



Course Outline

- Basic concepts
- Where timing constraints come from?
- Concurrency control
- Task scheduling
- Feasibility Analysis
- Handling shared resources









































































• Architecture

- cache, pipelining, interrupts, DMA
- Operating system
 - scheduling, synchronization, communication
- Language
 - lack of explicit support for time
- Design methodologies
 - lack of analysis and verification techniques









- 3. Difficult software maintainability
 - Complex appl.s consists of millions lines of code
 - Code understanding takes more that re-writing
 - But re-writing is VERY expensive and bug prone
- 4. Difficult to verify timing constraints without explicit support from the OS and the language

















► n



























Task Criticality

HARD tasks

All jobs must meet their deadlines. Missing a deadline may cause catastrophical effects.

SOFT tasks

Missing deadlines is not desired but causes only a performace degradation.

An operating system able to handle hard tasks is called a hard real-time system.

65

Typical HARD tasks

- sensory acquisition
- low-level control
- sensory-motor planning

Typical SOFT tasks

- reading data from the keyboard
- user command interpretation
- message displaying
- graphical activities














• Timing constraints

- activation, completion, jitter.

• Precedence constraints

- they impose an ordering in the execution.

Resource constraints

they enforce a synchronization in the access of mutually exclusive resources.









$T_d = 1 s$
$T_r = 1 s$
$T_b = v/(\mu g)$

For $\begin{cases} v = 50 \text{ Km/h} = 14 \text{ m/s} \\ \mu = 0.5 \end{cases}$ we have: $T_f = 2.8 \text{ s}$

77

Time to stop the car from the yellow: $D \cong 5 s$















Length covered by the robot after the contact:
$$\mathcal{L} = \mathcal{V} \mathcal{T}_s + \mathcal{X}_f$$
 $\alpha_f = \int_0^{\infty} \mathcal{V}(t) dt = \int_0^{\infty} \mathcal{V}_0 e^{-t/\tau_d} dt = -\mathcal{V}_0 \tau_d (e^{-\infty} - e^0) = \mathcal{V}_0 \tau_d$ $\mathcal{L} = \mathcal{V}(\mathcal{T}_s + \tau_d)$ There on the robot tool:(K = elastic coefficient) $\mathcal{L} = \mathcal{L} = \mathcal{L} \mathcal{V}(\mathcal{T}_s + \tau_d) < \mathcal{L}_{max}$





































- The operating system is the part most responsible for a predictable behavior.
- Concurrency control must be enforced by:
 - appropriate scheduling algorithms
 - appropriate syncronization protocols
 - efficient communication mechanisms
 - predictable interrupt handling















Timeline Scheduling (cyclic scheduling)

It has been used for 30 years in military systems, navigation, and monitoring systems.

Examples

- Air traffic control
- Space Shuttle
- Boeing 777

113

Timeline Scheduling

Method

- The time axis is divided in intervals of equal length (*time slots*).
- Each task is statically allocated in a slot in order to meet the desired request rate.
- The execution in each slot is activated by a timer.







Advantages

- Simple implementation (no real-time operating system is required).
- Low run-time overhead.
- It allows jitter control.









































How can we verify feasibility?

• Each task uses the processor for a fraction of time: $U_{\cdot} = \frac{C_{i}}{C_{i}}$

$$V_i = \frac{C_i}{T_i}$$

• Hence the total **processor utilization** is:

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

• U_p is a misure of the **processor load**



If $U_p > 1$ the processor is overloaded hence the task set cannot be schedulable.

However, there are cases in which $U_p < 1$ but the task is not schedulable by RM.










A sufficient condition

If $U_p \leq U_{lub}$ the task set is certainly schedulable with the RM algorithm.

NOTE

If $U_{lub} < U_p \le 1$ we cannot say anything about the feasibility of that task set.

145





































Dynamic Priority

EDF with $D \le T$

Schedule based on absolute deadlines

Schedulability Analysis

Processor Demand Criterion [Baruah '90]

In any interval, the computation demanded by the task set must be no greater than the available time.

163

























Schedulability Analysis			
	$D_i = T_i$	$D_i \leq T_i$	
RM	$\begin{array}{llllllllllllllllllllllllllllllllllll$	$\begin{array}{l} \textbf{pseudo-polynomial}\\ \text{Response Time Analysis}\\ \forall \text{i} R_i \leq D_i\\ R_i = C_i + \sum_{k=1}^{i-1} \left\lceil \frac{R_i}{T_k} \right\rceil C_k \end{array}$	
EDF	polynomial: $O(n)$ $\Sigma U_i \leq 1$	pseudo-polynomialProcessor Demand Analysis $\forall L > 0, g(0,L) \leq L$	
Graduate Course on Embedded Control Systems - Pisa 8-12 June 2009 176			























However, EDF offers the following advantages with respect to RM:

- Better processor utilization (100%)
- Less overhead due to preemptions;
- More flexible behavior in overload situations;
- More uniform jitter control;
- Better aperiodic responsiveness.

Graduate Course on Embedded Control Systems - Pisa 8-12 June 2009

187

Handling shared resources Problems caused by mutual exclusion







Mars Mission Pathfinder + Sojourner

July 4 th , '97	landing on Mars (after 6 months)	
July 15 th , '97	the system self-resets for timeout: the robot stops working	
July 16 th , '97	News spread all over the world: engineers start working	
July 17 th , '97	the error is identified: probability = 10 ⁻⁶	
		192









Priority Inversion

A high priority task is blocked by a lower- priority task a for an unbounded interval of time.

Solution

Introduce a concurrency control protocol for accessing critical sections.

197

Resource Access Protocols

- Non Preemptive Protocol (NPP)
- Highest Locker Priority (HLP)
- Priority Inheritance Protocol (PIP)
- Priority Ceiling Protocol (PCP)
- Stack Resource Policy (SRP)

198




























































Stack Resource Policy [Baker 1990]

This protocol satisfies the same PCP properties:

- it avoids unbounded priority inversion;
- it prevents deadlocks and chained blocking;
- each task can be blocked at most once;

In addition:

- SRP allows using multi-unit resources;
- it can be used under fixed and dynamic priorities;
- it allows tasks to share the same stack space.

229







































Non preemptive scheduling

Problem Statement

Preemption has a cost, often not negleagible:

- Scheduler cost: due to the time taken by the scheduler to suspend the running task, switch the context, and dispatch the new incoming task.
- Pipeline cost: due to the time taken to flush the pipeline when the task is interrupted and the time taken to refill the pipeline when the task is resumed.
- <u>Cache-related cost</u>: due to the time taken to reload the cache lines evicted by the preempting task.





























A simple bound for Q_i

$$\forall i \ \max_{P_j < P_i} \{q_j\} \le \beta_i$$

$$\begin{cases} i = 1 \ max \{q_2, q_3, q_4\} \le \beta_1 \\ i = 2 \ max \{q_3, q_4\} \le \beta_2 \\ i = 3 \ q_4 \le \beta_3 \end{cases}$$

$$\begin{cases} i = 1 \ q_2 \le \beta_1 \\ i = 2 \ q_3 \le \min\{\beta_1, \beta_2\} \\ i = 3 \ q_4 \le \min\{\beta_1, \beta_2, \beta_3\} \end{cases}$$















Response Time Analysis (FPP)













Remarks

- ➤ Preemption Thresholds are easy to specify, but it is difficult to predict the number of preemptions and where they occur ⇒ large preemption overhead
- Deferred Preemption allows bounding the number of preemptions but it is difficult to predict where they occur. Note that the analysis assumes $q_i^{last} = 0$
- Fixed Preemption Points allow more control on preemptions and can be selected on purpose (e.g., to minimize overhead, stack size, and reduce WCETs).
- A large final chunk in τ_i reduces the interference from hp-tasks (hence R_i), but creates more blocking to hp-tasks.