



REVIEW: Multiprocessor Scheduling

Given

- n equivalent processors,
- a finite set M of aperiodic/periodic tasks

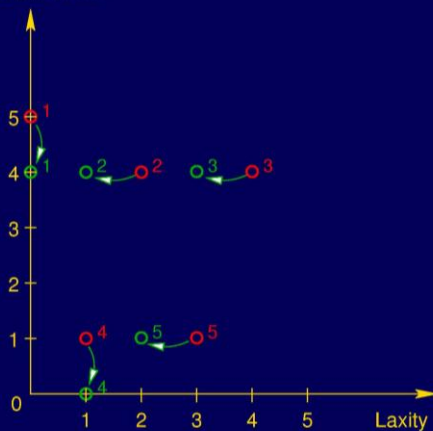
find a schedule such that each task always meets its deadline.

Assumptions:

- Tasks can freely be migrated between processors
 - at any integer time instant, without overhead
 - however: no task may run on two processors simultaneously
- All tasks are preemptable
 - at any integer time instant, without overhead

REVIEW: LLF (Least Laxity First)

Remaining computation time



When tasks are released, they are inserted into the game board according to their WCET and laxity (= deadline – remain. comp. time).


In every time scheduling step / turn of the game:
 – at most n nodes go down by 1
 – the rest moves 1 to the left

LLF is optimal.

BF - ES

- 3 -

REVIEW: Schedulability



Within a set M of aperiodic tasks, we identify three classes with respect to the next k time units starting at time t :

1. Tasks that have to be **fully run** within the next k time units:

$$FR(t, k) = \{i \in M \mid D_i(t) \leq k\}$$

2. Tasks that have to be **partially run** within the next k time units:

$$PR(t, k) = \{i \in M \mid L_i(t) \leq k \wedge D_i(t) > k\}$$

3. Tasks that **need not be run** within the next k time units:

$$NN(t, k) = \{i \in M \mid L_i(t) > k\}$$

BF - ES

- 4 -

REVIEW: Surplus computing power

$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

Lemma: $SCP(0, k) \geq 0$ for all $k > 0$ is a **necessary** condition for schedulability.

- Executing $FR(0, k)$ require $\sum_{i \in FR(0, k)} C_i$
- Executing $PR(0, k)$ require $\sum_{i \in PR(0, k)} (k - L_i(0))$
- kn total computing power in k steps
horizontal lines downward during k steps

BF - ES

- 5 -

Surplus computing power

$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

Theorem: If all tasks are released at time 0, then $SCP(0, k) \geq 0$ for all $k > 0$ is a **necessary and sufficient** condition for schedulability.

We show that no deadlines are missed using LLF.
Proof by induction on t :

1) $SCP(t, k) \geq 0 \quad \forall k$
 \Rightarrow # of tokens on the $-$ axis is $\leq n$

2) $SCP(t, k) \geq 0 \quad \forall k$
 $\Rightarrow SCP(t+1, k) \geq 0 \quad \forall k$
 # token on C axis

Ad 1) $SCP(t, 1) \geq 0$
 $n \geq \sum_{i \in FR(t, 1)} C_i(t) + \sum_{i \in PR(t, 1)} (1 - L_i(t))$

BF - ES


- 6 -

$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

Ad 2)

Claim: $\forall k \geq 0 \exists k'$ s.t.
 $SCP(t+1, k) \geq SCP(t, k')$

Case 1: All tokens left of $L=k$ have arrived by vertical move



pick $k' = k$: $\Delta SCP = SCP(t+1, k) - SCP(t, k)$

- Token in FR: contributes 1 to ΔSCP
- Token in PR: contributes 0 to ΔSCP
- Token that moves from PR to FR, hitting $L+C=k$


$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

Contribution to ΔSCR : $-C_i(t+1) + (k - L_i(t))$
 $= -(k - L_i(t+1)) + (k - L_i(t)) = 0$
 $= L_i(t)$

- Token that moved from NM to line $L=k$:
 Contribution to ΔSCR : $-(k - L_i(t+1)) = 0$

Case 2: Some of the tokens left of $L=k$ have arrived by horizontal move

pick $k' := k+1$



$\Delta SCP = SCP(t+1, k) - SCP(t, k+1)$

$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

By LLF, a token must have more demand left of $L = k+1$.

Consider a token that more demand

in PR: contribution to $\Delta SCP =$
 $-(k - L_i(t+1)) + (k+1 - L_i(t))$

$= 1$
 in FR: contribution to $\Delta SCP =$
 $-C_i(t+1) + C_i(t) = 1$

$$SCP(t, k) = \underbrace{kn}_{\text{avail. computing power}} - \underbrace{\sum_{i \in FR(t, k)} C_i(t)}_{\text{needed for full runs}} - \underbrace{\sum_{i \in PR(t, k)} (k - L_i(t))}_{\text{needed for partial runs}}$$

Consider a token that more demand horizontally:

LLF chooses token in FR first

\Rightarrow token in PR:


Contribution to ΔSCP : $-(k - L_i(t+1)) + (k+1 - L_i(t))$

$L_i(t+1) = L_i(t) + 1$


$= 0$
 $\Delta SCP = k \cdot h - (k+1) \cdot h + h$
 $= 0$

Online scheduling?

Theorem: There can be no optimal scheduling algorithm if the release times are not known a priori.



Case 1: Scheduler selects $A+B$




New tasks D, E arrive
with $L=0, C=1$ at $t=1$
 \Rightarrow no schedule

Feasible schedule: $B+C, D+E, A$

BF - ES

- 11 -

Case 2: Scheduler selects $B+C$



New tasks D, E arrive with
 $L=0, C=2$ at $t=2$

\hookrightarrow no schedule

Feasible schedule: $A+B, A+C, D+E, D+E$

BF - ES

- 12 -

Periodic tasks

Theorem: A necessary and sufficient condition for the schedulability of periodic tasks is that $U \leq n$.

Scheduling idea

1. Divide the time line into time slices such that each period of each process is divided into an integral number of time slices.

Slice length $T = \text{GCD}(T_1, \dots, T_n)$.

2. Within each time slice, allocate processor time in proportion to the utilization $U_i = \frac{C_i}{T_i}$ originating from the various tasks.

Processing time per slice $r_i = T U_i = T \frac{C_i}{T_i}$.

Hence, each task runs $\frac{T_i}{T} r_i = \frac{T_i}{T} T \frac{C_i}{T_i} = C_i$ time units within its period.

3. Allocate r_i according to the following algorithm
 - (a) Look for the first processor proc_j that has free capacity in its time slices.
 - (b) Allocate that portion of r_i to proc_j that proc_j can accommodate.
 - (c) If all of r_i has been allocated then proceed with the next task (goto step a).
 - (d) Otherwise allocate the remainder of r_i to proc_{j+1} .
 proc_{j+1} has enough spare capacity as it has not previously been used and $r_i \leq T$ due to $U_i \leq 1$. Furthermore, due to $r_i \leq T$, we don't generate temporal overlap between the two partial runs of task i .

Example (2 processors)

	τ_1	τ_2	τ_3
T_i	4	8	6
C_i	2	8	3

$$U = \frac{2}{4} + \frac{8}{8} + \frac{3}{6} = 2$$


$$T = \text{gcd}(4, 8, 6) = 2$$

In each slice,

$$T_1 \text{ has } 2 \cdot \frac{2}{4} = 1 \text{ unit}$$

$$T_2 \text{ has } 2 \cdot \frac{8}{8} = 2 \text{ units}$$

$$T_3 \text{ has } 2 \cdot \frac{3}{6} = 1 \text{ unit}$$



BF - ES

- 15 -

Scheduling idea

This scheme works if

- the load isn't too high:

$$U = \sum_{i \in M} \frac{C_i}{T_i} \leq n$$

and

- the time slices allocated have integral length:

$$r_i = T U_i = T \frac{C_i}{T_i} \in \mathbb{N} \text{ for each } i \in M$$

BF - ES

- 16 -

Rescheduling fractional parts

- Let $X_i = T \cdot C_i / T_i - \lfloor T \cdot C_i / T_i \rfloor$
- In each period,
 - allocate in $X_i \cdot T_i / T$ slices: $\lfloor T \cdot C_i / T_i \rfloor + 1$ units
 - and in all other slices: $\lfloor T \cdot C_i / T_i \rfloor$ units
- This can be done without allowing any task to miss its deadline: use EDF!

BF - ES

- 17 -

Example (2 processors)

	τ_1	τ_2	τ_3
T_i	4	6	4
C_i	2	4	3

$$U = \frac{2}{4} + \frac{4}{6} + \frac{3}{4} = \frac{23}{12} < 2$$

$$T = \text{gcd}(4, 6, 4) = 2$$

In each slice

$$T_1 \text{ has } \lfloor 2 \cdot \frac{2}{4} \rfloor = 1 \text{ time unit}$$

$$T_2 \text{ has } \lfloor 2 \cdot \frac{4}{6} \rfloor = \lfloor \frac{8}{3} \rfloor = 1 \text{ time unit}$$

$$T_3 \text{ has } \lfloor 2 \cdot \frac{3}{4} \rfloor = \lfloor \frac{6}{4} \rfloor = 1 \text{ time unit}$$


T_2, T_3 have fractional parts

$\Rightarrow T_2$ needs 1 extra unit every 3 slices

T_3 needs 1 extra unit every 2 slices

BF - ES

- 18 -



Extension: Task migration time

Theorem: A **necessary** and **sufficient** condition for scheduling periodic tasks on n processors is $U \leq n$, if the task migration time is one unit.

Extension: Task migration time

Lemma: If $U \leq n$, then **within each time slice** the tasks can meet the migration time requirement without missing deadlines, if the task migration time is one unit.

- Sort tasks according to non-increasing computation time
- If computation block = T
 \Rightarrow allocate processor exclusively.
- If computation time is $< T$:
 - allocate completely on 1 processor if possible \Rightarrow no migration
 - allocate part of computation at the end of proc i , rest at beginning of proc $i+1$
 \Rightarrow gap of at least 1.

BF - ES

- 21 -

Extension: Task migration time

Lemma: If $U \leq n$, then **between time slices** the tasks can meet the migration time requirement without missing deadlines, if the task migration time is one unit.

- For each slice, sort tasks according to non-increasing computation blocks
- If computation block = $T \rightarrow$ find processor that executed task at the end of the previous slice
 \rightarrow no migration
- if no such processor exists \Rightarrow assign to same left-over processor at the end (migration time already accounted for in previous slice)

BF - ES

- 22 -

• If complete block $< T$

=> Find processor j that executed task at the end of the previous slice

Assign as much as possible to current processor;

if insufficient, use j from beginning.

(no migration at beginning, ≥ 1 unit gap within slice)

If no such processor exists => assign task later (migration time already accounted for in the previous slice).