Multiprocessor Scheduling of Elastic Tasks

James Orr
Washington University in St. Louis
james.orr@wustl.edu

Sanjoy Baruah
Washington University in St. Louis
baruah@wustl.edu

ABSTRACT

The elastic task model enables the adaptation of recurrent real-time tasks under uncertain or potentially overloaded conditions. The model was originally defined for sequential tasks executing upon a preemptive uniprocessor platform; it was later extended to include tasks with internal parallelism executing on multiple processors. This paper bridges a gap in the theory of elastic task scheduling by considering the multiprocessor scheduling of sequential tasks (i.e., tasks with no internal parallelism). We define algorithms for scheduling sequential elastic tasks under the global and partitioned paradigms of multiprocessor scheduling, and provide a simulation-based comparison of the different approaches.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; Real-time systems; • **Software and its engineering** → *Multiprocessing / multiprogramming / multitasking*.

KEYWORDS

elastic scheduling, multi-processor scheduling, real-time systems

ACM Reference Format:
James Orr and Sanjoy Baruah. 2019. Multiprocessor Scheduling of Elastic Tasks. In *27th International Conference on Real-Time Networks and Systems (RTNS 2019), November 6–8, 2019, Toulouse, France.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3356401.3356403

1 INTRODUCTION

Buttazzo et al. introduced the *elastic task model* as a way of modeling recurrent real-time tasks, such as multimedia players or adaptive control systems, whose periods can change depending on the stress on the system [4]. The authors compare real-time tasks to physical springs, where changing a task's period (and therefore processor utilization) is analogous to changing the length of the spring, and keeping the system-wide processor utilization below a certain value is analogous to compressing multiple contiguous springs to below a cumulative length. As originally presented, elastic scheduling seeks to schedule a task set on a single preemptive processor. Each such task has a worst-case execution time and a range of acceptable periods, rather than a single period parameter (as in the original Liu and Layland recurrent task model [18]). Each task must be assigned a period within its acceptable range such that the overall task set utilization remains below a desired value. To determine the appropriate period value to assign each task, every task also has an *elastic coefficient* which acts as an indicator of the task's resistance to increasing its period from the minimum (and desired) period, analogous to a spring's resistance to being compressed.

In the decades since the elastic task model was introduced, real-time systems have increasingly utilized multiple processors, thereby enabling the exploitation of both inter-task and intra-task parallelism — it is appropriate that the elastic task model should also be extended to consider multiprocessors. We have previously extended the elastic task model to include scheduling of tasks with intra-task parallelism on heterogeneous multi-core systems under the federated scheduling paradigm [23]. In this paper we focus on the scheduling of sequential tasks on homogeneous multi-core systems. We present algorithms for scheduling systems of such tasks upon a homogeneous multiprocessor platform under both the global and partitioned paradigms of multiprocessor scheduling. We compare the effectiveness of different algorithms via an extensive series of simulation experiments; based upon the outcomes of these simulations, we make some recommendations regarding the choice of algorithms for the multiprocessor scheduling of sequential elastic tasks.

The remainder of this paper is structured as follows. Section 2 presents our task model. Sections 3 and 4 discuss the global and partitioned scheduling of tasks respectively. Section 5 details our experimental evaluation of the different schemes. Section 6 describes related work, and Section 7 concludes and provides future direction.

2 TASK MODEL AND ASSUMPTIONS

In the model proposed by Buttazzo et al. [4], each elastic task τ_i is characterized by a worst-case execution time (WCET) C_i, a minimum (and preferred) period $T_i^{(\min)}$, a maximum acceptable period $T_i^{(\max)}$, and an elasticity coefficient E_i. The elasticity coefficient is a measure of a task's resistance to changing its period, like a spring's resistance to changing its length in the above analogy. A higher elasticity coefficient indicates a more elastic task. In this work we seek to schedule a set of n such independent sequential elastic tasks $\Gamma = \tau_1 \ldots \tau_n$ on m homogeneous processors.

We refer to elastic tasks represented using this model as *period-elastic tasks*. In prior work [22] we had considered an alternative model of *computationally-elastic tasks* that specifies a range of possible computation times and a fixed period, rather than a range of periods with a fixed execution time, per task. Although the scheduling mechanisms described in this paper are also applicable for computationally-elastic tasks, for simplicity and readability we discuss only period-elastic tasks, as introduced by Buttazzo et al. [4], in the remainder of this paper.

RTNS 2019, November 6–8, 2019, Toulouse, France

ACM ISBN 978-1-4503-7223-7/19/11…$15.00
https://doi.org/10.1145/3356401.3356403

As stated above, a period-elastic elastic task τ_i is characterized by the parameters

$$\tau_i = \left(C_i, T_i^{(\max)}, T_i^{(\min)}, E_i\right) .$$

All the scheduling approaches that we will be considering in this paper have *utilization-based* schedulability conditions: only the utilization parameters of tasks appear in these schedulability conditions. We therefore find it convenient to convert the period parameters of each task (the $T_i^{(\min)}$ and $T_i^{(\max)}$ parameters) to corresponding utilization parameters $U_i^{(\max)}$ and $U_i^{(\min)}$ respectively:

$$U_i^{(\max)} = C_i/T_i^{(\min)}$$
$$U_i^{(\min)} = C_i/T_i^{(\max)}$$

In the remainder of this manuscript, each task τ_i is therefore characterized by the parameters

$$\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)$$

Letting U_i denote the actual utilization "allocated" to τ_i, the desired elasticity property defined by Buttazzo et al. [4] is equivalent to specifying that the amounts by which tasks' utilizations are reduced from their desired maximums be in proportion to their E_i ("elasticity") coefficients:

$$\forall\ i, j,\ \left(\frac{U_i^{(\max)} - U_i}{E_i}\right) = \left(\frac{U_j^{(\max)} - U_j}{E_j}\right) \tag{1}$$

Letting λ denote the desired equilibrium value such that for all tasks $\lambda = \left((U_i^{(\max)} - U_i)/E_i\right)$, Expression 1 suggests

$$U_i \leftarrow U_i^{(\max)} - \lambda E_i$$

However, we also require $U_i \geq U_i^{(\min)}$; hence we choose

$$U_i(\lambda) \leftarrow \max\left(U_i^{(\max)} - \lambda E_i, U_i^{(\min)}\right) \tag{2}$$

Note that for a given value of λ, an elastic task $\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)$ is just a "regular" Liu and Layland task with utilization $U_i(\lambda)$ as given by Expression 2 above.

The problem considered. For each of the multiprocessor scheduling strategies we will study in this paper, the question we ask is: given an n-task system

$$\Gamma = \left\{\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)\right\}_{i=1}^{n}$$

that is to be scheduled upon an m-processor platform, what is the smallest value of λ for which the Liu and Layland task system comprising n tasks with utilizations $U_1(\lambda), U_2(\lambda), \ldots, U_n(\lambda)$ is successfully schedulable by that particular scheduling strategy?

3 GLOBAL SCHEDULING

Under the global paradigm of multiprocessor scheduling for recurrent tasks, individual tasks are not restricted to executing upon specific processors. Instead, a newly-arrived job of a task may begin execution upon any available processor and a preempted job may resume execution at a later point in time upon any processor, not just the one it had been executing upon prior to preemption. We consider three different global scheduling algorithms: fluid (Section 3.1), Earliest Deadline First (Section 3.2), and an algorithm called PriD [13] that can be thought of as a generalization of EDF (Section 3.3).

3.1 Fluid Scheduling

The *fluid scheduling* paradigm of multiprocessor real-time scheduling permits that individual tasks be assigned a fraction $f, 0 \leq f \leq 1$, of a processor at each instant in time (in contrast to non-fluid schedules, in which each task may execute either upon zero processors or upon a single processor at each instant). Fluid scheduling is a convenient abstraction that considerably simplifies many multiprocessor real-time scheduling problems; techniques are known (see, e.g, [2, 14, 17, 20]) for converting fluid schedules to non-fluid ones for many problems and under a wide range of conditions and circumstances.

Fluid scheduling of Liu and Layland tasks – a review. Consider some Liu and Layland task system Γ, and let U_i denote the utilization of $\tau_i \in \Gamma$. It has been shown [14] that a necessary and sufficient condition for Γ to be fluid-schedulable upon a multiprocessor platform comprising m unit-speed processors is that

$$max_{\tau_i \in \Gamma}\{U_i\} \leq 1 \tag{3}$$

and

$$\left(\sum_{\tau_i \in \Gamma} U_i\right) \leq m \tag{4}$$

Any task system satisfying Conditions 3 and 4 can be fluid-scheduled by simply assigning each job of τ_i a fraction U_i of one of the m processors at each instant between its release date and its deadline.

Extension to period-elastic tasks. In the original elastic scheduling paper [4], Buttazzo et al. present an iterative ($\Theta(n^2)$) algorithm called Task_Compress(Γ,U_d) for assigning a period to each task in a system Γ of period-elastic tasks such that the total system utilization stays below a desired value U_d — this algorithm is reproduced in this paper as Algorithm 1. It is evident that Algorithm 1 is, in essence, determining the smallest value of λ for which

$$\left(\sum_{i=1}^{n} U_i(\lambda)\right) \leq U_d ,$$

where the $U_i(\lambda)$'s are as defined according to Expression 2. Observe, too, that Algorithm 1 never *increases* the actual utilization assigned to any any task τ_i to beyond $U_i^{(\max)}$ — this follows from the observation that in Line 19, the value assigned to the actual utilization —the parameter U_i— is obtained by *subtracting* a positive quantity from $U_1^{(\max)}$. Hence given an elastic task system Γ of sequential tasks that is to be fluid-scheduled upon m unit-speed processors, we can determine the effective utilizations of the individual tasks that satisfy Conditions 3 and 4, and therefore bear witness to the fluid-schedulability of Γ, by simply calling the procedure Task_Compress(Γ,U_d) of Algorithm 1 with $U_d \leftarrow m$. The instance Γ can then be fluid-scheduled by assigning each job of each $\tau_i \in \Gamma$ a fraction of a processor equal to this effective utilization at each instant between its release date and its deadline.

Algorithm 1 Task_Compress(Γ,U_d)

$U^{(\max)} = \sum_{i=1}^{n} C_i/T_i^{(\min)}$
$U^{(\min)} = \sum_{i=1}^{n} C_i/T_i^{(\max)}$
if $U_d < U^{(\min)}$ **then**
return INFEASIBLE
end if
ok= 0
while ok == 0 **do**
  $U_f = E_v = 0$
  for *each* τ_i **do**
    if E_i == 0 **or** $T_i == T_i^{(\max)}$ **then**
      $U_f = U_f + U_i$
    else
      $E_v = E_v + E_i$
    end if
  end for
  ok= 1
  for *each* $\tau_i \in \Gamma_v$ **do**
    if $E_i > 0$ **and** $T_i < T_i^{(\max)}$ **then**
      $U_i = U_i^{(\max)} - (U^{(\max)} - U_d + U_f) * E_i/E_v$
      $T_i = C_i/U_i$
      if $T_i > T_i^{(\max)}$ **then**
        $T_i = T_i^{(\max)}$
        ok= 0
      end if
    end if
  end for
end while
return FEASIBLE

3.2 Global EDF

While the fluid scheduling model is a convenient abstraction for considering multiprocessor scheduling, it is not in general directly implementable. As mentioned above, techniques are known for converting fluid schedules to non-fluid ones under a variety of conditions; however, most such conversions yield schedules with a large number of preemptions and inter-processor migrations. In environments in which there is a considerable overhead associated with each preemption and/ or inter-processor migration, this approach of obtaining a fluid schedule and then converting to a non-fluid one may incur unacceptably high overhead costs.

Review of results for Liu and Layland tasks. The global Earliest Deadline First (EDF) scheduling algorithm has the property that the total number of preemptions and inter-processor migrations in a schedule is bounded from above at the number of jobs in the schedule. (This is easily seen by observing that a job may preempt an already-executing one only upon its arrival, if it happens to have an earlier deadline; such preemption may later lead to an inter-processor migration if the preempted job resumes upon a different processor.) Global EDF may therefore be a more appropriate algorithm to use in environments characterized by significant preemption/ migration overhead costs. Goossens et al. showed [13, Theorem 5] that a system Γ of Liu & Layland tasks is scheduled by global EDF to meet all deadlines upon m unit-speed processors if the following condition holds:

$$\sum_{\tau_i \in \Gamma} U_i \leq m - (m-1) \times \max_{\tau_i \in \Gamma} \{U_i\} \tag{5}$$

(This condition was also shown [13, Theorem 6] to be tight from a utilization-based perspective: there are systems in which $\left(\sum_{\tau_i \in \Gamma} U_i\right)$ is greater than $\left(m - (m-1) \times \max_{\tau_i \in \Gamma} \{U_i\}\right)$ by an arbitrarily small amount, upon which global EDF misses deadlines.)

Extension to period-elastic tasks. Given a system Γ of period-elastic tasks

$$\Gamma = \left\{\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)\right\}_{i=1}^{n}$$

that is to be scheduled upon an m-processor platform, our objective is to find the smallest value of λ such that the Liu & Layland task system with the following utilizations

$$U_i \leftarrow \left\{\max\left(U_i^{(\max)} - \lambda E_i, U_i^{(\min)}\right)\right\}_{i=1}^{n} \tag{6}$$

is schedulable using global EDF. We have chosen to solve this problem by iterating through the possible values of λ — see Algorithm 2. This algorithm steps through the range $[0, \Phi]$ with a "granularity" ϵ (Line 1 of Algorithm 2), where Φ is the maximum value among all tasks of the equation $\left(\frac{U_i^{(\max)} - U_i^{(\min)}}{E_i}\right)$. The algorithm seeks the smallest value of λ or which the Liu & Layland task system of Expression 6 above is global EDF-schedulable according to Expression 5. Once this smallest value of λ is determined and returned by Algorithm 2, we can convert the period-elastic task system to a regular Liu & Layland task system by computing the effective utilizations of the tasks according to Expression 2, and then schedule the Liu & Layland task system so obtained by global EDF. Algorithm 2 is $\Theta(\frac{\Phi}{\epsilon} \times n)$.

Algorithm 2 Global EDF(Γ,m)

1: $\epsilon \leftarrow 0.05 \times \Phi$ ▷ "Granularity" of the test...
2: **for** $\lambda \leftarrow 0$ **to** Φ **by** ϵ **do**
3: $S \leftarrow 0.0$ ▷ Total utilization of compressed tasks
4: $M \leftarrow 0.0$ ▷ Max. utilization amongst compressed tasks
5: **for** $i \leftarrow 1$ **to** $|\Gamma|$ **do**
6: $tmp \leftarrow \max\left(U_i^{(\max)} - \lambda E_i, U_i^{(\min)}\right)$
7: $S \leftarrow S + tmp$
8: $M \leftarrow \max(M, tmp)$
9: **end for**
10: **if** $(S \leq m - (m-1) \times M)$ **then**
11: ▷ By Eqn. 5, the compressed tasks are global-EDF schedulable,
12: **return** λ
13: **end if**
14: **end for**
15: **return** *(global EDF fails)*

Algorithm PriD (Γ, m)
The Liu & Layland task system $\Gamma = \{\tau_1, \tau_2, \ldots \tau_n\}$ is to be scheduled on m processors
Assume the tasks are indexed according to utilization: $U_i \geq U_{i+1}$ for all i, $1 \leq i < n$
for $i = 1$ **to** m **do**
if $\{\tau_{i+1}, \tau_{i+2}, \ldots, \tau_n\}$ is global-EDF schedulable upon $(m - i)$ processors
then
During run-time $\{\tau_1, \tau_2, \ldots, \tau_i\}$'s jobs will be assigned highest priority and $\{\tau_{i+1}, \tau_{i+2}, \ldots, \tau_n\}$'s jobs will be assigned EDF-priority
return success
return failure *// Not schedulable by PriD*

Figure 1: Algorithm PriD priority-assignment rule

3.3 Algorithm PriD

It was observed [1] that global EDF tends to under-perform when there is even a single task with high utilization. This is easily explained by examining the utilization-based global-EDF schedulability condition of Inequality 5: observe the presence of the

$$\left((m-1) \times \max_{\tau_i \in \Gamma} \{U_i\}\right)$$

term on the right-hand side. Since this term is *subtracted* from the total computing capacity of the platform (i.e., m), the consequence is that a capacity of $(m-1)$ times the largest individual utilization becomes unavailable due to the presence of this large-utilization task. This phenomenon can be looked upon a consequence of the well-known *Dhall effect* [9, 10] which has been widely studied in multiprocessor real-time scheduling theory. Several results have been obtained within the real-time scheduling theory community for dealing with such utilization loss; below we first review some of these results and then seek to extend their applicability to incorporate period-elasticity.

Review of results for Liu and Layland tasks. Recall that one major advantage of EDF-generated schedules over those obtained by converting a fluid-based one is the reduced number of preemptions and inter-processor migrations: the total number of preemptions and migrations in an EDF-generated is no greater than the number of jobs that are scheduled. It turns out that this property is in fact enjoyed by an entire class of algorithms: all those in which each job is assigned a single fixed priority and at each instant during run-time the highest-priority jobs that are eligible to execute are selected for execution. Algorithms in this class are referred to as *Fixed Job Priority* (FJP) [7] scheduling algorithms. The algorithm fpEDF was proposed [1] as an FJP algorithm that circumvents the utilization loss caused by the Dhall effect. Under the fpEDF run-time scheduling algorithm, jobs of tasks with utilization > 0.5 are statically assigned highest priority while priorities to jobs of the remaining tasks are assigned according to EDF. It has been shown [1, Theorem 4] that a task system Γ is scheduled by fpEDF to meet all deadlines upon m unit-speed processors if the following condition holds:

$$\sum_{\tau_i \in \Gamma} U_i \leq \frac{m+1}{2} \tag{7}$$

A pragmatic improvement to fpEDF, called Algorithm PriD (for "priority driven") was proposed by Goossens et al. [13] — this is the algorithm that we will be adapting below for period-elastic tasks. Algorithm PriD is presented in pseudo-code form in Figure 1. Algorithm PriD, like fpEDF, seeks to circumvent the Dhall effect by assigning greatest priority to jobs of tasks with high utilization; however, while fpEDF designates all tasks with utilization > 0.5 to be "high-utilization" ones, Algorithm PriD determines which tasks are "high-utilization" based on the characteristics of the task system under consideration. It is shown [13] that Algorithm PriD strictly dominates fpEDF: all instances that are deemed schedulable by fpEDF are also deemed schedulable by PriD while the converse of this statement is not true – there are instances deemed schedulable by Algorithm PriD that will not pass the fpEDF schedulability test of Expression 7.

Extension to period-elastic tasks. Our adaptation of Algorithm PriD to period-elastic tasks is similar to our adaptation of global EDF: given an instance of periodic-elastic tasks

$$\Gamma = \left\{\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)\right\}_{i=1}^{n}$$

to be scheduled upon m unit-speed processors, we iterate through possible values of λ between 0 and Φ, seeking the smallest value such that the Liu & Layland task system with utilizations

$$U_i \leftarrow \left\{\max(U_i^{(\max)} - \lambda E_i, U_i^{(\min)}\right\}_{i=1}^{n}$$

is deemed schedulable by Algorithm PriD upon m unit-speed processors. (The pseudo-code for this algorithm is very similar to the pseudo-code in Algorithm 2, and hence omitted.)

Algorithm PriD is $\Theta\Big(n \times log(n) + m\Big)$. Therefore, the overall complexity of iterating over λ values for elastic tasks to be scheduled under Algorithm PriD is $\Theta\Big(\frac{\Phi}{\epsilon} \times (n \times log(n) + m)\Big)$

4 PARTITIONED SCHEDULING

The partitioned scheduling of Liu & Layland task systems is known to be equivalent to the bin-packing problem[15, 16], and hence NP-hard in the strong sense. Several polynomial-time heuristics have been proposed for solving this problem approximately: most of these heuristic algorithms for partitioning have the following common structure. First, they specify an order in which the tasks are to be considered. Then in considering each task (in the order chosen), they specify the order in which to consider upon which

processor to attempt to allocate the task. A task is successfully allocated upon a processor if it is observed to "fit" upon the processor; within the context of the partitioned EDF-scheduling, a task fits on a processor if the task's utilization does not exceed the processor capacity minus the sum of the utilizations of all tasks previously allocated to the processor. The algorithm declares success if all tasks are successfully allocated; otherwise, it declares failure.

Lopez et al. [19] have extensively compared several widely-used heuristic algorithms that fit this overall structure. They define the concept of a *Reasonable Allocation* (RA) partitioning algorithm: an RA algorithm is one that fails to allocate a task to a multiprocessor platform only when the task does not fit into any processor upon the platform. All the heuristic algorithms considered by Lopez et al. [19] are RA ones — indeed, there seems to be no reason why a system designer would ever consider using a non-RA partitioning algorithm. Within the RA algorithms, Lopez et al. [19] compared heuristics that

(1) use three different ways for ordering the tasks to consider: arbitrary, in order of increasing utilization, and in order of decreasing utilization; and
(2) also use three different heuristics for ordering the processors to consider: "first fit" (assign a task to the first processor upon which it fits), "worst fit" (assign a task to the processor with the maximum remaining capacity), and and "best fit" (assign a task to the processor with the minimum remaining capacity that exceeds the task's utilization).

Extension to period-elastic tasks. Any of the partitioning heuristics can be adapted for period-elastic tasks in a manner that is very similar in structure to the manner in which global EDF and PriD were adapted for elastic tasks. That is, given an instance of periodic-elastic tasks

$$\Gamma = \left\{\tau_i = (U_i^{(\max)}, U_i^{(\min)}, E_i)\right\}_{i=1}^{n}$$

to be scheduled upon m unit-speed processors, we iterate through possible values of λ between 0 and Φ, seeking the smallest value such that the Liu & Layland task system with utilizations

$$U_i \leftarrow \left\{\max(U_i^{(\max)} - \lambda E_i, U_i^{(\min)}\right\}_{i=1}^{n}$$

is deemed schedulable by upon m unit-speed processors by the partitioning heuristic. (The pseudo-code for doing so is again very similar to the pseudo-code in Algorithm 2, and hence omitted.)

We note that after partitioning tasks onto processors, it is highly unlikely that all processors are fully utilized (i.e. the assigned utilizations of the partitioned tasks sum to 1.0). The procedure Task_Compress(Γ,U_d) of Buttazzo et al. [4] (reproduced here as Algorithm 1) can be applied to each processor with $U_d \leftarrow 1.0$ to perhaps increase system utilization while still guaranteeing schedulability,

Sorting tasks and partitioning them is $\Theta\Big(n \times log(n) + n \times m\Big)$. Therefore, the overall complexity of iterating over λ values for elastic tasks to be scheduled under partitioned scheduling is $\Theta\Big(\frac{\Phi}{\epsilon} \times (n \times log(n) + n \times m)\Big)$

5 SIMULATION EXPERIMENTS

We have performed a simulation-based comparison of the various algorithms presented in Sections 3 and 4 for the multiprocessor scheduling of sequential period-elastic tasks; we report on the findings of this comparison below. We describe the setup for these simulation experiments in Section 5.1 and present our findings in Section 5.2; based upon these findings, we draw some high-level conclusions in Section 5.3.

5.1 Experimental Setup

We randomly generate sets of sequential period-elastic tasks and attempt to schedule them upon a given number of processors m using the different scheduling algorithms – fluid, global EDF, PriD, and partitioned – described in Sections 3 and 4 above. Specifically,

- We separately consider multiprocessor platforms containing $m = 4$, 8, and 16 identical processors.
- For each of these values for m, we consider task sets with $n = 2 \times m$, $2.5 \times m$, $3 \times m$, $n = 4 \times m$, and $n = 8 \times m$ tasks.
- For each combination of values of m and n, we also vary the maximum utilization value any individual task is allowed to be assigned, denoted α. This value can directly impact schedulability of a task set, particularly when using the global EDF and PriD algorithms. We study values of of $\alpha = \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$.
- For each selected combination of values of m and n, we generate task sets in which the maximum utilizations of the tasks (i.e., their $U_i^{(\max)}$ parameters) sum to $1.1 \times m \times \alpha$, $1.5 \times m \times \alpha$, and $1.9 \times m \times \alpha$.

Hence a total of $3 \times 5 \times 6 \times 3 = \mathbf{270}$ different combinations of values of m, n, α, and $\left(\sum_i U_i^{(\max)}\right)$ are considered. For each such combination, we generate 1000 task sets in the following manner. We generate the individual $U^{(\max)}$ values using the *Randfixedsum* algorithm [11] to provide an unbiased distribution of maximum utilizations. The corresponding individual task minimum utilization values $U_i^{(\min)}$ are uniformly generated over the range $(0, U_i^{(\max)})$. In the case that a task set's $U_i^{(\min)}$ values sum to more than m (i.e. the task set is not schedulable under fluid scheduling, or therefore, any other scheduling algorithm), we repeatedly generated new $U_i^{(\min)}$ values for each task until their sum is sufficiently low. Tasks' elastic coefficients is chosen uniformly randomly over the range [1, 5]. For all algorithms a "granularity" of $\epsilon = \frac{\Phi}{1,000}$ was used.

We attempt to schedule each task set generated as described above using the four algorithms discussed in Sections 3 and 4: fluid, global EDF, PriD, and partitioned. For partitioned, we first sort the tasks in order of decreasing utilization (their $U_i(\lambda)$ parameters), and attempt to assign them to the available processors using the the "first-fit," "worst-fit," and "best-fit" heuristics. We return the first λ value that deems the task set schedulable by any of these heuristics. We note that the ability to partition tasks onto processors in an efficient manor is an advantage over the global scheduling algorithms considered in this paper, as it is infeasible to carry out full simulation of global EDF or PriD.

5.2 Observations

In our experiments, we noted (i) the fraction of task-sets that were determined to be schedulable by each of our four algorithms; and (ii) for those task-sets that were deemed schedulable by all the algorithms, the minimum λ needed to achieve schedulability by each algorithm. Our results are presented in graphical form in Figures 2–4. In these graphs we show results of both the average minimum normalized λ value ($\frac{\lambda}{\Phi}$–this gives a value on the interval $[0, 1]$ and is needed to compare λ values across task sets) needed to achieve schedulability for a given scheduling algorithm, and the percentage of the 1,000 task sets that each algorithm deemed scheduleable. To ensure a consistent comparison, we only compare lambda values for task sets deemed schedulable by all scheduling algorithms.

As mentioned in Subsection 5.1, there are 270 combinations of m, n, α, and $\left(\sum_i U_i^{(\max)}\right)$ considered in this simulation. Due to spatial constraints, this paper includes the representative subset of $\alpha = \{0.6, 0.8, 1.0\}$. The remaining data in which $\alpha = \{0.5, 0.7, 0.9\}$ and follows the same patterns discussed in this section can be found online [21].

Figures 2, 3, and 4 show both the λ values and percentage of schedulable task sets for all four scheduling algorithms (fluid, global EDF, PriD, and partitioned) for $\alpha = 0.6, 0.8$, and 1.0, respectively. Some trends can be noticed across all graphs. We note that fluid scheduling is an idealized optimal scheduling algorithm; not surprisingly, therefore, it schedules the largest percentage of task-sets and returns the smallest λ value. This is seen consistently across all results. It serves as an upper bound for achievable simulation results for the other scheduling algorithms. We also note that partitioned scheduling consistently dominates algorithm PriD and global EDF in both λ value and in percentage of schedulable task sets. This is consistent with prior observations [3] regarding global versus partitioned multiprocessor scheduling; in essence, this is likely a reflection of the fact that while global scheduling algorithms like PriD apply schedulability tests that are utilization-based and incorporate considerable pessimism since they must consider "worst-case" task-sets with the same utilization parameters as running full simulations is infeasible, partitioned schedulability tests actually attempt to perform a partition and hence do not necessarily pay the price in terms of such analysis-based pessimism.

Note that global EDF scheduling always requires the highest *lambda* value, and that the percentage of task sets deemed schedulable under global EDF decreases as $\left(\sum_i U_i^{(\max)}\right)$ increases. This is a manifestation of the Dhall effect, and our experiments revealed that this worsens as the number of processors and tasks increase: for some combinations of m, n, and $\left(\sum_i U_i^{(\max)}\right)$ that we considered, global EDF fails to schedule a single task set out of 1000. In such cases the reported λ is off the chart. The Dhall effect can be observed to worsen as α increases.

Our experiments also reveal that it becomes more difficult to schedule tasks (in terms of both λ value and schedulability percentage) for all the scheduling algorithms as $\left(\sum_i U_i^{(\max)}\right)$ increases. The same is true as the number of processors increases but the ratio of processors to tasks remains the same. On a constant number of processors, fluid and partitioned scheduling can return a lower λ value with more tasks in the task set, and a higher percentage of task sets are deemed schedulable under partitioned scheduling (while PriD and global-EDF seem less affected). We believe this improvement seen to be a reduction in the Dhall effect: as more tasks are introduced into the system the largest single task is more likely to decrease. Naturally fluid scheduling always deems 100% of tasks to be schedulable.

5.3 Recommendation

Based on our observations in the previous subsection and the graphs in Figures 2–4, we recommend that in the absence of specific knowledge regarding task characteristics that may advocate in favor of PriD or global EDF, partitioned scheduling be used for the scheduling of sequential period-elastic tasks on uniform multiprocessor systems, particularly in systems with a large number of tasks. Among the realistic scheduling algorithms considered in this paper, it 1) consistently returns the lowest value of λ (and therefore compresses tasks the least) and 2) schedules the highest percentage of task sets.

6 RELATED WORK

Buttazzo et al. first introduced the elastic task model for sequential tasks on a preemptive uniprocessor [4]. The sequential model was later extended to include resource sharing [5] and unknown computational loads [6]. Chantem et al. proved Buttazzo's initial scheduling algorithm to be equivalent to solving a quadratic optimization problem and introduced a period-based optimization problem scheme for period selection [8]; they further extended the model to include constrained deadlines [8]. Our prior work introduced elastic scheduling of tasks with internal parallelism under the federated scheduling paradigm [23] and the concept of computational elasticity [22]. Recent work by Gill et al., has applied elastic scheduling to mixed-criticality systems [12]. We leave the sequential multi-core scheduling extensions to all of these problems as future work.

7 CONCLUSION

In this paper we have introduced elastic scheduling for sequential tasks on multiprocessor systems. We have introduced algorithms for scheduling such tasks under both global (in a variety of manners) and partitioned scheduling paradigms. We ran an extensive simulation to compare these methods and conclude that partitioned scheduling should be used if possible.

ACKNOWLEDGMENTS

This research was supported in part by NSF grants CSR-1911460 Medium "Resource Efficient Implementation of Mixed-Criticality Systems" and CSR-1814739 Small "Dynamically Customizable Safety Critical Embedded Systems".

REFERENCES

[1] Sanjoy Baruah. 2004. Optimal utilization bounds for the fixed-priority scheduling of periodic task systems on identical multiprocessors. *IEEE Trans. Comput.* 53, 6 (2004).

[2] Sanjoy Baruah, Neil Cohen, Greg Plaxton, and Don Varvel. 1996. Proportionate progress: A Notion of Fairness in Resource Allocation. *Algorithmica* 15, 6 (June 1996), 600–625.

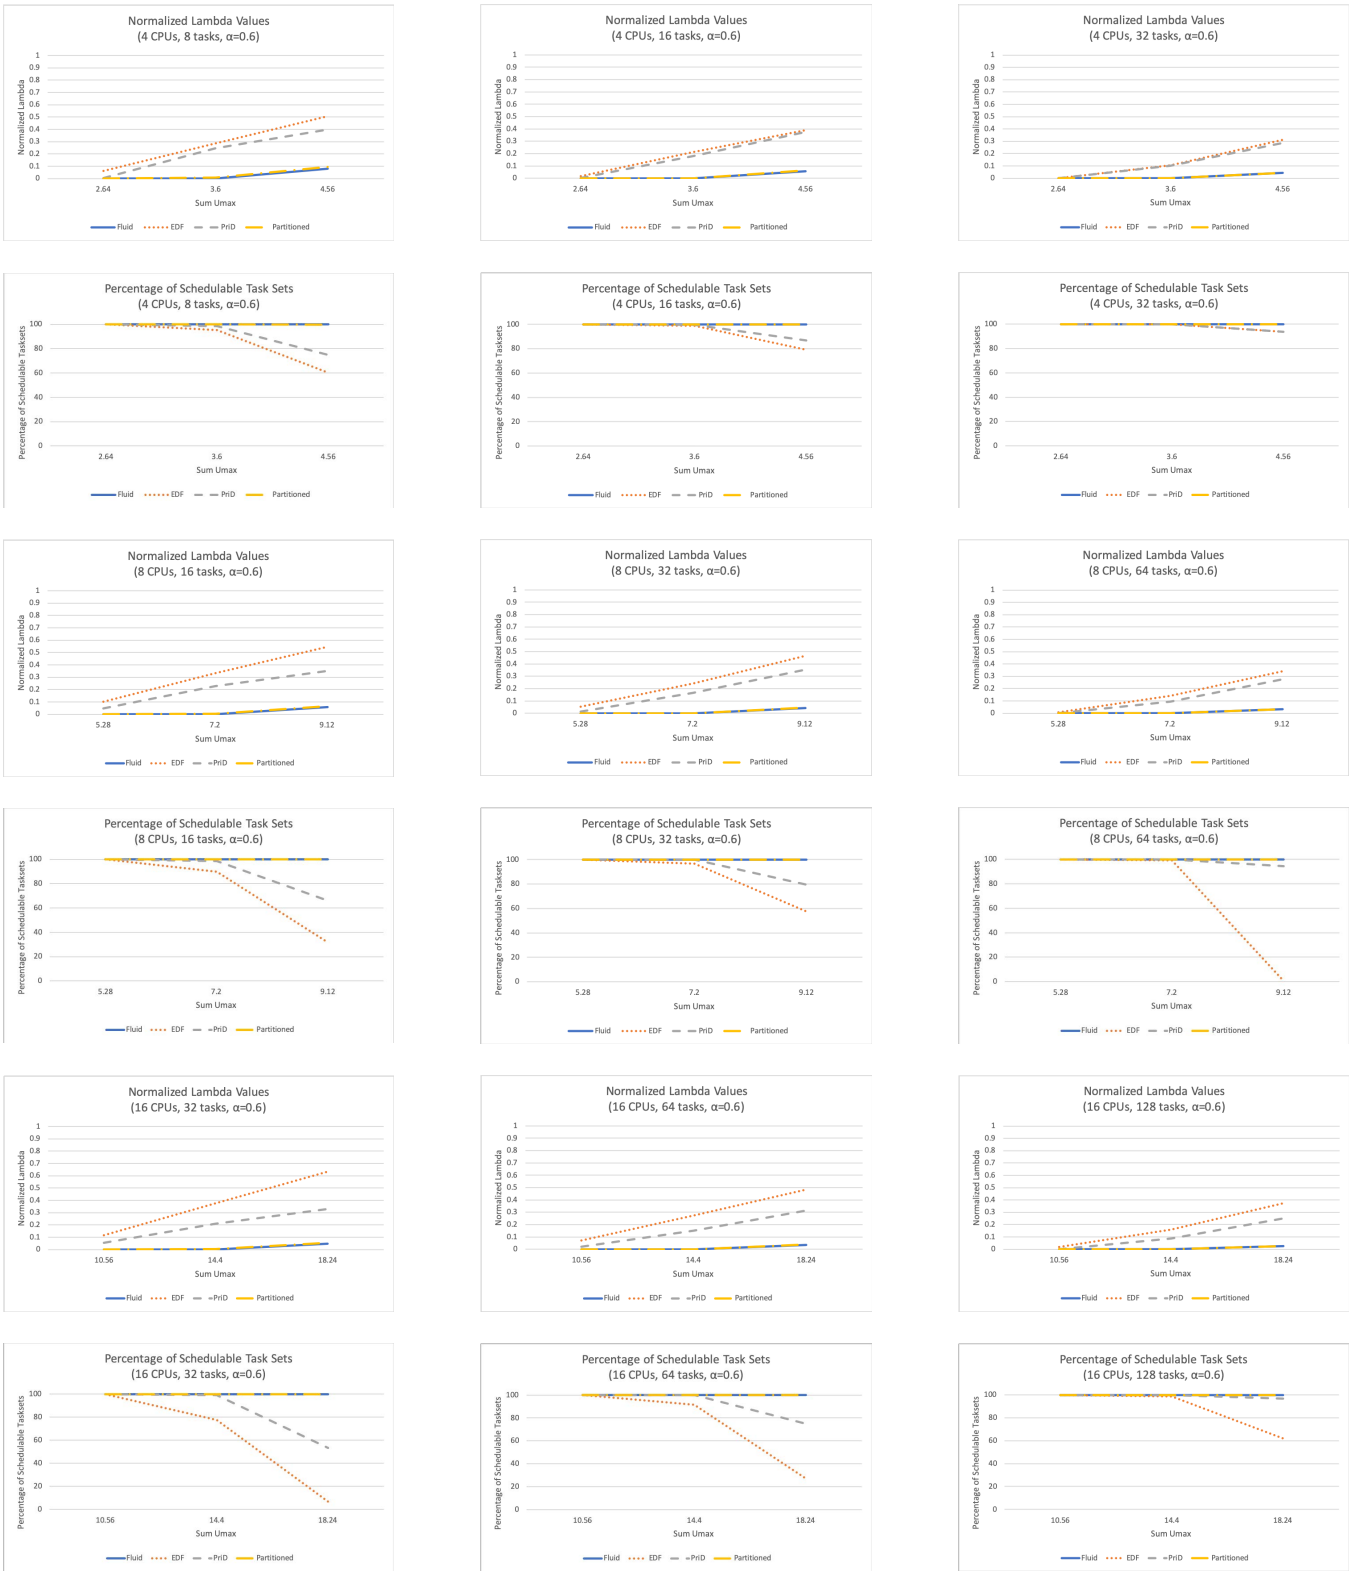

Figure 2: Lambda Values and Schedulability ($\alpha = 0.6$)

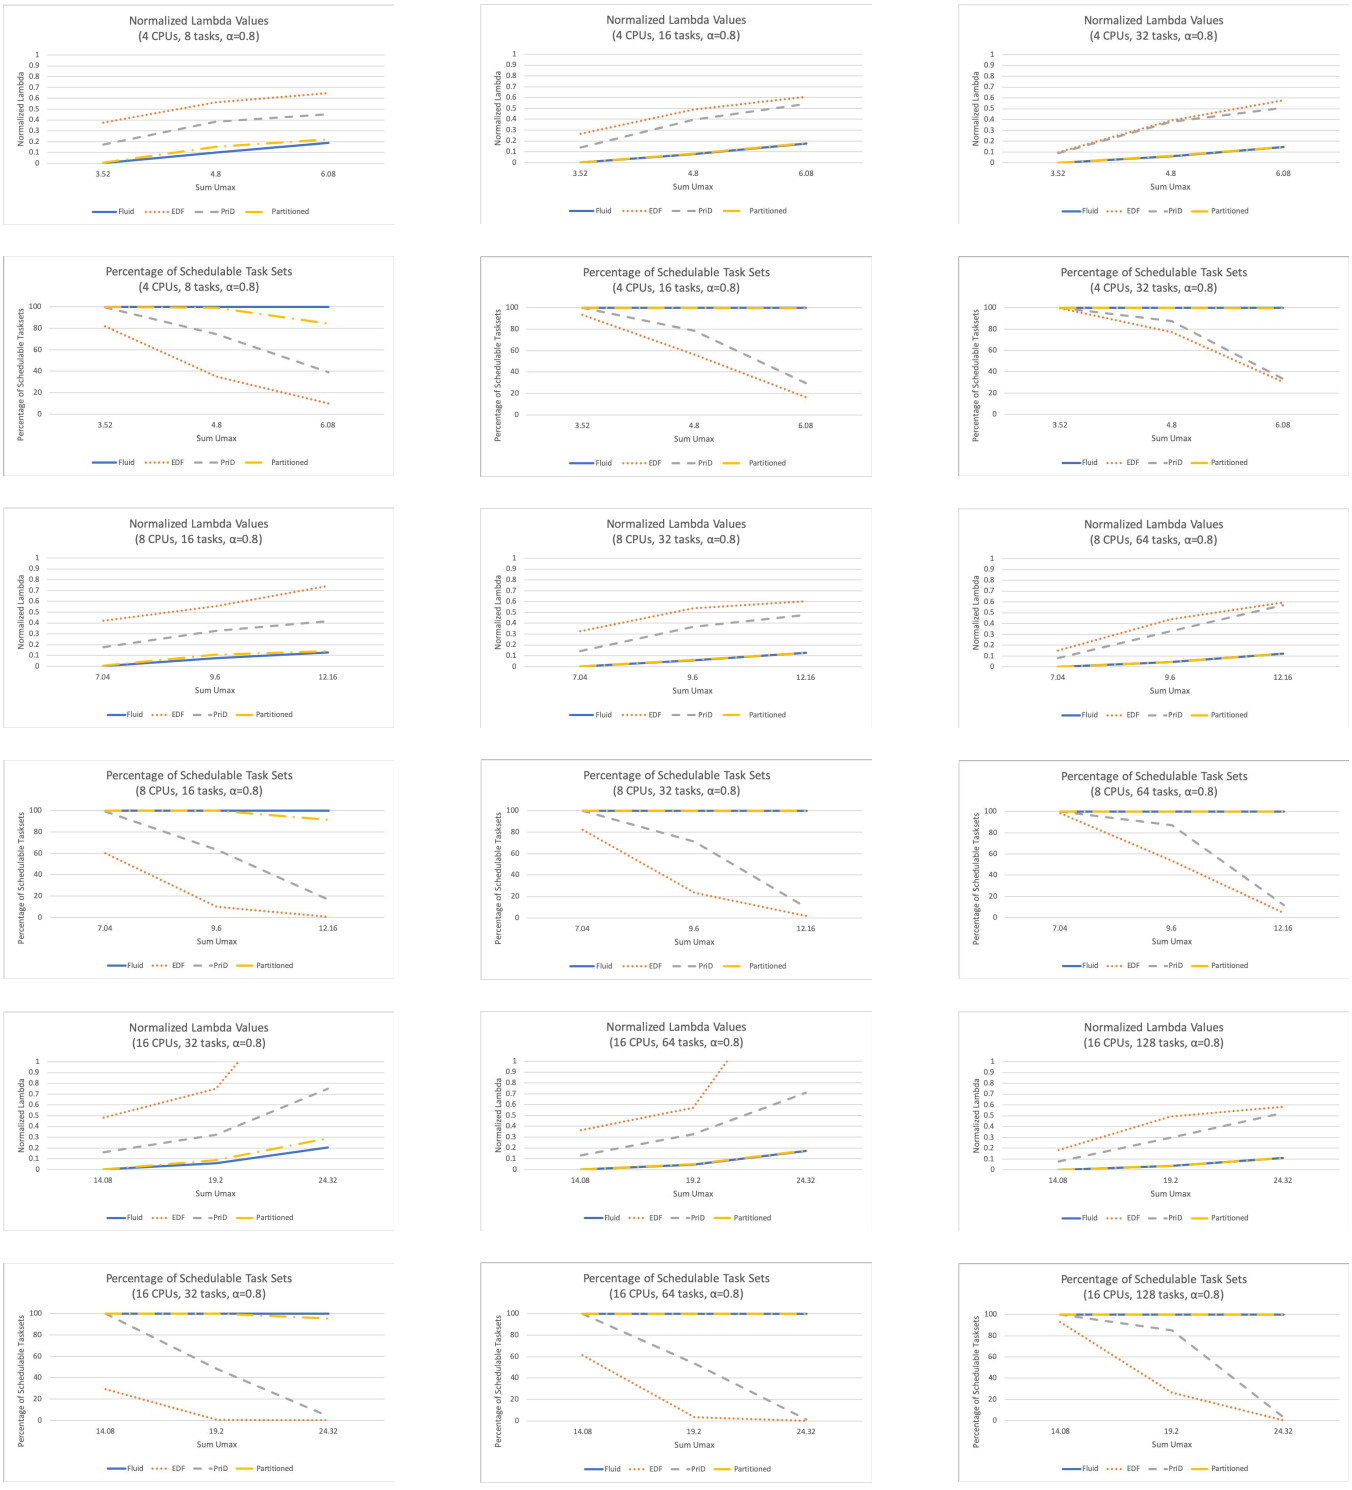

Figure 3: Lambda Values and Schedulability ($\alpha = 0.8$)

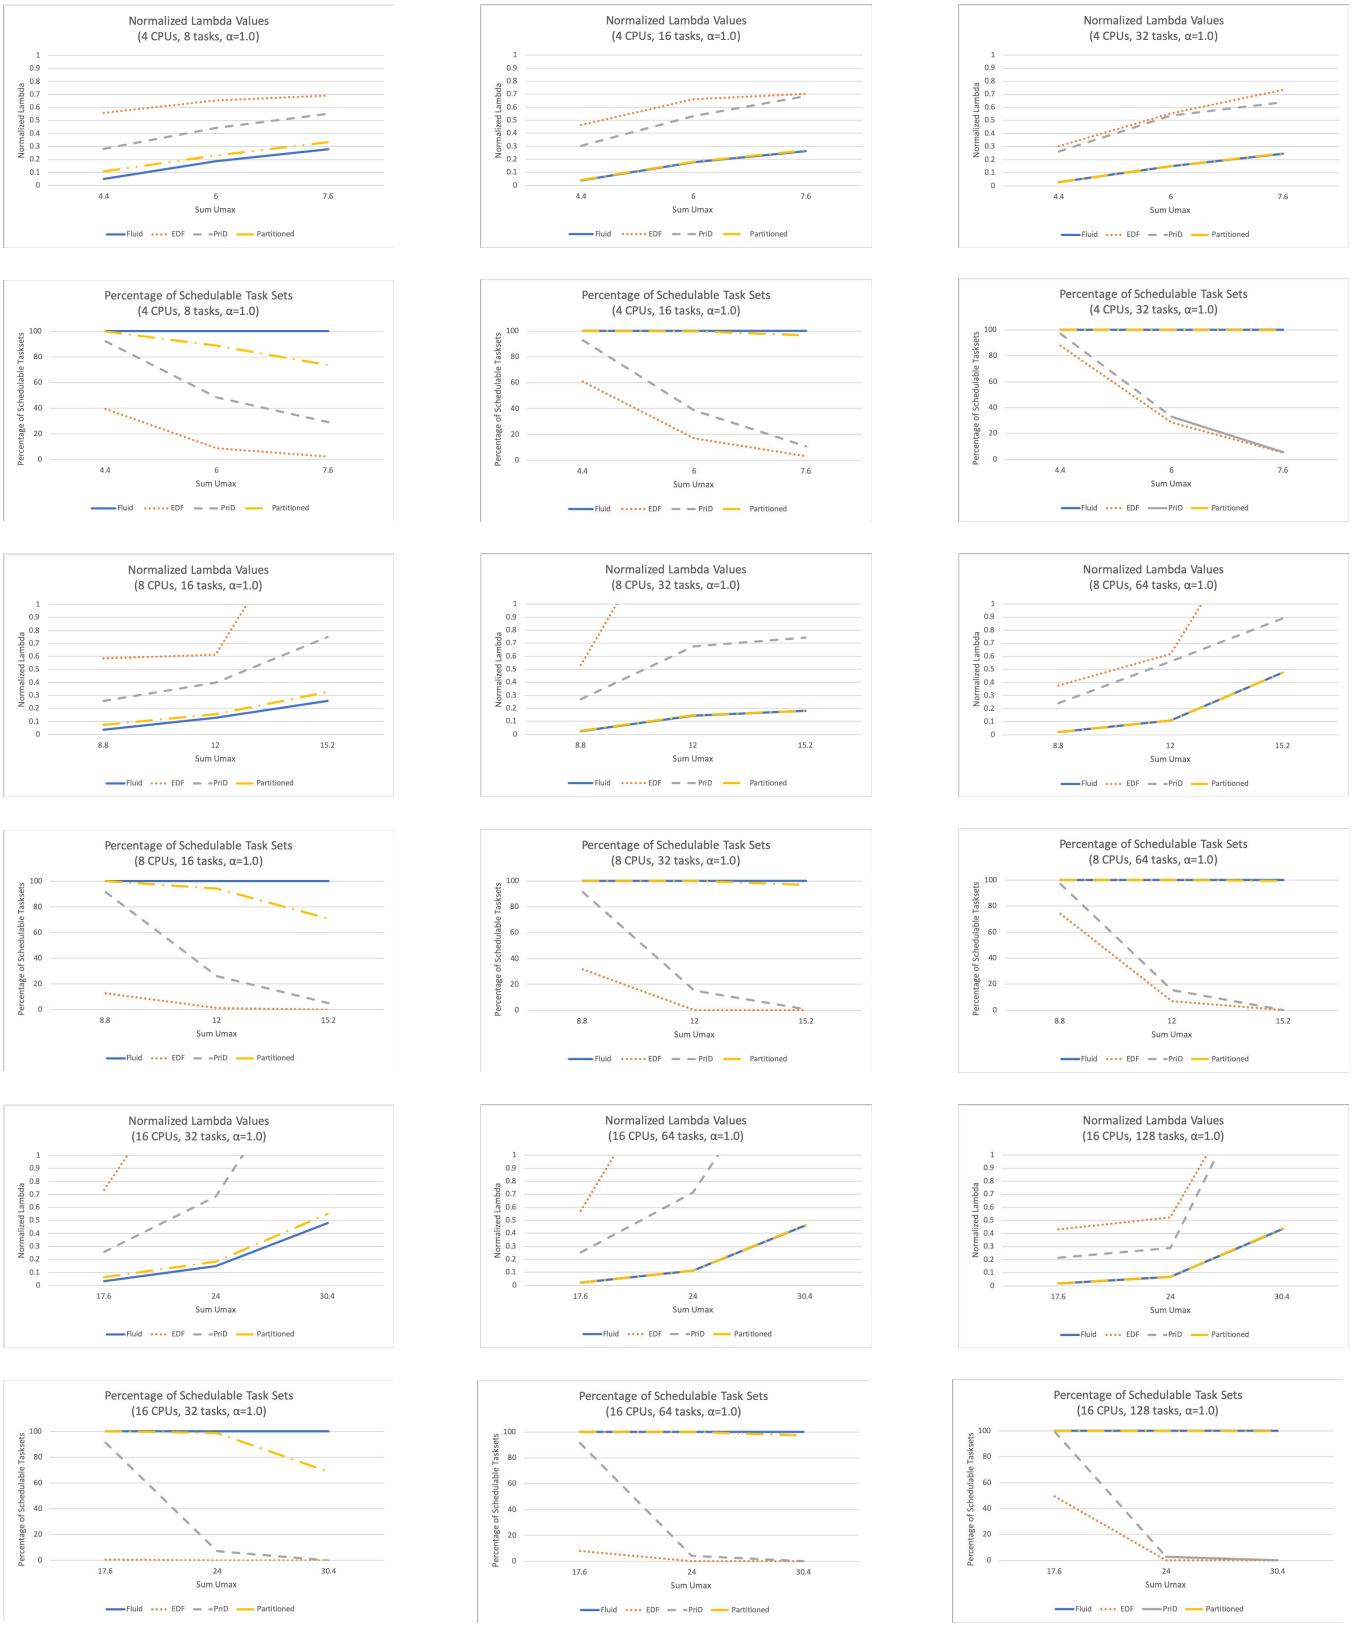

Figure 4: Lambda Values and Schedulability ($\alpha = 1.0$)

[3] A. Bastoni, B. Brandenburg, and J. Anderson. 2010. An Empirical Comparison of Global, Partitioned, and Clustered Multiprocessor Real-Time Schedulers. In *Proceedings of the Real-Time Systems Symposium*. IEEE Computer Society Press, San Diego, CA, 14–24.

[4] Giorgio C. Buttazzo, Giuseppe Lipari, and Luca Abeni. 1998. Elastic Task Model for Adaptive Rate Control. In *IEEE Real-Time Systems Symposium (RTSS)*.

[5] Giorgio C. Buttazzo, Giuseppe Lipari, Marco Caccamo, and Luca Abeni. 2002. Elastic Scheduling for Flexible Workload Management. *IEEE Trans. Comput.* 51, 3 (March 2002), 289–302. https://doi.org/10.1109/12.990127

[6] M. Caccamo, G. Buttazzo, and Lui Sha. 2000. Elastic feedback control. In *Proceedings 12th Euromicro Conference on Real-Time Systems. Euromicro RTS 2000.* 121–128. https://doi.org/10.1109/EMRTS.2000.853999

[7] John Carpenter, Shelby Funk, Phil Holman, Anand Srinivasan, Jim Anderson, and Sanjoy Baruah. 2003. A Categorization of Real-time Multiprocessor Scheduling Problems and Algorithms. In *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, Joseph Y.-T Leung (Ed.). CRC Press LLC.

[8] T. Chantem, X. Hu, and M. Lemmon. 2009. Generalized Elastic Scheduling for Real-Time Tasks. *IEEE Trans. Comput.* 58, 4 (April 2009), 480–495. https://doi.org/10.1109/TC.2008.175

[9] Sudarshan Dhall. 1977. *Scheduling Periodic Time-Critical Jobs on Single Processor and Multiprocessor Systems.* Ph.D. Dissertation. Department of Computer Science, The University of Illinois at Urbana-Champaign.

[10] S. K. Dhall and C. L. Liu. 1978. On a Real-Time Scheduling Problem. *Operations Research* 26 (1978), 127–140.

[11] P Emberson, R Stafford, and R.I. Davis. 2010. Techniques For The Synthesis Of Multiprocessor Tasksets. *WATERS'10* (01 2010).

[12] C. Gill, J. Orr, and S. Harris. 2018. Supporting Graceful Degradation through Elasticity in Mixed-Criticality Federated Scheduling. In *Proceedings of the 6th International Workshop on Mixed Criticality Systems (WMC).*

[13] Joel Goossens, Shelby Funk, and Sanjoy Baruah. 2003. Priority-driven Scheduling of Periodic Task Systems On Multiprocessors. *Real Time Systems* 25, 2–3 (2003), 187–205.

[14] W. Horn. 1974. Some simple scheduling algorithms. *Naval Research Logistics Quarterly* 21 (1974), 177–185.

[15] David Johnson. 1974. Fast algorithms for bin packing. *Journal of Computer and Systems Science* 8, 3 (1974), 272–314.

[16] D. S. Johnson. 1973. *Near-optimal Bin Packing Algorithms.* Ph.D. Dissertation. Department of Mathematics, Massachusetts Institute of Technology.

[17] Jaewoo Lee, Kieu-My Phan, Xiaozhe Gu, Jiyeon Lee, A. Easwaran, Insik Shin, and Insup Lee. 2014. MC-Fluid: Fluid Model-Based Mixed-Criticality Scheduling on Multiprocessors. In *Real-Time Systems Symposium (RTSS), 2014 IEEE.* 41–52.

[18] C. Liu and J. Layland. 1973. Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment. *J. ACM* 20, 1 (1973), 46–61.

[19] J. M. Lopez, J. L. Diaz, and D. F. Garcia. 2004. Utilization Bounds for EDF Scheduling on Real-Time Multiprocessor Systems. *Real-Time Systems: The International Journal of Time-Critical Computing* 28, 1 (2004), 39–68.

[20] R McNaughton. 1959. Scheduling with Deadlines and Loss Functions. *Management Science* 6 (1959), 1–12.

[21] James Orr. 2019. Multiprocessor Elastic Scheduling Simulation Data. https://osf.io/395bk/

[22] James Orr, Chris Gill, Kunal Agrawal, Sanjoy Baruah, Christian Cianfarani, Phyllis Ang, and Christopher Wong. 2018. Elasticity of Workloads and Periods of Parallel Real-Time Tasks. In *Proceedings of the 26th International Conference on Real-Time Networks and Systems (RTNS '18).* ACM, New York, NY, USA, 61–71. https://doi.org/10.1145/3273905.3273915

[23] James Orr, Chris Gill, Kunal Agrawal, Jing Li, and Sanjoy Baruah. 2019. Elastic Scheduling for Parallel Real-Time Systems. *Leibniz Transactions on Embedded Systems* 6, 1 (2019), 05–1–05:14. https://doi.org/10.4230/LITES-v006-i001-a005