




REVIEW: Aperiodic scheduling



- **Given:**
 - A set of non-periodic tasks $\{J_1, \dots, J_n\}$ with
 - arrival times a_i , deadlines d_i , computation times C_i
 - precedence constraints
 - resource constraints
 - Class of scheduling algorithm:
 - Preemptive, non-preemptive
 - Off-line / on-line
 - Optimal / heuristic
 - One processor / multi-processor
 - ...
 - Cost function:
 - Minimize maximum lateness
 - ...
- **Find:**
 - Feasible schedule
 - Optimal schedule according to given cost function

REVIEW: EDD – Earliest Due Date

EDD: execute the tasks in **order of non-decreasing deadlines**

- **Lemma:**

If arrival times are **synchronous**, then preemption does not help, i.e. if there is a preemptive schedule with maximum lateness L_{\max} , then there is also a non-preemptive schedule with maximum lateness L_{\max} .

- **Theorem (Jackson '55):**

Given a set of n independent tasks with **synchronous** arrival times, any algorithm that executes the tasks in **order of non-decreasing deadlines** is **optimal with respect to minimizing the maximum lateness**.

REVIEW: EDF – Earliest Deadline First

- **EDF:** At every instant execute the task with the earliest absolute deadline among all the ready tasks.

- **Theorem (Horn '74):**

Given a set of n independent task **with arbitrary arrival times**, any algorithm that at every instant executes the task with the earliest absolute deadline among all the ready tasks is optimal with respect to minimizing the maximum lateness.



REVIEW: Non-preemptive version

- **Theorem** (Jeffay et al. '91): EDF is an optimal **non-idle** scheduling algorithm also in a **non-preemptive** task model.
- Non-preemptive scheduling with **idle schedules allowed** is **NP-hard**
- Possible approaches:
 - Heuristics
 - Bratley's algorithm: Branch-and-bound

REVIEW: Scheduling with precedence constraints

- Non-preemptive scheduling with non-synchronous arrival times, deadlines and precedence constraints is **NP-hard**.
- LDF for **synchronous** arrival times (all tasks arrive at 0)

REVIEW: LDF – Latest Deadline First



BF - ES


- 7 -

LDF

Theorem (Lawler 73):

LDF is optimal wrt. maximum lateness.

- Task set $J = \{J_1, \dots, J_n\}$
- $T \subseteq J$ subset without successors
- $J_e \in T$ with latest deadline



We show that we can move J_e to the end with

- 1) no violation of the precedences
- 2) with no increase in max. lateness

BF - ES

- 8 -

1) precedence is not violated
(J_e does not have successors)

$$2) L'_{\max} = \max \{ L'_{\max}(A), L'_{\max}(B), L'_k, L'_e \}$$

$L'_{\max}(A) = L_{\max}(A)$ nothing changed

$L'_{\max}(B) \leq L_{\max}(B)$ starts & ends earlier

$L'_k < L_k$ starts & ends earlier

$$L'_e = \sum_{i \in h} C_i - d_e < \sum_{i \in h} C_i - d_k = L_k$$

Remove J_e and continue.

Preemptive

- Non-preemptive scheduling with non-synchronous arrival times, deadlines and precedence constraints is **NP-hard**.
- Modified EDF for **preemptive** scheduling, **arbitrary arrival times**

EDF with precedence constraints

1. Modify arrival times

- For any **initial** node J_i of the precedence graph, set $a_i^* := a_i$.
- For any task J_i such that all predecessors have been processed, set $a_i^* := \max \{a_i, a_h^* + C_h \mid J_h \rightarrow J_i\}$

2. Modify deadlines

- For any **terminal** node J_i of the precedence graph, set $d_i^* := d_i$.
- For any task J_i such that all successors have been processed, set $d_i^* := \min \{d_i, d_h^* - C_h \mid J_i \rightarrow J_h\}$



($J_h \rightarrow J_i$: J_h is a direct predecessor of J_i)

BF - ES

- 11 -

Example

	J_1	J_2	J_3	J_4	J_5	J_6
A_i	1	2	3	3	3	4
C_i	1	1	1	1	1	1
d_i	2	3	5	7	4	6



BF - ES

- 12 -

EDF with precedence constraints

Theorem: The given task set is schedulable such that the precedence constraints are met if and only if the modified task set is schedulable under EDF.

" \Rightarrow " In any feasible schedule that meets the precedence constraints, we have that

$$s_i \geq \max \{a_i, a_h^* + C_h \mid j_h \rightarrow j_i\}$$

and

$$f_i \leq \min \{d_i, d_h^* - C_h \mid j_i \rightarrow j_h\}$$

$\Rightarrow \sigma$ is feasible for the modified task set.

" \Leftarrow " For any $j_i < j_j$, we have that

$$a_i^* < a_j^* \quad \text{and} \quad d_i^* < d_j^*$$

When j_j arrives then j_i and j_i has the earlier deadline has already arrived

$\Rightarrow j_j$ cannot start before j_i
 j_i cannot preempt j_j

\Rightarrow precedence constraints are satisfied.


$a_i^* \geq a_i, d_i^* \leq d_i \Rightarrow$ timing constraints satisfied.

Optimal scheduling algorithms for *periodic* tasks

BF - ES

- 15 -

Periodic scheduling



▪ **Given:**

- A set of periodic tasks $\Gamma = \{\tau_1, \dots, \tau_n\}$ with
 - phases Φ_i (arrival times of first instances of tasks),
 - periods T_i (time difference between two consecutive activations)
 - relative deadlines D_i (deadline relative to arrival times of instances)
 - computation times C_i

$\Rightarrow j$ th instance $\tau_{i,j}$ of task τ_i with

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

▪ **Find a feasible schedule**

- start time $s_{i,j}$ and
- finishing time $f_{i,j}$

BF - ES

- 16 -

Assumptions

- A.1. Instances of periodic task τ_i are regularly activated with constant period T_i .
- A.2. All instances have same worst case execution time C_i .
- A.3. All instances have same relative deadline D_i , here in most cases equal to T_i (i.e., $d_{i,j} = \Phi_i + j \cdot T_i$)
- A.4. All tasks in Γ are independent.
- A.5. Overhead for context switches is neglected, i.e. assumed to be 0 in the theory.


- Basic results based on these assumptions form the core of scheduling theory.
- For practical applications, assumptions A.3. and A.4. can be relaxed, but results have to be extended.

BF - ES

- 17 -

Examples for periodic scheduling (1)

	τ_1	τ_2
Φ_i	0	0
T_i	2	4
C_i	1	2
D_i	1	4




- Schedulable, but only preemptive schedule possible.

BF - ES

- 18 -

Examples for periodic scheduling (2)

	τ_1	τ_2
Φ_i	0	0
T_i	2	4
C_i	1	2
D_i	2	4



- Schedulable with non-preemptive schedule.

BF - ES

- 19 -

Examples for periodic scheduling (3)

	τ_1	τ_2
Φ_i	0	0
T_i	3	4
C_i	2	2
D_i	3	4

$$T_1 \cdot T_2 = 12$$

Within 12 units:

$$\frac{12}{3} = 4 \text{ executions of } T_1$$

$$\frac{12}{4} = 3 \text{ executions of } T_2$$

- No feasible schedule for single processor.

$$\begin{array}{r}
 4 \times 2 = 8 \text{ units by } T_1 \\
 3 \times 2 = 6 \text{ units by } T_2 \\
 \hline
 14 \text{ units capacity within} \\
 12 \text{ units impossible!}
 \end{array}$$

BF - ES

- 20 -

Processor utilization

Definition:

Given a set Γ of n periodic tasks, the **processor utilization U** is given by

$$U = \sum_{i=1}^n \frac{C_i}{T_i}$$

$$U = \frac{2}{3} + \frac{2}{4} = \frac{14}{12} > 1$$

Processor utilization as a schedulability criterion

- Given: a scheduling algorithm A
- Define $U_{\text{bnd}}(A) = \inf \{U(\Gamma) \mid \Gamma \text{ is not schedulable by algorithm } A\}$.
- If $U_{\text{bnd}}(A) > 0$ then a **simple, sufficient criterion for schedulability by A can be based on processor utilization**:
 - If $U(\Gamma) < U_{\text{bnd}}(A)$ then Γ is schedulable by A .
 - However, if $U_{\text{bnd}}(A) < U(\Gamma) \leq 1$, then Γ may or may not be schedulable by A .
- **Question:**
Does a scheduling algorithm A exist with $U_{\text{bnd}}(A) = 1$?

Processor utilization

- **Question:**
Does a scheduling algorithm A exist with $U_{\text{bnd}}(A) = 1$?

- **Answer:**
 - No, if $D_i < T_i$ allowed.
 - Example:

	τ_1	τ_2
Φ_i	0	0
T_i	2	2
C_i	1	1
D_i	1	1

- Yes, if $D_i = T_i$ (or $D_i \geq T_i$)) **Earliest Deadline First (EDF)**
- In the following: assume $D_i = T_i$

BF - ES

- 23 -

Earliest Deadline First (EDF)

- EDF is applicable to both periodic and aperiodic tasks.
- If there are only periodic tasks, priority-based schemes like “rate monotonic scheduling (RM)” (see later) are often preferred, since
 - They are simpler due to fixed priorities
⇒ use in “standard OS” possible
 - sorting wrt. to deadlines **at run time** is not needed

BF - ES

- 24 -

EDF and processor utilization factor

- Theorem:** A set of periodic tasks τ_1, \dots, τ_n with $D_i = T_i$ is schedulable with EDF iff $U \leq 1$.

" \Rightarrow " . Let $T = T_1 \cdot \dots \cdot T_n$

$$\cdot \sum_{i=1}^n \frac{T}{T_i} \cdot C_i \quad \text{time taken by task set within } T$$

$$= \sum_{i=1}^n \frac{C_i}{T_i} \cdot T = U \cdot T$$

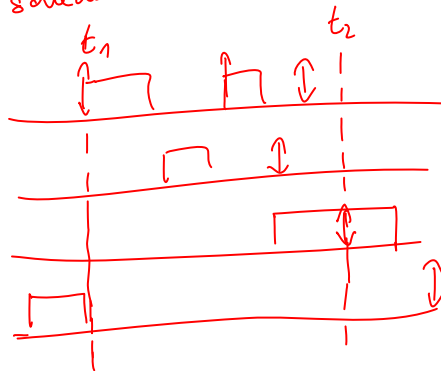
Assume $U > 1$. Then $UT > T$ and task set is not schedulable.

BF - ES

- 25 -

" \Leftarrow " , Assume $U \leq 1$ and the task set is not EDF-schedulable.

- Let t_2 be the earliest time in the EDF-schedule when a task misses its deadline



Let $[t_1, t_2]$ be the largest interval s.t.

- no idle times
- only instances with deadline $\leq t_2$ are executed.

BF - ES

- 26 -

Claim: The tasks executed in $[t_1, t_2]$ have arrival times $\geq t_1$

Case 1: The processor was idle directly before t_1
 \Rightarrow No unfinished tasks with arrival $< t_1$

Case 2: The task running directly before t_1 has a deadline $\leq t_2$
 \Rightarrow Contradiction to maximality of $[t_1, t_2]$

Case 3: The task running directly before t_1 has a deadline $> t_2$
Due to EDF, no task with arrival $< t_1$ and deadline $\leq t_2$ left

BF - ES

- 27 -

Since all tasks in $[t_1, t_2]$ have deadline $\leq t_2$
 \Rightarrow all tasks in $[t_1, t_2]$ have arrival $\geq t_1$.

Time overflow at t_2 :

$$(t_2 - t_1) < \sum_{\substack{a_{i,j} \geq t_1 \\ d_{i,j} \leq t_2}} C_i$$

$$\begin{aligned} &= \sum_{i=1}^n \left\lfloor \frac{t_2 - t_1}{T_i} \right\rfloor \cdot C_i \\ &\leq \sum_{i=1}^n \frac{t_2 - t_1}{T_i} \cdot C_i = (t_2 - t_1) \cdot \sum_{i=1}^n \frac{C_i}{T_i} \\ &= (t_2 - t_1) \cdot u \end{aligned}$$

$$\Rightarrow u > 1$$

BF - ES

- 28 -

Rate monotonic scheduling (RM)

- Rate monotonic scheduling (RM) (Liu, Layland '73):
 - Assign **fixed priorities** to tasks τ_i :
 - $\text{priority}(\tau_i) = 1/T_i$
 - I.e., priority **reflects release rate**
 - **Always execute ready task with highest priority**
 - Preemptive: currently executing task is preempted by newly arrived task with shorter period.


BF - ES

- 29 -

Example for RM (1)

	τ_1	τ_2	τ_3
Φ_i	0	0	0
T_i	4	6	12
C_i	2	1	4
D_i	4	6	12

priority (τ_1) > priority (τ_2)
> priority (τ_3)




BF - ES

- 30 -

Example for RM (2)

	τ_1	τ_2	τ_3
Φ_i	0	0	0
T_i	4	5	10
C_i	2	2	1
D_i	4	5	10

$$U = \frac{2}{4} + \frac{2}{5} + \frac{1}{10} = 1$$



BF - ES

- 31 -

Optimality of Rate Monotonic Scheduling

- **Theorem (Liu, Layland, 1973):**
RM is **optimal among all fixed-priority** scheduling algorithms.
- **Def.:** The **response time** $R_{i,j}$ of an instance j of task i is the time (measured from the arrival time) at which the instance is finished: $R_{i,j} = f_{i,j} - a_{i,j}$.
- The critical instant of a task is the time at which the arrival of the task will produce the largest response time.

BF - ES

- 32 -

Response times and critical instants

- **Observation:**

For RM, the critical instant t of a task τ_i is given by the time when $\tau_{i,j}$ arrives together with all tasks $\tau_1, \dots, \tau_{i-1}$ with higher priority.

