A Note on the EDF Preemption Behavior in "Rate Monotonic vs. EDF: Judgment Day"

Adrien Brun, Chunhui Guo, *Student Member, IEEE*, and Shangping Ren, *Senior Member, IEEE*

***Abstract*—In [1], the author empirically compared EDF (Earliest Deadline First) and RM (Rate Monotonic) scheduling algorithms and made a few EDF preemption behavior observations based on data obtained from the first 1,000 time units of scheduling activities. However, based on test settings given in [1], the first 1,000 time units occupies only a small percentage of the entire task set's hyper-period. We extend EDF preemption behavior study by extending scheduling activities from the first small percentage of a hyper-period of a given task to the entire hyper-period. The extended experiments indicate that the number of preemptions occurred at the beginning of a task set's hyper-period does not necessarily represent the trend for the entire hyper-period. Hence, comparisons and conclusions made based on a small percentage of a scheduling interval over a task set's hyper-period may not be accurate. Second, the total number of preemptions within a task set's hyper-period does not decrease when the task set total utilization increases which is different from the observation obtained in [1]. We also investigate the impact of execution time differences among tasks on the preemption behavior.**

***Index Terms*—Real-Time Scheduling; Earliest Deadline First; Preemption**

I. Introduction

In [1], Buttazzo empirically studied the relationship between the number of preemptions and task set utilization. His experiment settings are: (1) a task set with 10 tasks; (2) each task's period is randomly selected within $[10, 100]$; (3) utilization of the task set ranges from 0.5 to 0.95 with a step size of 0.05; and (4) the system executes for 1000 time units and the total number of preemptions within $[1, 1000]$ is counted. The results are: when task set utilization increases from 0.5 to 0.85, the total number of preemptions increases; while when the utilization increases from 0.85 to 0.95, the total number of preemptions decreases. Buttazzo has used a concrete example to explain why the latter case might happen. In particular, for a periodic task set $\Gamma_1 = \{(6,1), (10,3), (16,4)\}$ and $\Gamma_2 = \{(6,1), (10,3), (16,8)\}$, where $U(\Gamma_2) = 0.97 > U(\Gamma_1) = 0.72$, there are four preemptions in Γ_1 during the first 20 time units, but only three preemptions in Γ_2. Each periodic task τ is represented by a 2-tuple (T, C), where T is the period and C is the worst-case execution time (WCET) [2].

We argue that the preemption behavior of a given task set within different time intervals of a fixed size does not represent the preemption behavior within the task set's hyper-period. Hence, the observation obtained within the task set's first 1,000 time units does not represent a general case about EDF preemption behavior.

The authors are with Computer Science Department, Illinois Institute of Technology, Chicago, IL 60616, USA

Emails: {abrun, cguo13}@hawk.iit.edu, ren@iit.edu

The research is supported in part by NSF under grant number CAREER 0746643, CNS 1018731, and CNS 1035894.

II. Preemption within Hyper-Period

A. Utilization Impact

We take the same two task sets given in Section I from [1] and rather than using EDF to schedule the two task sets for the first 20 time units, we execute the task set for the entire hyper-period, i.e., 240 time units. Fig. 1 shows the number of preemptions of the two task sets in each ten percent of their hyper-period time intervals.

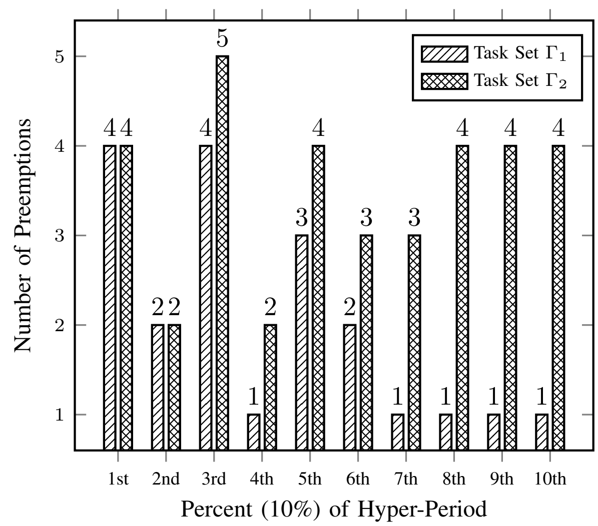

Fig. 1. Number of Preemptions within Each Ten Percent (10%) Hyper-Period Intervals

From Fig. 1, we have the following observations:

Obs 1: The number of preemptions within different time intervals of the task set's hyper-period is different.

Obs 2: The highest number of preemptions **does not** always occur in the interval that contains critical instants, i.e., the instant when all higher priority tasks are released at the same time.

Obs 3: The preemption behavior during the first ten percent of the hyper-period **does not** reflect the preemption behavior during the entire hyper-period.

Obs 4: The number of preemptions **does not** decrease when task set utilization increases.

To investigate if these observations obtained from an individual example represent a general case, we design an experiment with the following settings[1]: (1) task number: 10; (2) task period range: [10, 250]; (3) total task utilization range: [0.3 1.0] with step 0.05; and (4) testing time units: 510,510. Task periods are products of two different, randomly selected prime numbers from $\{2, 3, 5, 7, 11, 13, 17\}$. To have a meaningful comparison, we ensure that all generated task sets have the same hyper-period of 510,510, which is the product of the given seven prime numbers. We also ignore periods not in the range [10, 250]. For each utilization, similar to [1], we use UUniform [3] to uniformly distribute the utilization among ten tasks. We repeat the experiment 1,000 times and obtain the average number of preemptions divided by the average number of total preemptions for the entire hyper-period, defined as the Normalized Average Number of Preemptions. Fig. 2, Fig. 3, and Fig. 4 depict the results.

To investigate the EDF preemption behavior when the utilization among tasks is skewed instead of uniformly distributed, we use UFitting, whose major characteristic is asymmetrical on task utilization distribution [3], to randomly distribute task utilizations, and repeat the above experiments. The results are also shown in Fig. 4.

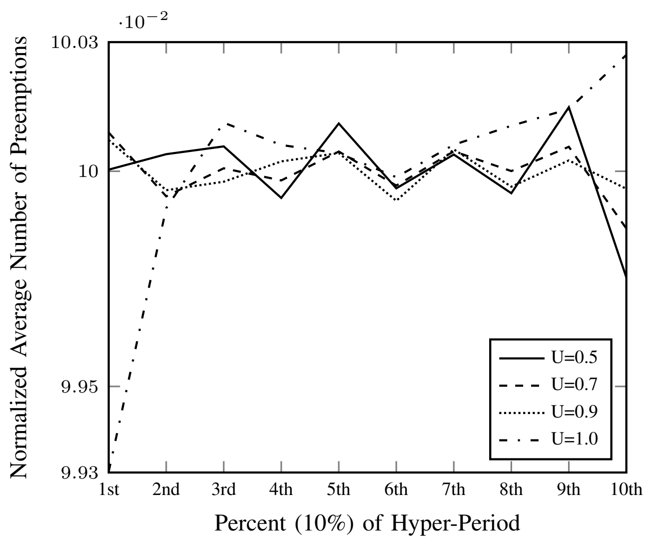

Fig. 2. Normalized Average Number of Preemptions within Each Ten Percent (10%) of Task Set Hyper-Period

From Fig. 2, Fig. 3, and Fig. 4, we can make an empirically-based conclusion, i.e., in general, observations **Obs 1** to **Obs 4** obtained from a single case is *true*. **Obs 4** is different from the conclusion made in [1]. Fig. 4 also shows that the task utilization distribution does not impact the general EDF preemption behavior.

B. Task Number Impact

To evaluate the task number impact on the EDF preemption behavior, we design an experiment with the same experiment settings in Section II-A except: (1) the total task utilization is fixed to be 0.9; and (2) the task number varies in range [4, 14]. We also use both UUniform and UFitting to distribute task utilizations and repeat the experiment 1,000 times. Fig. 5 depicts the results.

[1]The test settings are intended to match Buttazzo's test settings as much as possible.

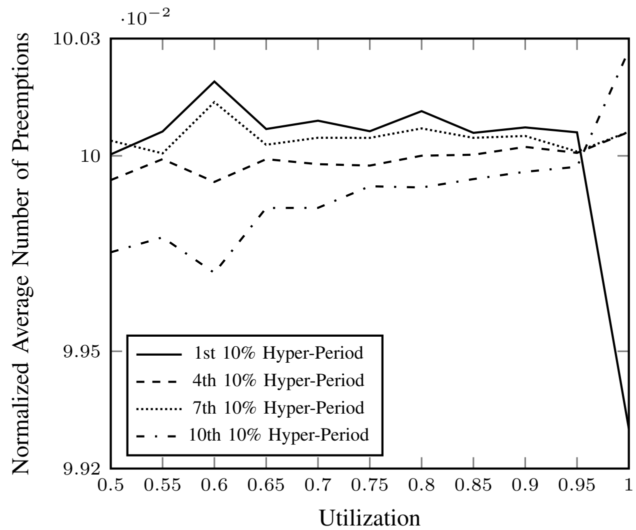

Fig. 3. Normalized Average Number of Preemptions with Different Task Set Utilizations

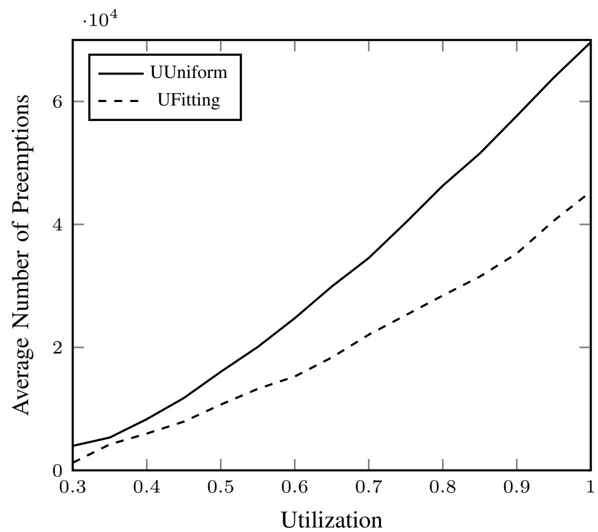

Fig. 4. Average Total Number of Preemptions in the Hyper-Period with Fixed Number of Tasks

From Fig. 5, we can make similar conclusions in [1], i.e., the number of preemptions increases for small task sets and decreases for larger task sets. Fig. 5 also shows similar results in Fig. 4 that the task utilization distribution does not impact the general EDF preemption behavior.

III. What May Impact Total Number of Preemptions

Consider the three periodic task sets shown in TABLE I, they have the same task periods and $U(\Gamma_1) = 0.52 < U(\Gamma_2) = 0.56 < U(\Gamma_3) = 0.66$. When each task set is executed for its entire hyper-period, i.e., 180 time units, the total number of preemptions are 8, 3, and 10 for task set Γ_1, Γ_2, and Γ_3, respectively. This example reveals that although the general trend is that the larger the utilization, the higher the total number of preemptions within the task set's hyper-period, for individual cases, increased utilization does

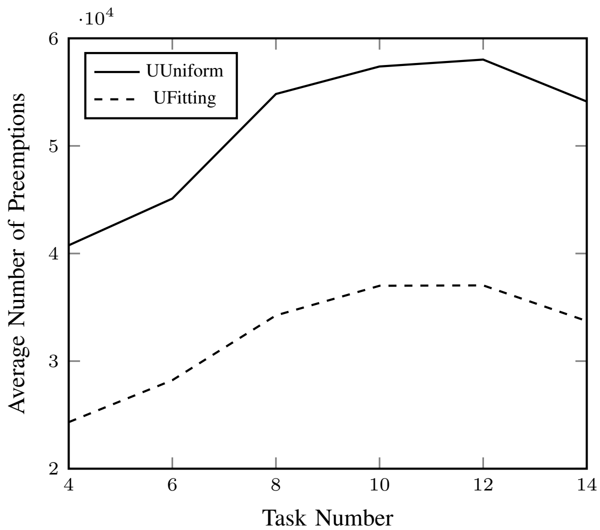

Fig. 5. Average Total Number of Preemptions in the Hyper-Period with Fixed Utilization

not necessarily result in increased number of preemptions or decreased number of preemptions.

TABLE I
TASK SET INFORMATION

Task Set	Period	WCET	Utilization	# Preemption
Γ_1	4	1	0.52	8
	10	1		
	18	3		
Γ_2	4	1	0.56	3
	10	2		
	18	2		
Γ_3	4	1	0.66	10
	10	3		
	18	2		

Taking a closer look at the differences among the three task sets, the execution time differences between high priority and low priority tasks seem to play an important role with respect to the total number of preemptions within a hyper-period. For instance, the three task sets have execution times of $\{1,1,3\}$, $\{1,2,2\}$, and $\{1,3,2\}$. The differences among the task execution times within each task set seem to be correlated with the number of preemptions incurred with EDF scheduling.

To evaluate the execution time difference impact on the preemption behavior by experiments, we first define the execution time difference (ETD) of a task set as follows.

Definition 1: Given a task set $\Gamma = \{\tau_1, \tau_2, \ldots, \tau_n\}$, where $\tau_i = (T_i, C_i)$, T_i is the period, and C_i is the worst-case execution time, the execution time difference (ETD) of Γ is defined as the average value of the absolute execution time differences between each task and all its higher priority tasks, i.e.,

$$\mathtt{ETD} = \frac{\sum_{i=1}^{n} \sum_{j \in \mathtt{hp}(i)} |C_i - C_j|}{\binom{n}{2}} \tag{1}$$

where $\mathtt{hp}(i)$ is the set of tasks with higher priority than τ_i. $\square$

To evaluate the execution time difference impact on the EDF preemption behavior, we design an experiment with the same experiment settings in Section II-A except that we fix the the total task utilization to be 0.8. In the experiment, we randomly generate 100 task periods sets (each set contains 10 task periods). For each task periods set, we fix each task's period and adjust its execution time to generate 100 task sets with fixed utilization 0.8. We run these 100×100 test cases and combine the results based on execution time differences (defined by Eq. (1)) rounded to the nearest tenth.

Fig. 6 depicts the average number of preemptions under different execution time differences. As shown in Fig. 6, in general, when the execution time difference increases, the number of preemptions also increases. When the execution time difference is above 10, variations increase.

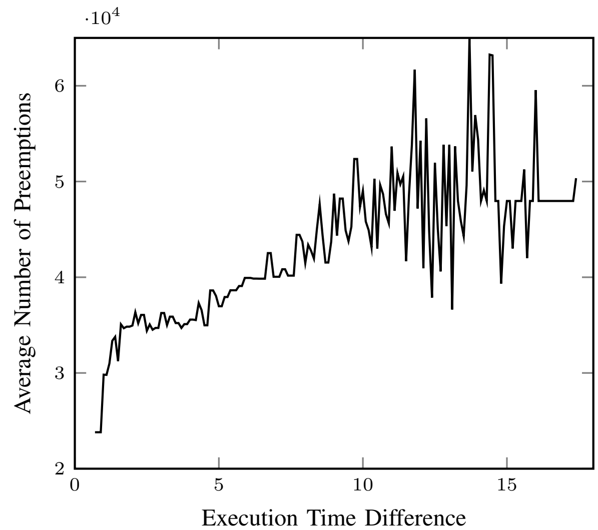

Fig. 6. Impact of Execution Time Difference on Preemption Behavior

IV. CONCLUSION

We extend the empirical study on EDF preemptions initiated in [1] and find that the number of preemptions that occur at the beginning of a task set's hyper-period does not represent the preemption behavior during the entire hyper-period. Hence, the conclusions based on a partial hyper-period may not be accurate or may not represent a general case. Though the average case in a large number of tests shows that the larger the task set utilization, the higher the number of preemptions, for individual cases, it's possible that increased utilization can result in either an increase or decrease in the number of preemptions. Further, we observe that the execution time differences among tasks seem to have some correlation to the total number of preemptions. We also use experiments to verify the observation of the execution time difference impact on the preemption behavior.

REFERENCES

[1] Giorgio C. Buttazzo. Rate monotonic vs. edf: Judgment day. *Real-Time Systems*, 29(1):5–26, January 2005.

[2] C. L. Liu and James W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1):46–61, jan 1973.

[3] Enrico Bini and Giorgio C. Buttazzo. Measuring the performance of schedulability tests. *Real-Time Systems*, 30(1-2):129–154, May 2005.