

Lecture 7

EARLIEST DEADLINE FIRST

The Earliest Deadline First (EDF) algorithm is a dynamic scheduling rule that selects tasks according to their absolute deadlines. Specifically, tasks with earlier deadlines will be executed at higher priorities. Since the absolute deadline of a periodic task depends on the current j th instance as

$$d_{i,j} = \Phi_i + (j - 1)T_i + D_i,$$

EDF is a dynamic priority assignment. Moreover, it is typically executed in preemptive mode, thus the currently executing task is preempted whenever another periodic instance with earlier deadline becomes active.

Note that EDF does not make any specific assumption on the periodicity of the tasks; hence, it can be used for scheduling periodic as well as aperiodic tasks.

A set of periodic tasks is schedulable with EDF if and only if

$$\sum (C_i/P_i) \leq 1 \text{ for all task } i$$

Example:

Consider the following periodic task set

Task	C_i or E_{x_i}	Period
T1	2	5
T2	4	7


Processor utilization factor is

$$U = 2/5 + 4/7 = 34/35 \sim 0.97$$

This means that 97 percent of the processor time is used to execute the periodic tasks, whereas the CPU is idle in the remaining 3 percent. Being $U > 2(\sqrt{2}-1) \sim 0.83$, the schedulability of the task set cannot be guaranteed under RM, whereas it is guaranteed under EDF. Indeed, as shown in Figure below, RM generates a deadline miss at time $t = 7$, whereas EDF completes all tasks within their deadlines.

Another important difference between RM and EDF concerns the number of preemptions occurring in the schedule. As shown in Figure 4.13, under RM every instance of task τ_2 is preempted, for a total number of five preemptions in the interval $T = T_1 T_2$. Under EDF, the same task is preempted only once

in the same interval. The smaller number of preemptions in EDF is a direct consequence of the dynamic priority assignment, which at any instant privileges the task with the earliest deadline, independently of tasks periods.



(a)



(b)

پابلسین حمو