



# Parallel & Distributed Real-Time Systems

## Lecture #6

Risat Pathan

Department of Computer Science and Engineering  
Chalmers University of Technology

# Administrative issues

## Group registration:

- Please register to a homework assignment group via PingPong. There are 20 groups available.
- HW#1 will be available on April 19, on Friday.

# Feasibility testing

## What techniques for feasibility testing exist?

- Hyper-period analysis (for static and dynamic priorities)
  - In a simulated schedule no task execution may miss its deadline
- Guarantee bound analysis (for static and dynamic priorities)
  - The fraction of processor time that is used for executing the task set must not exceed a given bound
- Response time analysis (for static priorities)
  - The worst-case response time for each task must not exceed the deadline of the task
- Processor demand analysis (for dynamic priorities)
  - The accumulated computation demand for the task set under a given time interval must not exceed the length of the interval

# Feasibility testing

## What techniques for feasibility testing exist?

- Hyper-period analysis (for static and dynamic priorities)
  - In a simulated schedule no task execution may miss its deadline
- Guarantee bound analysis (for static and dynamic priorities)
  - The fraction of processor time that is used for executing the task set must not exceed a given bound
- Response time analysis (for static priorities)
  - The worst-case response time for each task must not exceed the deadline of the task
- Processor demand analysis (for dynamic priorities)
  - The accumulated computation demand for the task set under a given time interval must not exceed the length of the interval

# Guarantee bound analysis

## Basic principle:

- If the accumulated utilization  $U$  of all tasks in the system does not exceed a guarantee bound, all timing constraints will be met.
- The guarantee bound  $U_{GB}$  is expressed as a fraction of the available processing capacity of the system.  
(= 100% multiplied by the number of processors)
- The utilization  $U_i$  of a task is expressed as the fraction of processing capacity used for executing the task.

Thus, guarantee bound analysis will have a polynomial time complexity

$$\text{task utilization} = \frac{C_i}{T_i}$$

$$\text{accumulated utilization} = \sum_{i=1}^n \frac{C_i}{T_i}$$

# Guarantee bound analysis

A good guarantee bound ...

... enables prediction of required processing capacity, e.g. # and speed of processors, of the hardware (when software is known)

... enables derivation of timing parameters, e.g. periods of tasks, in the software (when hardware implementation is known)


=

A good guarantee bound ...

... enables prediction of how “strong” the hardware implementation must be (when the software “load” is known)

... enables prediction of how high the software “load” is allowed to be (when the “strength” of the hardware implementation is known)

# Guarantee bound analysis



A good guarantee bound ...

... enables prediction of how “strong” the hardware implementation must be (when the software “load” is known)

... enables prediction of how high the software “load” is allowed to be (when the “strength” of the hardware implementation is known)

# Guarantee bound analysis

Guarantee bound analysis for RM: (Liu & Layland, 1973)

- The guarantee bound for RM scheduling is

$$U_{\text{GB-RM}} = n \left( 2^{1/n} - 1 \right)$$

- A conservative lower limit on the guarantee bound can be derived by letting  $n \rightarrow \infty$

$$\lim_{n \rightarrow \infty} n \left( 2^{1/n} - 1 \right) = \ln 2 \approx 0.693$$



# Guarantee bound analysis

Guarantee bound analysis for RM: (Liu & Layland, 1973)

- A sufficient condition for RM scheduling is

$$U = \sum_{i=1}^n \frac{C_i}{T_i} \leq n(2^{1/n} - 1)$$

The test is only valid if all of the following conditions apply:

1. Single-processor system
2. Synchronous task sets
3. Independent tasks
4. Periodic or sporadic tasks
5. Tasks have deadlines equal to period ( $D_i = T_i$ )

# Guarantee bound analysis

## Guarantee bound analysis for RM: (Liu & Layland, 1973)

- The proof of the condition uses the fact that the worst-case response time for a task occurs at a critical instant (where the task arrives at the same time as all higher-priority tasks)
- The feasibility test is derived using an analysis of this special case
- The proof also shows that if the task set is schedulable for the critical instant case, it is also schedulable for any other case
- The proof is given in Krishna and Shin (Section 3.2.1)  
Highly recommended reading!



# Guarantee bound analysis

Guarantee bound analysis for EDF: (Liu & Layland, 1973)

- A sufficient and necessary condition for EDF scheduling is

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

$$U_{\text{GB-EDF}} = 1$$

The test is only valid if all of the following conditions apply:

1. Single-processor system
2. Synchronous task sets
3. Independent tasks
4. Periodic tasks
5. Tasks have deadlines equal to period ( $D_i = T_i$ )

# Feasibility testing

## What techniques for feasibility testing exist?

- Hyper-period analysis (for static and dynamic priorities)
  - In a simulated schedule no task execution may miss its deadline
- Guarantee bound analysis (for static and dynamic priorities)
  - The fraction of processor time that is used for executing the task set must not exceed a given bound
- Response time analysis (for static priorities)
  - The worst-case response time for each task must not exceed the deadline of the task
- Processor demand analysis (for dynamic priorities)
  - The accumulated computation demand for the task set under a given time interval must not exceed the length of the interval

# Response-time analysis

## Response time:

- The response time  $R_i$  for a task  $\tau_i$  represents the worst-case completion time of the task when execution interference from other tasks are accounted for.

The response time for a task  $\tau_i$  consists of:

$C_i$  The task's uninterrupted execution time (WCET)

$I_i$  Interference from higher-priority tasks

$$R_i = C_i + I_i$$

# Response-time analysis

## Interference:

- For static-priority scheduling, the interference term is

$$I_i = \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i}{T_j} \right\rceil C_j$$

where  $hp(i)$  is the set of tasks with higher priority than  $\tau_i$ .

- The response time for a task  $\tau_i$  is thus:

$$R_i = C_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i}{T_j} \right\rceil C_j$$

# Response-time analysis

## Response-time calculation:

- The equation does not have a simple analytic solution.
- However, an iterative procedure can be used:

$$R_i^{n+1} = C_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i^n}{T_j} \right\rceil C_j$$

- The iteration starts with a value that is guaranteed to be less than or equal to the final value of  $R_i$  (e.g.  $R_i^0 = C_i$ )
- The iteration completes at convergence ( $R_i^{n+1} = R_i^n$ ) or if the response time exceeds the deadline  $D_i$

# Response-time analysis

Schedulability test: (Joseph & Pandya, 1986)

- An exact condition for static-priority scheduling is

$$\forall i: R_i \leq D_i$$

The test is only valid if all of the following conditions apply:

1. Single-processor system
2. Synchronous task sets
3. Independent tasks
4. Periodic tasks
5. Tasks have deadlines not exceeding the period ( $D_i \leq T_i$ )



# Response-time analysis

Time complexity:

Response-time analysis has pseudo-polynomial time complexity

Proof:

calculating the response-time for task  $\tau_i$  requires no more than  $D_i$  iterations

since  $D_i \leq T_i$  the number of iterations needed to calculate the response-time for task  $\tau_i$  is bounded above by  $T_i$

the procedure for calculating the response-time for all tasks is therefore of time complexity  $O(\max\{T_i\})$

the longest period of a task is also the largest number in the problem instance

# Response-time analysis

## Accounting for blocking:

- Blocking caused by critical regions
  - Blocking factor  $B_i$  represents the length of critical region(s) that are executed by processes with lower priority than  $\tau_i$
- Blocking caused by non-preemptive scheduling
  - Blocking factor  $B_i$  represents largest WCET (not counting  $\tau_i$ )

$$R_i = C_i + B_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_j}{T_j} \right\rceil C_j$$

Observation: the feasibility test is now only sufficient since the worst-case blocking will not always occur at run-time.

# Response-time analysis

## Accounting for blocking: (using PCP or ICPP)

When using priority ceiling a task  $\tau_i$  can only be blocked once by a task with lower priority than  $\tau_i$ .

This occurs if the lower-priority task is within a critical region when  $\tau_i$  arrives, and the critical region's ceiling priority is higher than or equal to the priority of  $\tau_i$ .

Blocking now means that the start time of  $\tau_i$  is delayed (= the blocking factor  $B_i$ )

As soon as  $\tau_i$  has started its execution, it cannot be blocked by a lower-priority task.

# Response-time analysis

Accounting for blocking: (using PCP or ICPP)

Determining the blocking factor for  $\tau_i$

1. Determine the ceiling priorities for all critical regions.
2. Identify the tasks that have a priority lower than  $\tau_i$  and that calls critical regions with a ceiling priority equal to or higher than the priority of  $\tau_i$ .
3. Consider the times that these tasks lock the actual critical regions. The longest of those times constitutes the blocking factor  $B_i$ .



**(this page intentionally left blank)**

# Processor-demand analysis

## Processor demand:

- The processor demand for a task  $\tau_i$  in a given time interval  $[0, L]$  is the amount of processor time that the task needs in the interval in order to meet the deadlines that fall within the interval.

Let  $N_i^L$  represent the number of instances of  $\tau_i$  that must complete execution before  $L$ .

The total processor demand up to  $L$  is


$$C_P(0, L) = \sum_{i=1}^n N_i^L C_i$$

# Processor-demand analysis

## Number of relevant task arrivals:

- We can calculate  $N_i^L$  by counting how many times task  $\tau_i$  has arrived during the interval  $[0, L - D_i]$

We can ignore instance of the task that has arrived during the interval  $[L - D_i, L]$  since  $D_i > L$  for these instances.



# Processor-demand analysis

Processor-demand analysis:

- We can express  $N_i^L$  as

$$N_i^L = \left\lfloor \frac{L - D_i}{T_i} \right\rfloor + 1$$

The total processor demand is thus

$$C_P(0, L) = \sum_{i=1}^n \left( \left\lfloor \frac{L - D_i}{T_i} \right\rfloor + 1 \right) C_i$$



# Processor-demand analysis

Schedulability test: (Baruah et al., 1990)

- A sufficient and necessary condition for EDF scheduling is

$$\forall L \in K : C_p(0, L) \leq L$$

The test is only valid if all of the following conditions apply:

1. Single-processor system
2. Synchronous task sets
3. Independent tasks
4. Periodic tasks
5. Tasks have deadlines not exceeding the period ( $D_i \leq T_i$ )

# Processor-demand analysis

Schedulability test: (Baruah et al., 1990)

- The set of control points  $K$  is

$$K = \left\{ D_i^k \mid D_i^k = kT_i + D_i, D_i^k \leq L_{\max}, 1 \leq i \leq n, k \geq 0 \right\}$$

$$L_{\max} = \max \left\{ D_1, \dots, D_n, \frac{\sum_{i=1}^n (T_i - D_i) U_i}{1 - U} \right\}$$

Observation:

$$L_{\max} \leq \max \left\{ \max \{ D_i \}, \frac{U}{1 - U} \max \{ T_i - D_i \} \right\} \leq \max \left\{ \max \{ T_i \}, \frac{U}{1 - U} \max \{ T_i \} \right\}$$

# Processor-demand analysis

Time complexity:

Processor-demand analysis has pseudo-polynomial time complexity if total task utilization is less than 100%

Proof:

the number of control points needed to check the processor demand is bounded above by

$$Q_L^{\max} = \max \left\{ \max \{ T_i \}, \frac{U}{1-U} \max \{ T_i \} \right\} = \max \left\{ 1, \frac{U}{1-U} \right\} \square \max \{ T_i \}$$

since  $U/(1-U)$  is a constant the procedure for calculating the processor demand is therefore of time complexity  $O(\max \{ T_i \})$

the longest period of a task is also the largest number in the problem instance

# Processor-demand analysis

Accounting for blocking: (using Stack Resource Policy)

Tasks are assigned static preemption levels:

The preemption level of task  $\tau_i$  is denoted  $\pi_i$

Task  $\tau_i$  is not allowed to preempt another task  $\tau_j$  unless  $\pi_i > \pi_j$

If  $\tau_i$  has higher priority than  $\tau_j$  and arrives later, then  $\tau_i$  must have a higher preemption level than  $\tau_j$ .

Note:

- The preemption levels are static values, even though the tasks priorities may be dynamic.
- For EDF scheduling, suitable levels can be derived if tasks with shorter relative deadlines get higher preemption levels, that is:

$$\pi_i > \pi_j \iff D_i < D_j$$

# Processor-demand analysis

Accounting for blocking: (using Stack Resource Policy)

Resources are assigned dynamic resource ceilings:

Each shared resource is assigned a ceiling that is always equal to the maximum preemption level among all tasks that may be blocked when requesting the resource.

The protocol keeps a system-wide ceiling that is equal to the maximum of the current ceilings of all resources.

A task with the earliest deadline is allowed to preempt only if its preemption level is higher than the system-wide ceiling.

Note:

The original priority of the task is not changed at run-time.

The resource ceiling is a dynamic value calculated at run-time as a function of current resource availability.

# Processor-demand analysis

Accounting for blocking: (using Stack Resource Policy)

Blocking factor  $B_i$  represents the length of critical / non-preemptive regions that are executed by tasks with lower preemption levels than  $\tau_i$

Tasks are indexed in the order of increasing preemption levels, that is:  $\pi_i > \pi_j \Leftrightarrow i < j$

$$\forall L \in K, \forall i \in [1, n]: C_P^i(0, L) \leq L$$

$$C_P^i = \sum_{k=1}^i \left( \left\lfloor \frac{L - D_k}{T_k} \right\rfloor + 1 \right) C_k + \left( \left\lfloor \frac{L - D_i}{T_i} \right\rfloor + 1 \right) B_i$$

# Processor-demand analysis

Accounting for blocking: (using Stack Resource Policy)

Determining the blocking factor for  $\tau_i$

1. Determine the worst-case resource ceiling for each critical region, that is, assume the run-time situation where the corresponding resource is unavailable.
2. Identify the tasks that have a preemption level lower than  $\tau_i$  and that calls critical regions with a worst-case resource ceiling equal to or higher than the preemption level of  $\tau_i$ .
3. Consider the times that these tasks lock the actual critical regions. The longest of those times constitutes the blocking factor  $B_i$ .

---

# End of lecture #6