

# **Real-Time Scheduling**

Scheduling of Reactive Systems

Priority-Driven Scheduling

[Some parts of this lecture are based on a real-time systems course  
of Colin Perkins

<http://csperkins.org/teaching/rtes/index.html>]

# Current Assumptions

- ▶ Single processor
- ▶ Fixed number,  $n$ , of *independent periodic* tasks  
i.e. there is no dependency relation among jobs
  - ▶ Jobs can be preempted at any time and never suspend themselves
  - ▶ No aperiodic and sporadic jobs
  - ▶ No resource contentions

Moreover, unless otherwise stated, we assume that

- ▶ **Scheduling decisions take place precisely at**
  - ▶ release of a job
  - ▶ completion of a job

(and nowhere else)

- ▶ Context switch overhead is negligibly small  
i.e. assumed to be zero
- ▶ There is an unlimited number of priority levels

# Fixed-Priority vs Dynamic-Priority Algorithms

A priority-driven scheduler is on-line

i.e. it does not precompute a schedule of the tasks

- ▶ It assigns priorities to jobs after they are released and places the jobs in a ready job queue in the priority order with the highest priority jobs at the head of the queue
- ▶ At each scheduling decision time, the scheduler updates the ready job queue and then schedules and executes the job at the head of the queue  
i.e. one of the jobs with the highest priority

**Fixed-priority** = all jobs in a task are assigned the same priority

**Dynamic-priority** = jobs in a task may be assigned different priorities

# Priority-Driven Algorithms

## Dynamic-priority:

- ▶ **EDF** = at the time of a scheduling decision, the job queue is ordered by the earliest deadline

## Fixed-priority:

- ▶ **RM** = assigns priorities to tasks based on their periods
- ▶ **DM** = assigns priorities to tasks based on their relative deadlines

(In all cases, ties are broken arbitrarily.)

We consider the following questions:

- ▶ Are the algorithms optimal?
- ▶ How to efficiently (or even online) test for schedulability?

To measure abilities of scheduling algorithms and to get fast online tests of schedulability we use a notion of **utilization**

# Utilization

- ▶ *Utilization  $u_i$  of a periodic task  $T_i$*  with period  $p_i$  and execution time  $e_i$  is defined by  $u_i := e_i/p_i$   
 $u_i$  is the fraction of time a periodic task with period  $p_i$  and execution time  $e_i$  keeps a processor busy
- ▶ *Total utilization  $U^{\mathcal{T}}$  of a set of tasks  $\mathcal{T} = \{T_1, \dots, T_n\}$*  is defined as the sum of utilizations of all tasks of  $\mathcal{T}$ , i.e. by

$$U^{\mathcal{T}} := \sum_{i=1}^n u_i$$

- ▶  $U$  is a *schedulable utilization* of an algorithm ALG if all sets of tasks  $\mathcal{T}$  satisfying  $U^{\mathcal{T}} \leq U$  are schedulable by ALG.  
*Maximum schedulable utilization  $U_{\text{ALG}}$  of an algorithm ALG* is the *supremum of schedulable utilizations of ALG*.
  - ▶ If  $U^{\mathcal{T}} < U_{\text{ALG}}$ , then  $\mathcal{T}$  is schedulable by ALG.
  - ▶ If  $U > U_{\text{ALG}}$ , then there is  $\mathcal{T}$  with  $U^{\mathcal{T}} \leq U$  that is not schedulable by ALG.

## Utilization – Example

- ▶  $T_1 = (2, 1)$  then  $u_1 = \frac{1}{2}$
- ▶  $T_1 = (11, 5, 2, 4)$  then  $u_1 = \frac{2}{5}$   
(i.e., the phase and deadline do not play any role)
- ▶  $\mathcal{T} = \{T_1, T_2, T_3\}$  where  $T_1 = (2, 1)$ ,  $T_2 = (6, 1)$ ,  $T_3 = (8, 3)$   
then

$$U^{\mathcal{T}} = \frac{1}{2} + \frac{1}{6} + \frac{3}{8} = \frac{25}{24}$$

# **Real-Time Scheduling**

Priority-Driven Scheduling

Dynamic-Priority

# Optimality of EDF

## Theorem 1

Let  $\mathcal{T} = \{T_1, \dots, T_n\}$  be a set of independent, preemptable periodic tasks with  $D_i \geq p_i$  for  $i = 1, \dots, n$ . The following statements are equivalent:

1.  $\mathcal{T}$  can be feasibly scheduled on one processor
2.  $U^{\mathcal{T}} \leq 1$
3.  $\mathcal{T}$  is schedulable using EDF

(i.e., in particular,  $U_{EDF} = 1$ )

## Proof.

1. $\Rightarrow$ 2. We prove that  $U^{\mathcal{T}} > 1$  implies that  $\mathcal{T}$  is not schedulable (whiteb.)
2. $\Rightarrow$ 3. Next slides and whiteboard ...
3. $\Rightarrow$ 1. Trivial





## Proof of 2. $\Rightarrow$ 3.

**Notation:** Given a set of tasks  $\mathcal{L}$ , we denote by  $\cup \mathcal{L}$  the set of all jobs of the tasks in  $\mathcal{L}$ .

We prove  $\neg 3. \Rightarrow \neg 2.$  assuming that  $D_i = p_i$  for  $i = 1, \dots, n$  (note that the general case immediately follows).

Assume that  $\mathcal{T}$  is not schedulable by EDF. We show that  $U^{\mathcal{T}} > 1$ .

Suppose that a job  $J_{i,k}$  of  $T_i$  misses its deadline at time  $t = r_{i,k} + p_i$ .

Let  $\mathcal{T}'$  be the set of all tasks whose jobs are released in  $[r_{i,k}, t]$  (i.e., a task belongs to  $\mathcal{T}'$  iff at least one job of the task is released in  $[r_{i,k}, t]$ ).

Let  $t_-$  be the end of the *latest* interval before  $t$  in which either jobs of  $\cup(\mathcal{T} \setminus \mathcal{T}')$  are executed, or the processor is idle.

Then  $r_{i,k} \geq t_-$  since all jobs of  $\cup(\mathcal{T} \setminus \mathcal{T}')$  waiting for execution during  $[r_{i,k}, t]$  have deadlines later than  $t$  (thus have lower priorities than  $J_{i,k}$ ).

## Proof of 2. $\Rightarrow$ 3. (cont.)

It follows that

- ▶ no job of  $\cup(\mathcal{T} \setminus \mathcal{T}')$  is executed in  $[t_-, t]$ ,  
(by definition of  $t_-$ )
- ▶ all jobs of  $\cup \mathcal{T}'$  executed in  $[t_-, t]$  are released in  $[t_-, t]$  and have their deadlines in  $[t_-, t]$ ,  
(since no job of  $\cup \mathcal{T}'$  executes just before  $t_-$ , and jobs with deadlines after  $t$  have lower priorities than  $J_{i,k}$ )
- ▶ the processor is fully utilized in  $[t_-, t]$ .  
(by definition of  $t_-$ )

Let  $G$  be the set of all jobs that are released in  $[t_-, t]$  and have their deadlines in  $[t_-, t]$ .

Note that  $J_{i,k} \in G$  since  $r_{i,k} \geq t_-$ . Since  $J_{i,k}$  misses its deadline, the processor should not be able to complete all jobs of  $G$  before  $t$ .

Denote by  $E_G$  the sum of all execution times of all jobs in  $G$  (the total execution time of  $G$ ).

It follows that  $E_G > t - t_-$  because otherwise, all jobs of  $G$  (in particular,  $J_{i,k}$ ) would complete in  $[t_-, t]$ .

## Proof of 2. $\Rightarrow$ 3. (cont.)

How to compute  $E_G$ ?

For  $T_\ell \in \mathcal{T}'$ , denote by  $R_\ell$  the earliest release time of a job in  $T_\ell$  during the interval  $[t_-, t]$ .

For every  $T_\ell \in \mathcal{T}'$ , exactly  $\left\lfloor \frac{t-R_\ell}{p_\ell} \right\rfloor$  jobs of  $T_\ell$  belong to  $G$ . (For every  $T_\ell \in \mathcal{T} \setminus \mathcal{T}'$ , exactly 0 jobs belong to  $G$ .)

Thus

$$E_G = \sum_{T_\ell \in \mathcal{T}'} \left\lfloor \frac{t-R_\ell}{p_\ell} \right\rfloor e_\ell$$

As argued above:

$$t-t_- < E_G = \sum_{T_\ell \in \mathcal{T}'} \left\lfloor \frac{t-R_\ell}{p_\ell} \right\rfloor e_\ell \leq \sum_{T_\ell \in \mathcal{T}'} \frac{t-t_-}{p_\ell} e_\ell \leq (t-t_-) \sum_{T_\ell \in \mathcal{T}'} u_\ell \leq (t-t_-) U^{\mathcal{T}}$$

which implies that  $U^{\mathcal{T}} > 1$ .

# Density and EDF

What about tasks with  $D_i < p_i$  ?

*Density of a task  $T_i$*  with period  $p_i$ , execution time  $e_i$  and relative deadline  $D_i$  is defined by

$$e_i / \min(D_i, p_i)$$

*Total density  $\Delta^{\mathcal{T}}$  of a set of tasks  $\mathcal{T}$*  is the sum of densities of tasks in  $\mathcal{T}$

Note that if  $D_i < p_i$  for some  $i$ , then  $\Delta^{\mathcal{T}} > U^{\mathcal{T}}$

## Theorem 2

*A set  $\mathcal{T}$  of independent, preemptable, periodic tasks can be feasibly scheduled on one processor if  $\Delta^{\mathcal{T}} \leq 1$ .*

Note that this is NOT a necessary condition! (Example whiteb.)

# Schedulability Test For EDF

**The problem:** Given a set of independent, preemptable, periodic tasks  $\mathcal{T} = \{T_1, \dots, T_n\}$  where each  $T_i$  has a period  $p_i$ , execution time  $e_i$ , and relative deadline  $D_i$ , decide whether  $\mathcal{T}$  is schedulable by EDF.

**Solution using utilization and density:**

If  $p_i \leq D_i$  for each  $i$ , then it suffices to decide whether  $U^{\mathcal{T}} \leq 1$ .

Otherwise, decide whether  $\Delta^{\mathcal{T}} \leq 1$ :

- ▶ If yes, then  $\mathcal{T}$  is schedulable with EDF
- ▶ If not, then  $\mathcal{T}$  does not have to be schedulable

Note that

- ▶ Phases of tasks do not have to be specified
- ▶ Parameters may vary: increasing periods or deadlines, or decreasing execution times does not prevent schedulability

# Schedulability Test for EDF – Example

Consider a digital robot controller

- ▶ A control-law computation
  - ▶ takes no more than 8 ms
  - ▶ the sampling rate: 100 Hz, i.e. computes every 10 ms

Feasible? Trivially yes ....

- ▶ Add Built-In Self-Test (BIST)
  - ▶ maximum execution time 50 ms
  - ▶ want a minimal period that is feasible (max one second)

With 250 ms still feasible ....

- ▶ Add a telemetry task
  - ▶ maximum execution time 15 ms
  - ▶ want to minimize the deadline on telemetry  
period may be large

Reducing BIST to once a second, deadline on telemetry  
may be set to 100 ms ....