# Real-Time Systems

Anno accademico 2009/10 Laurea magistrale in informatica Dipartimento di Matematica Pura e Applicata Università di Padova Tullio Vardanega



## Outline

- Introduction 1.
- Dependability issues 2.
- Scheduling issues 3. 4.
- More on fixed-priority scheduling
- Task interactions and blocking 5.
- System issues 6.
- 7. Multi-cores and distribution
- Bibliography

2009/10 UniPD, T. Vardanega

J. Liu, "Real-Time Systems", Prentice Hall, 2000 A. Burns, A. Wellings, "Concurrent and Real-Time Programming in Ada", Cambridge University Press, 2007 A. Burns, A. Wellings, "Real Time Systems and Programming Languages: Ada 95, Real-Time Java and Real-Time C/POSIX", Addison-Wesley, 2009

Real-time systems

2 of 37

4 of 37

4. More on Fixed-Priority Scheduling

> Credits to A. Burns and A. Wellings Rented to the second

#### Simple workload model

- The application is assumed to consist of a fixed set of tasks
- All tasks are periodic with known periods D This defines the periodic workload model
- The tasks are completely <u>independent</u> of each other
- All system overheads (context-switch times and so on) are ignored
- Assumed to have zero cost or otherwise negligible
- All tasks have a deadline <u>equal</u> to their period Each task must complete before it is next released
- All tasks have a fixed WCET
- Operation modes are not considered

2009/10 UniPD, T. Vardanega

#### Standard notation

- B: Worst-case blocking time for the task (if applicable)
- C: Worst-case computation time (WCET) of the task
- D: Deadline of the task
- The interference time of the task I:
- Release jitter of the task I:
- N: Number of tasks in the system
- P: Priority assigned to the task (if applicable)
- R: Worst-case response time of the task
- T: Minimum time between task releases (task period)

Real-time systems

U: The utilization of each task (equal to C/T)

The name of a task a-Z:

5 of 37

## Fixed-priority scheduling (FPS)

- Currently this is the most widely used approach And it is the distinct focus of this segment
- Each task has a fixed, static, priority which is computed off-line
- The ready tasks are dispatched to execution in the order determined by their priority
- In real-time systems the "priority" of a task is derived from its temporal requirements, not its importance to the correct functioning of the system or its integrity (!)

```
2009/10 UniPD, T. Vardanega
```

2009/10 UniPD, T. Vardanega

#### Preemption and non-preemption – 1

- With priority-based scheduling, a high-priority task may be released during the execution of a lower priority one
- In a preemptive scheme, there will be an immediate switch to the higher-priority task
- With non-preemption, the lower-priority task will be allowed to complete before the other may execute
- Preemptive schemes enable higher-priority tasks to be more reactive, hence they are preferred

### Preemption and non-preemption -2

- Alternative strategies allow a lower priority task to continue to execute for a bounded time
- These schemes are known as deferred preemption or . cooperative dispatching
- Schemes such as EDF can also take on a preemptive or non-preemptive form
- Value-based scheduling (VBS) can too VBS is useful when the system becomes overloaded and some adaptive scheme of scheduling is needed
  - UBS consists in assigning a value to each task and then employing an on-line value-based scheduling algorithm to decide which task to run next



## Rate-monotonic priority assignment

- Each task is assigned a (unique) priority based on its period
  - The shorter the period, the higher the priority
- Tasks are assigned distinct priorities (!) For any two tasks i and j
  - $T_i < T_j \Longrightarrow P_i > P_j$

9 of 37

11 of 37

- This assignment is <u>optimal</u> If any task set can be scheduled (using preemptive priority-based scheduling) with a fixed-priority assignment scheme, then the given task set can also be scheduled with a rate monotonic assignment scheme
  - This is termed rate monotonic scheduling Nomenclature
- Priority 1 as numerical value is the lowest (least) priority but the indices are still sorted highest to lowest (!)

2009/10 UniPD, T. Vardaneg Real-time systems

#### Utilization-based analysis

• A simple sufficient but not necessary schedulability condition exists for rate monotonic scheduling But only for task sets with D=T

$$U \equiv \sum_{i=1}^{N} \frac{C_i}{T_i} \le N (2^{1/N} - 1)$$

$$U \leq 0.69$$
 as  $N \to \infty$ 



#### Example: task set A

| Task | Period | Computation Time | Priority | Utilization |
|------|--------|------------------|----------|-------------|
|      | Т      | С                | Р        | U           |
| a    | 50     | 12               | 1 (low)  | 0.24        |
| b    | 40     | 10               | 2        | 0.25        |
| с    | 30     | 10               | 3 (high) | 0.33        |

■ The combined utilization is 0.82 (or 82%)

- This is above the threshold for three tasks (0.78), hence this task set fails the utilization test
- Then we have no a-priori answer

2009/10 UniPD, T. Vardanega

Real-time systems



Timeline for task set A



### Example: task set B

| Task | Period | Computation Time | Priority | Utilization |
|------|--------|------------------|----------|-------------|
|      | Т      | С                | Р        | U           |
| a    | 80     | 32               | 1 (low)  | 0.40        |
| b    | 40     | 5                | 2        | 0.125       |
| с    | 16     | 4                | 3 (high) | 0.25        |

• The combined utilization is 0.775

• This is below the threshold for three tasks (0.78), hence this task set will meet all its deadlines

2009/10 UniPD, T. Vandanega Real-time systems 13 of 57

### Example: task set C

| Task | Period | Computation Time | Priority | Utilization |
|------|--------|------------------|----------|-------------|
|      | Т      | С                | Р        | U           |
| a    | 80     | 40               | 1 (low)  | 0.50        |
| ь    | 40     | 10               | 2        | 0.25        |
| с    | 20     | 5                | 3 (high) | 0.25        |

The combined utilization is 1.0

• This is above the threshold for three tasks (0.78) but the task set will meet all its deadlines (!)

| 2009/10 UniPD, T. Vardanega |  | Real-time systems | 14 of 37 |
|-----------------------------|--|-------------------|----------|
|-----------------------------|--|-------------------|----------|

Timeline for task set C



### Critique of utilization-based tests

- They are not exact
- They are not general
- But they are Ω(N)
   Which makes them interesting for a large class of users
- The test is said to be sufficient but not necessary and as such falls in the class of "*schedulability tests*"

Real-time system

Response time analysis – 1

- The worst-case response time *R* of task *i* is calculated first and then checked (trivially) with its deadline
  - $R_i \leq D_i$

 $R_i = C_i + I_i$ 

Where I is the interference from higher priority tasks

2009/10 UniPD, T. Vardanega

17 of 37

## Calculating R

2009/10 UniPD, T. Vardanega

2009/10 UniPD, T. Vardaneş

During R, each higher priority task *j* will execute a number of times

Number of Releases = 
$$\left| \frac{R_i}{T_i} \right|$$

□ The ceiling function [ ] gives the smallest integer greater than the fractional number on which it acts
 ■ E.g., the ceiling of 1/3 is 1, of 6/5 is 2, and of 6/3 is 2
 ■ The total interference is given by [<u>R</u>]

Real-time systems

The total interference is given i

18 of 37

 $\overline{T_j}$ 

Response time equation

$$R_{i} = C_{i} + \sum_{j \in hp(i)} \left\lceil \frac{R_{i}}{T_{j}} \right\rceil C_{j}$$

Where hp(i) is the set of tasks with priority higher than task i
Solved by forming a recurrence relationship

$$w_i^{n+1} = C_i + \sum_{j \in hp(i)} \left[ \frac{w_i^n}{T_j} \right] C_j$$

• The set of values  $W_i^0$ ,  $w_i^1$ ,  $w_i^2$ ,  $\dots$ ,  $w_i^n$ ,  $\dots$  is monotonically non-decreasing when  $w_i^n = w_i^{n+1}$  the solution to the equation has been found, must not be greater than  $w_i^0$  (e.g. 0 or  $C_i$ )

| 2009/10 UniPD, T. Vardanega | Real-time systems | 19 of 37 |
|-----------------------------|-------------------|----------|

### Response time algorithm



#### Example: task set D



#### Revisiting task set C

| Task | Period | Computation Time | Priority | Response Time |
|------|--------|------------------|----------|---------------|
|      | Т      | С                | Р        | R             |
| a    | 80     | 40               | 1 (low)  | 80            |
| b    | 40     | 10               | 2        | 15            |
| с    | 20     | 5                | 3 (high) | 5             |

The combined utilization is 1.0

- This is above the utilization threshold for three tasks (0.78) hence the utilization-based schedulability test failed
- But response time analysis shows that the task set will meet all its deadlines

2009/10 UniPD, T. Vardanega

#### 23 of 37

# Example (cont'd) $\begin{cases} w_{c}^{0} = 5 \\ w_{c}^{1} = 5 + \left\lceil \frac{5}{7} \right\rceil 3 + \left\lceil \frac{5}{12} \right\rceil 3 = 11 \\ w_{c}^{2} = 5 + \left\lceil \frac{11}{7} \right\rceil 3 + \left\lceil \frac{11}{12} \right\rceil 3 = 14 \\ w_{c}^{3} = 5 + \left\lceil \frac{14}{7} \right\rceil 3 + \left\lceil \frac{14}{12} \right\rceil 3 = 17 \\ w_{c}^{4} = 5 + \left\lceil \frac{17}{7} \right\rceil 3 + \left\lceil \frac{17}{12} \right\rceil 3 = 20 \\ w_{c}^{5} = 5 + \left\lceil \frac{20}{7} \right\rceil 3 + \left\lceil \frac{20}{12} \right\rceil 3 = 20 \\ \boxed{R_{c} = 20}$ 2009/10 UniPD, T. Vardangs 22 of 37

#### Response time analysis – 2

- RTA is sufficient and necessary
   Hence it falls by right in the class of *feasibility tests*
- If the task set passes the test its tasks will meet all their deadlines
- If it fails the test then, at run time, a task will miss its deadline
  - Unless the computation time estimations (the WCET) themselves turn out to be pessimistic

Real-time systems

#### Sporadic tasks

- Sporadic tasks have a minimum inter-arrival time
   Which must be preserved at run time if schedulability is to be ensured, but how can it ?
- They also require D≤T
- The response time algorithm for fixed-priority scheduling works perfectly for D<T as long as the stopping criterion becomes</li>
- $W_i^{n+1} > D_i$ • Interestingly this also works perfectly well with *any* priority ordering

| 2009/10 UniPD, T. Vardanega Real-time systems 25 | 5 of 31 |
|--|---------|
|--|---------|

#### Hard and soft tasks

- In many situations the WCET for sporadic tasks are considerably higher than the average case
- Interrupts often arrive in bursts and an abnormal sensor reading may lead to significant additional computation
- Measuring schedulability with WCET may lead to very low processor utilizations being observed in the actual running system



#### General guidelines

- Rule 1 All tasks should be schedulable using average execution times and average arrival rates
  - There may therefore be situations in which it is not possible to meet all current deadlines
  - $\hfill\square$  This condition is known as a transient overload
- Rule 2
  - All hard real-time tasks should be schedulable using WCET and worst-case arrival rates of all tasks (including soft)
  - □ No hard real-time task will therefore miss its deadline
  - □ If Rule 2 incurs unacceptably low utilizations for "normal execution" then WCET values or arrival rates must be reduced

2009/10 UniPD, T. Vardanega Real-time systems

27 of 37

29 of 37

# Handing aperiodic tasks – 1

- These do not have minimum inter-arrival times
   But also no deadline
- However we may be interested in the system being responsive to themWe can run aperiodic tasks at a priority below the priorities
- assigned to hard tasks In a preemptive system they therefore cannot steal resources from the hard tasks
- This does not provide adequate support to soft tasks which will often miss their deadlines
- To improve the situation for soft tasks, a server can be employed
- Servers protect the processing resources needed by hard tasks but otherwise allow soft tasks to run as soon as possible

### Handing aperiodic tasks – 2

#### Polling server (PS)

- A fixed priority periodic task serves the aperiodic tasks requests
- It is given a fixed computing time quantum that uses to serve aperiodic task requests
- If no aperiodic tasks require execution the server time quantum is given over to execute periodic tasks
- The time quantum is reallocated to the server at the start of the new period



Real-time systems

# Handing aperiodic tasks – 3

#### Deferrable Server (DS)

2009/10 UniPD, T. Vardaneg

- High-priority periodic server handles aperiodic requestsSimilar in principle to PS
- However, if no aperiodic tasks require execution, the server retains its time quantum
  - Hence if an aperiodic task requires execution during the server period, it can be served immediately
     In the absence of pending requests the server does not sleep but
    - just waits for any incoming one
  - The time quantum is reallocated to the server at the start of the new period

Real-time systems

#### Handing aperiodic tasks – 4

#### • Priority Exchange (PE)

- □ High-priority periodic server serves aperiodic tasks, if any
- □ Similar in principle to DS
- □ If no aperiodic tasks require execution
  - PE exchanges its own priority with that of the pending (soft) periodic task with priority lower than that of itself (the server) and highest amongst all other pending periodic tasks
  - Hence the selected periodic task inherits a priority higher than its own

| 2009/10 UniPD, T. Vardanoga Real-time systems 31 of |
|---|
|---|

## Handing aperiodic tasks – 5

#### Sporadic Server (SS)

 High-priority periodic server activated (enabled) at a sufficiently high rate to server requests from sporadic tasks

#### $\square SS \neq DS$

- The time quantum is replenished only when exhausted, rather than at each server activation
- This places a tolerable bound on the overhead caused by the server
- □ The SS is the default server policy in POSIX



Task sets with D < T

- For D = T, Rate Monotonic priority assignment (a.k.a. ordering) is optimal
- For D < T, Deadline Monotonic priority assignment (ordering) is optimal

$$D_i < D_j \Longrightarrow P_i > P_j$$

Real-time syste

33 of 37

35 of 37

DMPO is optimal – 1

Deadline monotonic priority ordering (DMPO) is optimal

if any task set Q that is schedulable by priority-driven scheme W it is also schedulable by DMPO

- The proof of optimality of DMPO involves transforming the priorities of Q as assigned by W until the ordering becomes as assigned by DMPO
- Each step of the transformation will preserve schedulability

# DMPO is optimal – 2

2009/10 UniPD, T. Vardane

- Let i, j be two tasks with adjacent priorities in Q such that under W  $P_i > P_j \land D_i > D_j$
- Define scheme W' to be identical to W except that tasks i and j are swapped
- Now consider the schedulability of Q under W'

- All tasks with priorities greater than j will be unaffected by this change to lower-priority tasks
- All tasks with priorities lower than i will be unaffected as they will experience the same interference from i and j
- Task j, which was schedulable under W, now has a higher priority, suffers less interference, and hence must be schedulable under W'

2009/10 UniPD, T. Vardanega

Real-time systems

## DMPO is optimal – 3

All that is left is the need to show that task i, which has had its priority lowered, is still schedulable
Under W

$$R_i < D_i, D_i < D_i and D_i \leq T_i$$

Real-time systems

- Hence task j only interferes once during the execution of task i
- It follows that:

2009/10 UniPD, T. Vardaneş

$$R'_i = R_i \leq D_i < D$$

- Hence task i is still schedulable after the switch
- Priority scheme W' can now be transformed to W" by choosing two more tasks that are in the wrong order for DMP and switching them

2009/10 UniPD, T. Vardanega

# Summary

- A simple (periodic) workload model
- Delving into fixed-priority scheduling
- A (rapid) survey of schedulability tests
- Some extensions to the workload model
- Priority assignment techniques

2009/10 UniPD, T. Vardanega

Real-time systems