprocessor and Multiprocessor Systems," *USENIX Annual Technical Conference*, Apr 2005.
- [26] B. Avi-Itzhak and H. Levy, "On measuring fairness in queues," *Advances of Applied Probability*, 2004.
- [27] D. Raz, H. Levy, and B. Avi-Itzhak, "A Resource Allocation Queuing Fairness Measure," *ACM SIGMETRICS*, vol. 15, pp. 600–625, 2004.
- [28] D. Raz, B. Avi-Itzhak, and H. Levy, "Fair Operation of Multi-Server and Multi-Queue Systems," *ACM SIGMETRICS*, 2005.

- [29] N. Bansal and M. Harchol-Balter, "Analysis of SRPT Scheduling: Investigating Unfairness," *ACM SIGMETRICS*, 2001.
- [30] A. Wierman and M. Harchol-Balter, "Classifying Scheduling Policies with respect to Unfairness in an M/G/1," *ACM SIGMETRICS*, 2003.
- [31] A. Wierman, "Fairness and classifications," *ACM SIGMETRICS Performance Evaluation Review*, 2007.
- [32] S. Baruah, N. Cohen, C. Plaxton, and D. Varvel, "Proportionate progress: A notion of fairness in resource allocation," *Algorithmica*, vol. 15, pp. 600–625, 1996.
- [33] A. Takefusa, S. Matsuoka, H. Casanova, and F. Berman, "A Study of Deadline Scheduling for Client-Server Systems on the Computational Grid," *HPDC*, 2001.
- [34] E. Caron, P. K. Chouhan, and F. Desprez, "Deadline Scheduling with Priority for Client-Server Systems on the Grid," *ACM GRID*, Nov 2004.
- [35] A. K. F. Khattab and K. M. F. Elsayed, "Channel-Quality Dependent Earliest Deadline Due Fair Scheduling Schemes for Wireless Multimedia Networks," *ACM MSWiM*, 2004.
- [36] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems*, Jul 1982.
- [37] M. Castro and B. Liskov, "Practical Byzantine Fault Tolerance," *OSDI*, Feb 1999.
- [38] R. Bhagwan, K. Tati, Y.-C. Cheng, S. Savage, and G. Voelker, "Total Recall: System Support for automated availability management," *NSDI*, 2004.
- [39] A. Haeberlen, A. Mislove, and P. Druschel, "Glacier: Highly Durable, Decentralized Storage despite Massive Correlated Failures," *NSDI*, 2005.
- [40] S. Zhao, V. Lo, and C. GauthierDickey, "Result Verification and Trust-based Scheduling in Peer-to-Peer Grids," *P2P*, 2005.
- [41] J. Sonnek, M. Nathan, A. Chandra, and J. Weissman, "Reputation-Based Scheduling on Unreliable Distributed Infrastructures," *ICDCS*, Jul 2006.
- [42] M. Maheswaran, S. Ali, H. J. Siegel, D. Hensgen, and R. F. Freund, "Dynamic Matching and Scheduling of a Class of Independent Tasks onto Heterogeneous Computing Systems," *Heterogeneous Computing Workshop*, Apr 1999.