

Network of Excellence on Embedded Systems Design





Welcome to the Course on Real-Time Kernels for Microcontrollers: Theory and Practice



Real-Time Systems Laboratory

Scuola Superiore Sant'Anna, Pisa June 23-25, 2008

Course Program

Day 1: Monday, June 23

Morning :	RT scheduling & resource management	(G. Buttazzo)
	RT kernels for embedded systems	(P. Gal)
Afternoon:	The FLEX development board	(M. Marinoni)
	Erika kernel and the OSEK standard	(P. Gai)

Day 2: Tuesday, June 24

Morning:	Developing RT appl. with Erika	(P. Gai – M. Marinoni)
Afternoon:	Laboratory practice	(P. Gai – M. Marinoni)

Day 3: Wednesday, June 25

Morning: Embedded Systems and Wireless Communication

Afternoon: Laboratory practice

(P. Pagano – G. Franchino) (P. Pagano – G. Franchino) $\frac{2}{2}$

Lecture Schedule

- 09:00 Morning Lecture Part 1
- 11:00 Coffee Break
- 11:15 Morning Lecture Part 2
- 13:00 Lunch Break
- 14:30 Afternoon Lecture Part 1
- 16:15 Coffee Break
- 16:30 Afternoon Lecture Part 2
- 18:00 End of Lectures

Real-Time Scheduling and Resource Management

Giorgio Buttazzo E-mail: buttazzo@sssup.it



Real-Time Systems Laboratory

Scuola Superiore Sant'Anna



Goal

Provide some background of RT theory for implementing control applications:

- Background and basic concepts
- Modeling real-time activities
- Real-Time Task Scheduling
- > Timing analysis
- Handling shared resources

What's an Embedded System?

⇒ It is a computing system hidden in an object to control its functions, enhance its performance, manage the available resources and simplify the interaction with the user.



What's special in Embedded Systems?

Stringent constraints on space, weight, energy, cost

- ⇒ Scarce resources (processing power, memory)
- ⇒ Efficient resource usage at the OS level
- Interaction with the environment
 - ⇒ High responsiveness and timing constraints
 - ⇒ Schedulability analysis and predictable behavior (RTOS)
- Robustness (tolerance to parameter variations)
 - ⇒ Overload management and system adaptation, to cope with variable resource needs and high load variations.

... and many others

- mobile robot systems
- small embedded devices
 - \Rightarrow cell phones
 - \Rightarrow videogames
 - \Rightarrow smart sensors
 - \Rightarrow intelligent toys









Criticality



Timing constraints



QoS management	High performance	Safety critical	
soft	firm	hard	
	Timing constraints		

Importance of the RTOS

The Operating System is responsible for:

- managing the available resources in an efficient way (memory, devices, energy);
- Enforcing timing constraints on computational activities;
- Providing a standard programming interface to develop portable applications.
- Providing suitable monitoring mechanisms to trace the system evolution to support debugging.

Software Vision

processor







Activation modes

Periodic tasks:

(time driven)

tasks are automatically activated by the kernel at regular time intervals:



Aperiodic tasks:

(event driven)

tasks are activated upon the arrival of an event (interrupt or explicit activation)



OS support for periodic tasks

task τ_i









Periodic task model

 $au_i(\Phi_i, C_i, T_i, D_i)$



$$\begin{array}{c} \mathbf{r}_{i,k} \ = \ \Phi_i + (k-1) \ T_i \\ \mathbf{d}_{i,k} \ = \ \mathbf{r}_{i,k} + \mathbf{D}_i \end{array} \left(\begin{array}{c} \text{often} \\ \Phi_i = \mathbf{0} \\ \mathbf{D}_i = \mathbf{T}_i \end{array} \right) \end{array}$$

Aperiodic task model

- Aperiodic: $r_{i,k+1} > r_{i,k}$
- Sporadic: $r_{i,k+1} \ge r_{i,k} + T_i$



Periodic Task Scheduling

• We have *n* periodic tasks: $\{\tau_1, \tau_2 \dots \tau_n\}$

 $\tau_i(C_i\,,\,T_i\,,\,D_i\,)$

<u>Goal</u>

- Execute all tasks within their deadlines
- Verify feasibility before runtime

Assumptions

- Tasks are execute in a single processor
- Tasks are independent (do not block or self-suspend)
- Tasks are synchronous (all start at the same time)
- Relative deadlines are equal to periods $(D_i = T_i)$

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

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.

Example



 $\Delta = \text{GCD} \quad (\text{minor cycle})$

T = lcm (major cycle)



Implementation



Timeline scheduling

Advantages

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

Timeline scheduling

Disadvantages

- It is not robust during overloads.
- It is difficult to expand the schedule.
- It is not easy to handle aperiodic activities.

Problems during overloads

What do we do during task overruns?

- Let the task continue
 - we can have a *domino effect* on all the other tasks (timeline break)
- Abort the task
 - the system can remain in inconsistent states.

Expandibility

If one or more tasks need to be upgraded, we may have to re-design the whole schedule again.



Expandibility

 We have to split task B in two subtasks (B₁, B₂) and re-build the schedule:



Expandibility

If the frequency of some task is changed, the impact can be even more significant:



Example



31

Priority Scheduling

Method

- Each task is assigned a priority based on its timing constraints.
- We verify the feasibility of the schedule using analytical techniques.
- Tasks are executed on a priority-based kernel.

Priority Assignments



 $D_i = T_i$

- Rate Monotonic (RM): $p_i \propto 1/T_i$ (static)
- Earliest Deadline First (EDF): $p_i \propto 1/d_i \quad (\text{dynamic}) \qquad d_{i,k} \ = \ r_{i,k} + D_i$

Rate Monotonic (RM)

 Each task is assigned a <u>fixed</u> priority proportional to its rate.



How can we verify feasibility?

Each task uses the processor for a fraction of time:

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

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

A necessary condition

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.

An unfeasible RM schedule

$$U_p = \frac{3}{6} + \frac{4}{9} = 0.944$$


Utilization upper bound





NOTE: If C_1 or C_2 is increased, τ_2 will miss its deadline!

A different upper bound



The upper bound U_{ub} depends on the specific task set.

The least upper bound



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.

Basic results



In 1973, Liu & Layland proved that a set of *n* periodic tasks can be feasibly scheduled



RM bound for large n

$$U_{\rm lub}^{RM} = n(2^{1/n} - 1)$$

for
$$n \to \infty$$
 $U_{lub} \to ln 2$

Schedulability bound



A special case

If tasks have harmonic periods $U_{lub} = 1$.









Schedule





RM Optimality

RM is **optimal** among all <u>fixed priority</u> algorithms:

If there exists a fixed priority assignment which leads to a feasible schedule for Γ , then the RM assignment is feasible for Γ .

If Γ is not schedulable by RM, then it cannot be scheduled by any fixed priority assignment.

EDF Optimality

EDF is **optimal** among <u>all</u> algorithms:

If there exists a feasible schedule for Γ , then EDF will generate a feasible schedule.

If Γ is not schedulable by EDF, then it cannot be scheduled by any algorithm.

Critical Instant

For any task τ_i , the longest response time occurs when it arrives together with all higher priority tasks.



The Hyperbolic Bound

 In 2000, Bini et al. proved that a set of n periodic tasks is schedulable with RM if:







Extension to tasks with D < T



Scheduling algorithms

- Deadline Monotonic: $p_i \propto 1/D_i$ (static)
- Earliest Deadline First: $\mathbf{p_i} \propto \mathbf{1/d_i}$ (dynamic)

Deadline Monotonic



Problem with the Utilization Bound

$$U_p = \sum_{i=1}^n \frac{C_i}{D_i} = \frac{2}{3} + \frac{3}{6} = 1.16 > 1$$

but the task set is schedulable.

How to guarantee feasibility?



• Fixed priority: Response Time Analysis (RTA)

• EDF: Processor Demand Criterion (PDC)

Response Time Analysis [Audsley '90]

• For each task τ_i compute the interference due to higher priority tasks:

$$I_i = \sum_{D_k < D_i} C_k$$

compute its response time as

$$R_i = C_i + I_i$$

• verify if $R_i \leq D_i$



Interference of τ_k on τ_i in the interval [0, R_i]:

$$I_{ik} = \left\lceil \frac{R_i}{T_k} \right\rceil C_k$$

Interference of high priority tasks on τ_i :

$$I_i = \sum_{k=1}^{i-1} \left\lceil \frac{R_i}{T_k} \right\rceil C_k$$

Computing the response time

$$R_i = C_i + \sum_{k=1}^{i-1} \left\lceil \frac{R_i}{T_k} \right\rceil C_k$$

Iterative solution:

$$\begin{cases} R_i^0 = C_i \\ R_i^s = C_i + \sum_{k=1}^{i-1} \left[\frac{R_i^{(s-1)}}{T_k} \right] C_k \end{cases}$$

iterate until $R_i^s > R_i^{(s-1)}$

Processor Demand Criterion [Baruah, Howell, Rosier 1990]

For checking the <u>existence</u> of feasibile schedule and for EDF

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

$$\forall t_1, t_2 > 0, \quad g(t_1, t_2) \leq (t_2 - t_1)$$



The demand in $[t_1, t_2]$ is the computation time of those jobs started at or after t_1 with deadline less than or equal to t_2 :

$$g(t_1, t_2) = \sum_{\substack{r_i \ge t_1}}^{d_i \le t_2} C_i$$

Processor Demand

For synchronous task sets we can only analyze intervals [0,L]



Processor Demand Test



Question

How can we bound the number of intervals in which the test has to be performed?



Bounding complexity

 Since g(0,L) is a step function, we can check feasibility only at deadline points.



 If tasks are synchronous and U_p < 1, we can check feasiblity up to the hyperperiod H:

$$H = Icm(T_1, \dots, T_n)$$

Bounding complexity

• Moreover we note that: $g(0, L) \leq G(0, L)$

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

$$= \sum_{i=1}^{n} L \frac{C_i}{T_i} + \sum_{i=1}^{n} (T_i - D_i) \frac{C_i}{T_i}$$

$$= LU + \sum_{i=1}^n (T_i - D_i)U_i$$

Limiting L



Processor Demand Test

A set of *n* periodic tasks with $D \le T$ is schedulable by EDF if and only if

$$U < 1$$
 AND $\forall L > 0$ $\sum_{i=1}^{n} \left\lfloor \frac{L - D_i + T_i}{T_i} \right\rfloor C_i \leq L$

$$D = \{d_k \mid d_k \le \min(H, L^*)\}$$

$$\begin{cases} H = lcm(T_{1}, ..., T_{n}) \\ \sum_{i=1}^{n} (T_{i} - D_{i})U_{i} \\ L^{*} = \frac{1}{1 - U} \end{cases}$$

Summarizing: RM vs. EDF		
	$D_i = T_i$	$D_i \le T_i$
RM	Suff.: polynomial $O(n)$ LL: $\Sigma U_i \le n(2^{1/n} - 1)$ HB: $\Pi(U_i+1) \le 2$ Exact pseudo-polynomial RTA	pseudo-polynomial Response Time Analysis \forall_i $R_i \leq D_i$ $R_i = C_i + \sum_{k=1}^{i-1} \left\lceil \frac{R_i}{T_k} \right\rceil C_k$
EDF	polynomial: $O(n)$ $\Sigma U_i \leq 1$	pseudo-polynomialProcessor Demand Analysis $\forall L > 0, g(0,L) \leq L$

Inter-task communication mechanisms

- Shared memory
- Message passing ports
- Asynchonous buffers

Handling shared resources

Problems caused by mutual exclusion

Critical sections


Blocking on a semaphore



It seems that the maximum blocking time for τ_1 is equal to the length of the critical section of τ_2 , but ...

Priority Inversion



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

Resource Access Protocols

Under fixed priorities

- Non Preemptive Protocol (NPP)
- Highest Locker Priority (HLP)
- Priority Inheritance Protocol (PIP)
- Priority Ceiling Protocol (PCP)

Under EDF

• Stack Resource Policy (SRP)

Non Preemptive Protocol

- Preemption is forbidden in critical sections.
- Implementation: when a task enters a CS, its priority is increased at the maximum value.

ADVANTAGES:	simplicity
PROBLEMS:	high priority tasks that do not use CS may also block

Conflict on critical section



Schedule with NPP



$$\mathbf{P}_{\mathrm{CS}} = \max\{\mathbf{P}_1, \dots, \mathbf{P}_n\}$$

Problem with NPP



τ_1 cannot preemt, although it could

Highest Locker Priority

A task in a CS gets the highest priority among the tasks that use it.

FEATURES:

- Simple implementation.
- A task is blocked when attempting to preempt, not when entering the CS.

Schedule with HLP



$$P_{CS} = \max \{P_k \mid \tau_k \text{ uses } CS\}$$

 τ_2 is blocked, but τ_1 can preempt within a CS

Problem with HLP



83

Priority Inheritance Protocol [Sha, Rajkumar, Lehoczky, 90]

- A task in a CS increases its priority only if it blocks other tasks.
- A task in a CS inherits the highest priority among those tasks it blocks.

$$P_{CS} = \max \{P_k \mid \tau_k \text{ blocked on CS}\}$$

Schedule with PIP



Types of blocking

• Direct blocking

A task blocks on a locked semaphore

• Push-through blocking

A task blocks because a lower priority task inherited a higher priority.

BLOCKING:

a delay caused by a lower priority task

Identifying blocking resources

- A task τ_i can be blocked by those semaphores used by lower priority tasks and
 - directly shared with τ_i (direct blocking) or
 - shared with tasks having priority higher than τ_i (push-through blocking).

Theorem: τ_i can be blocked at most once by each of such semaphores

Theorem: τ_i can be blocked at most once by each lower priority task

Bounding blocking times

- If **n** is the number of tasks with priority less than τ_{i}
- and **m** is the number of semaphores on which τ_i can be blocked, **then**

Theorem: τ_i can be blocked at most for the duration of **min(n,m)** critical sections



- τ_1 can be blocked once by τ_2 (on A₂ or C₂) and once by τ_3 (on B₃ or D₃)
- τ_2 can be blocked once by τ_3 (on B₃ or D₃)
- τ₃ cannot be blocked

Schedule with PIP



Remarks on PIP

ADVANTAGES

- It is transparent to the programmer.
- It bounds priority inversion.

PROBLEMS

 It does not avoid deadlocks and chained blocking.

Chained blocking with PIP



by each lower priority task

Priority Ceiling Protocol

- Can be viewed as PIP + access test.
- A task can enter a CS only if it is free and there is no risk of chained blocking.

To prevent chained blocking, a task may stop at the entrance of a free CS (*ceiling blocking*).

Resource Ceilings

Each semaphore s_k is assigned a ceiling:

$$C(s_k) = max \{P_j: \tau_j \text{ uses } s_k\}$$

• A task τ_i can enter a CS only if

$$P_i > max \{C(s_k) : s_k \text{ locked by tasks} \neq \tau_i\}$$



 t_1 : τ_2 is blocked by the PCP, since $P_2 < C(s_1)$

PCP properties

Theorem 1

Under PCP, each task can block at most once.

Theorem 2

PCP prevents chained blocking.

Theorem 3

PCP prevents deadlocks.

Remarks on PCP

ADVANTAGES

- Blocking is reduced to <u>only one</u> CS
- It prevents deadlocks

PROBLEMS

 It is not transparent to the programmer: semaphores need ceilings

Typical Deadlock



Deadlock avoidance with PCP



Guarantee with resource constraints

- We select a scheduling algorithm and a resource access protocol.
- We compute the maximum blocking times (B_i) for each task.
- We perform the guarantee test including the blocking terms.



$$\forall i \quad \sum_{k=1}^{i-1} \frac{C_k}{T_k} + \frac{C_i + B_i}{T_i} \leq i (2^{1/i} - 1)$$



$$R_i = C_i + B_i + \sum_{k=1}^{i-1} \left[\frac{R_i}{T_k} \right] C_k$$

Resource Sharing under EDF

The protocols analyzed so far have been originally developed for fixed priority scheduling schemes. However:

- NPP can also be used under EDF
- PIP has been extended under EDF by Spuri (1997).
- PCP has been extended under EDF by Chen-Lin (1990)
- In 1990, Baker proposed a new protocol that works both under fixed and dynamic priorities.

Stack Resource Policy [Baker 1990]

- It works both with fixed and dynamic priority
- It limits blocking to 1 critical section
- It prevents deadlock
- It supports multi-unit resources
- It allows stack sharing
- It is easy to implement

Stack Resource Policy [Baker 90]

- For each resource R_k:
 - \Rightarrow Maximum units: N_k
 - \Rightarrow Available units: n_k



• For each task τ_i the system keeps:



Stack Resource Policy [Baker 90]

Resource ceiling

$$C_k(n_k) = \max_j \left\{ \pi_j : n_k < \mu_j(R_k) \right\}$$

System ceiling

$$\Pi_s = \max_k \{C_k(n_k)\}$$

SRP Rule

A job cannot preempt until p_i is the highest and $\pi_i > \prod_s$

Computing Resource Ceilings



Computing Resource Ceilings





Schedule with SRP



A task blocks when attempting to preempt
Schedule with PCP



A task is blocked when accessing a resource

Lemma

If $\pi_i > C_R(n_k)$ then there exist enough units of R

- 1. to satisfy the requirements of τ_i
- 2. to satisfy the requirements of all tasks that can make preemption on τ_i

Theorem 1

Under SRP, each task can block at most once.

Consider the following scenario where τ_1 blocks twice:



This is not possible, because τ_2 could not preempt τ_3 because, at time t*, $\pi_2 < \prod_s$

Theorem 2

If $\pi_i > \prod_s$ then τ_i will never block once started.

Proof

Since $\Pi_s = \max\{C_R(n_k)\}$, then there are enough resources to satisfy the requirements of τ_i and those of all tasks that can preempt τ_i .

Question

If a task can never block once started, can we get rid of the wait / signal primitives?

Theorem 3

SRP prevents deadlocks.

Proof

From Theorem 2, if a task can never block once started, then no deadlock can occur.

Deadlock avoidance with SRP



Schedulability Analysis under EDF

When $D_i = T_i$



 B_i can be computed as under PCP and refers to the length of longest critical section that can block τ_i .

EDF Guarantee: PD test $(D_i \leq T_i)$ τ_1 τ_i τ_k τ_n

Tasks are ordered by decreasing preemption level

Schedulability Analysis under EDF

When $D_i \leq T_i$

A task set is schedulable if U < 1 and $\forall L \in D$

$$\forall i \quad B_i + \sum_{k=1}^n \left\lfloor \frac{L + T_k - D_k}{T_k} \right\rfloor C_k \leq L$$

where
$$D = \{d_k \mid d_k \le \min(H, L^*)\}$$

 $H = lcm(T_1, ..., T_n)$ $L^* = \frac{\sum_{i=1}^n (T_i - D_i)U_i}{1 - U}$
118

Stack Sharing

Each task normally uses a private stack for

- saving context (register values)
- managing functions
- storing local variables



Stack Sharing

Why stack cannot be normally shared?

Suppose tasks share a resource:





Stack Sharing

Why stack can be shared under SRP?



Saving Stack Size

To really save stack size, we should use a small number of preemption levels.



NOTE on SRP

SRP for fixed priorities and single-unit resources is equivalent to Higher Locker Priority.

It is also referred to as Immediate Priority Ceiling



Non-preemtive scheduling

It is a special case of preemptive scheduling where all tasks share a single resource for their entire duration.



The max blocking time for task τ_i is given by the largest C_k among the lowest priority tasks:

$$B_i = max\{C_k : P_k < P_i\}$$

Advantages of NP scheduling

- Reduces runtime overhead
 - Less context switches
 - > No semaphores are needed for critical sections
- Reduces stack size, since no more than one task can be in execution.
- Preserves program locality, improving the effectiveness of
 - Cache memory
 - Pipeline mechanisms
 - Prefetch queues

Advantages of NP scheduling

- As a consequence, task execution times are
 - Smaller
 - More predictable



Advantages of NP scheduling

In fixed priority systems can improve schedulabiilty:



Disadvantages of NP scheduling

- In general, NP scheduling reduces schedulability.
- The utilization bound under non preemptive scheduling drops to zero:



Non preemptive scheduling anomalies





129

Trade-off solutions

Preemption thresholds

Each task has two priorities:

- Nominal priority (ready priority): used to enqueue the task in the ready queue
- Threshold priority: used for task execution



Preemption thresholds

- Nominal pr. = threshold: \Rightarrow fully preemptive
- \Rightarrow fully non preemptive • Threshold = P_{max}



Trade-off solutions

Tunable Preemptive Systems

- Compute the longest non-preemptive section that allows a feasible schedule.
- Allow preemption only in certain points in the code.

