

Lecture 12: An Overview of Scheduling Theory

[RTCS Ch 8]

- Introduction
- Execution Time Estimation
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - Earliest Deadline Scheduling
 - Scheduling of Aperiodic Tasks
- Alternative Scheduling Models
 - Reservation-Based Scheduling

1

Goal

Question to be answered:

- *How can we guarantee that a set of tasks meet their deadlines?*

2

Problem Formulation

Events

- events occur that require computations (interrupts)
- a task executes a piece of code in response to an event
- aperiodic (sporadic) events and periodic events

Worst-Case Execution Time

- an upper bound on the CPU time it takes to execute the task without any interfering tasks (alone on the CPU)

Deadline

- Maximum allowed time when the task should be completed

Scheduling

- the choice of which event to process at a given time, i.e. which task to execute

3

Schedulability Analysis

- For hard real-time systems the deadlines must always be met
- Off-line analysis (before the system is started) required to check so that there are no circumstances that could lead to missed deadlines
- A system is *unschedulable* if the scheduler will not find a way to switch between the tasks such that the deadlines are met
- The analysis is *sufficient* if, when it answers "Yes", all deadlines will be met
- The analysis is *necessary* if, when it answers "No", there really is a situation where deadlines could be missed
- The analysis is *exact* if it is both sufficient and necessary
- A sufficient analysis is an absolute requirement and we like it to be as close to necessary as possible

4

- Introduction
- **Execution Time Estimation**
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - Earliest Deadline Scheduling
 - Scheduling of Aperiodic Tasks
- Alternative Scheduling Models
 - Reservation-Based Scheduling

5

Execution Time Estimation

Basic Question:

- "How much CPU time does this piece of code need?"

Two major approaches:

1. Measuring execution times
2. Analyzing execution times

6

Measuring Execution Times

- the code is compiled and run with measuring devices (e.g. logical analyzer) connected, or ..
- ..., the OS provide execution time measurements
- a large set of test input data is used
- longest time required = longest time measured (+ safety margin)

General problem:

- No guarantees that we really have encountered the longest execution time

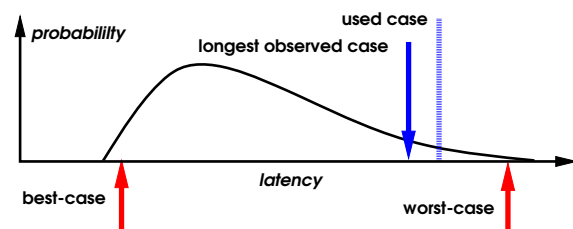
7

Problems:

- execution times are data dependent (e.g. a sensor reading)
- caching
 - memories have different speeds
 - a memory reference causing a cache miss takes much longer time than a reference inside the cache
- pipelining & speculative execution
- memory accesses for multiprocessor systems
- testing a real-time problem is difficult and time consuming
- garbage collection

Main Problem: No guarantees

8



9

Analyzing Execution Times

Aim:

- a tool that takes the source code and automatically and formally correct decides the longest execution time
- research area for the last 10-15 years

Problems:

- compiler dependent
 - different compilers generate different code
 - Remedy: work with the machine code

10

Approach:

- use the instruction time tables from the CPU manufacturer
- add up the instruction times of the individual statements

Problem:

- branching statements (IF, CASE)
 - how should we know which code that is executed

```

IF X > 5 THEN      MOVE (_X),D0      B1
  X := X + 1;      CMP D0,#5
ELSE              BGT L1
  X := X * 3;      MUL D0,#3
ENDIF;           MOVE D0      B2
                JMP L2
                L1: ADD D0,#1  B3
                L2: ...
IF B1 THEN       B3
ELSE            B2
ENDIF
    
```

Longest execution time = time(B1) + max(time(B2),time(B3))

11

Execution times of the basic blocks:

Operation	Number of CPU cycles
MOVE	8
CMP	4
BGT	4
MUL D0,#3	16 + 2 times # '1's
JMP	4
ADD	4

time(B1) = 8 + 4 + 4 = 16 cycles

time(B2) = 48 + 8 + 4 = 60 cycles (word length = 16 bits)

time(B3) = 4 cycles

⇒ time(if-statement) = 76 cycles

8MHz clock frequency ⇒ 1 cycle takes 125ns

⇒ time(if-statement) = 76 * 125ns = 9.5µs

12

Extended to more complex statements

```
IF X = 0 THEN
  IF X > 5 THEN
    X := X + 1;
  ELSE
    X := X * 3;
  ENDIF;
ELSE
  X := 1;
ENDIF;
```

13

Problems:

- Loops (IF, WHILE, ..)
 - How should we know how many times the code will loop?

```
WHILE X > 5 DO
  X := X - 1;
END;
```
 - Remedy: the programmer must annotate the source code with the maximum number of times the loop executes
- Recursion
 - difficult to know beforehand how deep the recursive call can get
 - Remedy: recursion not allowed
- allocation of dynamic memory
 - the time for the memory management often unknown
 - difficult for an analysis tool to handle

14

- goto statements
 - the data flow in a piece of code is difficult to work out
- caches and multi-threaded applications
 - caches with single-threaded applications can be handled reasonably well
 - caches with multi-threaded applications extremely pessimistic
 - * each context switch may cause a cache miss
- Main problem: *pessimism*
 - the actual longest execution time may be substantially smaller than what the analyzer says
 - however, if we want formal guarantees the analytical approach is the only choice

15

WCET Analysis Tools

Three phases:

1. Flow Analysis
 - calculates all possible execution paths in the program
 - in order to limit the number of times the instructions can be executed
2. Low-level Analysis
 - calculates the execution time of the different instructions on the given hardware
3. WCET Calculation
 - Combine step 1 and 2

For the uni-processor case with simple cache structures and without complex pipelines the obtained results is typically only 10-15 % larger than the true WCET for single-threaded applications.

However, for multi-threaded applications either on a uniprocessor or a multicore platform the pessimism is much larger.

16

- Introduction
- Execution Time Estimation
- **Basic Scheduling Approaches**
 - **Static Cyclic Scheduling**
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - Earliest Deadline Scheduling
 - Scheduling of Aperiodic Tasks
- Alternative Scheduling Models
 - Reservation-Based Scheduling

17

The Critical Instant

It can be shown that, in the uni-processor case, the worst situation, from a schedulability perspective, occurs when all tasks want to start their execution at the same time instant.

This is known as the *critical instant*.

If we can show that the task set is schedulable in this situation, it will also be schedulable in other situations.

If we can show that the task set is schedulable for the worst case execution times, then the task set will also be schedulable if the actual execution times are shorter.

Hence, all uni-processor scheduling analysis only need to check for this case.

18

Static Cyclic Scheduling

- off-line approach
- configuration algorithm generates an execution table or calendar
- many different algorithms (optimization)
- the table repeats cyclically \Rightarrow static cyclic scheduling
- works for both non-preemptive and preemptive scheduling
- the run-time dispatcher simply follows the table
 - sets up an hardware interrupt at the time when a context switch should be performed (preemptive)
 - starts the first task in the calendar
 - when the hardware interrupt arrives the first task is preempted and next task is run,
 - ...

19

Analysis:

- trivial, run through the table and check that all timing requirements are met

Limitations:

- can only handle periodic tasks
 - aperiodic tasks are made periodic through polling
- the calendar cannot be too large
 - shortest repeating cycle = the *least common multiple, LCM* of the task periods
 - periods 5,10,20 ms gives cycle of 20 ms
 - periods 7,13,23 ms gives cycle of 2093 ms
 - periods are made shorter than they need to be to reduce the calendar

20

- building a schedule is NP-Hard
 - we cannot expect an algorithm to always find a schedule even if one exists
 - good heuristic algorithms exist that can mostly find a solution if one exists

21

Advantages:

- A number of different task constraints can be handled
 - Exclusion constraints can be handled
 - Precedence constraints can be handled

Disadvantages:

- Inflexible
 - static design

22

Notation

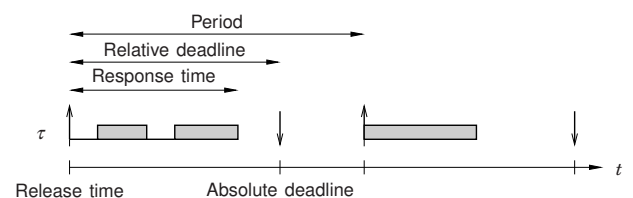
Notation	Description
C_i	Worst-case execution time of task i
T_i	Period of task i
D_i	Relative deadline of task i

CPU utilization U :

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

23

Notation



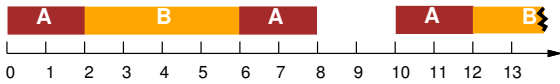
24

Example

Task name	T	D	C
A	5	5	2
B	10	10	4

Utilization: $2/5 + 4/10 = 0.8$

Schedule length: $\text{LCM}(5,10) = 10$



Worst case response time for task A, $R_A = 3 < D_A$

Worst case response time for task B, $R_B = 6 < D_B$

25

Implementation

```

CurrentTime(t);
LOOP
  A();
  B();
  A();
  IncTime(t,10);
  WaitUntil(t);
END;

```

Problem: Assume it only takes 2 time units to execute task B. Then task A will start before it should do.

26

Better implementation

```

CurrentTime(t);
LOOP
  A();
  IncTime(t,2);
  WaitUntil(t);
  B();
  IncTime(t,4);
  WaitUntil(t);
  A();
  IncTime(t,4);
  WaitUntil(t);
END;

```

27

- Introduction
- Execution Time Estimation
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - **Fixed Priority Scheduling**
 - * **Rate Monotonic Analysis**
 - Earliest Deadline Scheduling
 - Scheduling of Aperiodic Tasks
- Alternative Scheduling Models
 - Reservation-Based Scheduling

28

Fixed Priority Scheduling

- each task has a fixed priority
- the dispatcher selects the task with the highest priority
- preemptive
- used in most r-t kernels and RTOS

29

Rate Monotonic Priority Assignment

- a scheme for assigning priorities to processes
- priorities are set monotonically with rate (period)
- a task with a shorter period is assigned a higher priority
- introduced in

C.L Liu and J.W Layland, *Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment*, JACM, Vol. 20, Number 1, 1973

30

Rate Monotonic Analysis

Assumptions needed = model

Model:

- periodic tasks
- $D_i = T_i$
- tasks are not allowed to be blocked or suspend themselves
- priorities are unique
- task execution times bounded by C_i
- interrupts and context switches take zero time

31

Result:

If the task set has a utilization below a utilization bound then all deadlines will be met

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

Sufficient condition (if the utilization is larger than the bound the task set may still be schedulable)

As $n \rightarrow \infty$, the utilization bound $\rightarrow 0.693$

"If the CPU utilization is less than 69%, then all deadlines are met"

Alternative tighter test (Hyperbolic Bound):

$$\prod_{i=1}^{i=n} \left(\frac{C_i}{T_i} + 1 \right) \leq 2$$

32

Generalized Rate Monotonic Analysis

Since 1973 the models have become more flexible and the analysis better

M. Joseph and P. Pandaya, *Finding Response Times in a Real-Time System*, The Computer Journal, Vol. 29, No. 5, 1986

Notation:

Notation	Description
C_i	Worst-case execution time of task i
T_i	Period of task i
D_i	Relative deadline of task i
R_i	Worst-case response time of task i

Scheduling test: $R_i \leq D_i$ (necessary and sufficient)

Model:

- $D_i \leq T_i$

33

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

where $hp(i)$ is the set of tasks of higher priority than task i .

The function $\lceil x \rceil$ is the *ceiling function* that returns the smallest integer $\geq x$.

Recurrence relation, solved by iteration. The smallest solution is searched for.

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

Start with $R_i^0 = 0$

34

Example

Task set:

Task name	T	D	C	Priority
A	52	52	12	low
B	40	40	10	medium
C	30	30	10	high

Original (approximative) analysis:

$$\sum_{i=1}^{i=3} \frac{C_i}{T_i} = 0.814$$

$$3(2^{1/3} - 1) = 0.7798$$

Not schedulable

35

Hyperbolic bound:

$$\prod_{i=1}^{i=3} \left(\frac{C_i}{T_i} + 1 \right) = 2.0508$$

Not schedulable

36

Exact analysis:

$$\begin{aligned}
 R_A^0 &= 0, R_A^1 = C_A = 12, \\
 R_A^2 &= C_A + \left\lceil \frac{12}{T_B} \right\rceil C_B + \left\lceil \frac{12}{T_C} \right\rceil C_C = C_A + C_B + C_C = 32 \\
 R_A^3 &= \dots = 42, R_A^4 = \dots = 52, R_A^5 = \dots = 52 \\
 R_B^0 &= 0, R_B^1 = C_B = 10, \\
 R_B^2 &= C_B + \left\lceil \frac{10}{T_C} \right\rceil C_C = 20, \\
 R_B^3 &= \dots = 20 \\
 R_C^0 &= 0, R_C^1 = C_C = 10, R_C^2 = C_C = 10
 \end{aligned}$$

Task name	T	D	C	Priority	R
A	52	52	12	low	52
B	40	40	10	medium	20
C	30	30	10	high	10

$R_i \leq D_i \Rightarrow$ schedulable

37

Derivation of exact formulae

Task C has highest priority \rightarrow will not be interrupted and hence $R_C = C_C = 10$ (R_C^1)

Task B has medium priority. The response time will be at least equal to $C_B = 10$ (R_B^1). During that time B will be interrupted once by C. Hence, the response time will be extended by the execution time of C, i.e. $R_B^2 = 10 + 10 = 20$. During this time B will only be interrupted once by C and that has already been accounted for, i.e. $R_B^3 = 20$.

Task A has lowest priority. The response will be at least equal to $C_A = 12$ (R_A^1). During that time A will be interrupted once by C and once by B, i.e., $R_A^2 = 12 + 10 + 10 = 32$. During this time A will be interrupted twice by C and once by B, i.e., $R_A^3 = 32 + 10 = 42$. During this time A will be interrupted twice by C and twice by B, i.e., $R_A^4 = 42 + 10 = 52$. During this time no more unaccounted for interrupts will occur, i.e., $R_A^5 = 52$.

38

Limitation of the Exact Formula

If the response time is larger than the deadline then the quantitative value cannot be trusted

- Reason: The analysis does not take interference from previous jobs of the same task into account
- More advanced analysis exists

However, one still knows that the deadline won't be met, which is normally what one is interested in.

39

Best-Case Response Time

Under rate-monotonic priority assignment one can also calculate the best-case response time R_i^b of a task i .

$$R_i^b = C_i^{\min} + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i^b - T_j}{T_j} \right\rceil_0 C_j^{\min}$$

where C_i^{\min} is the best-case execution time of the task and $\lceil x \rceil_0 = \max(0, \lceil x \rceil)$.

Can be used to calculate the worst-case input-output latency of a control task.

40

Deadline Monotonic Scheduling

The rate monotonic policy is not very good when $D \leq T$.

An infrequent but urgent task would still be given a low priority.

The *deadline monotonic* ordering policy works better.

A task with a short relative deadline D gets a high priority.

This policy has been proved optimal when $D \leq T$ (if the system is unschedulable with the deadline monotonic ordering then it is unschedulable with *all* other orderings).

With $D \leq T$ we can control the jitter in control delay.

41

Deadline Monotonic Scheduling - Sufficient Condition

For a system with n tasks, all tasks will meet their deadlines if the total utilization of the system is below a certain bound.

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

42

Deadline Monotonic Scheduling - Exact Analysis

The response time calculations from the rate monotonic theory is also applicable to deadline monotonic scheduling.

Response time calculations does not make any assumptions on the priority assignment rule.

43

Extension: The Blocking Problem

How should interprocess communication be handled.

The analysis up to now does not allow tasks to share data under mutual exclusion constraints (e.g. no semaphores or monitors)

Main problem:

- a task i might want to lock a semaphore, but the semaphore might be held by a lower priority task
- task i is blocked

44

The *blocking factor*, B_i is the longest time a task i can be delayed by the execution of lower priority tasks

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

Priority inversion may cause unbounded blocking time if ordinary locks are used.

Different locking schemes have different blocking times.

- ordinary priority inheritance
- priority ceiling protocol
- immediate inheritance protocol

45

Further Extensions

- Release Jitter
 - the difference between the earliest and latest release of a task relative to the invocation of the task
- Context Switch Overheads
- Clock Interrupt Overheads
- Distributed systems using CAN

46

Overrun Behaviour - Fixed Priorities

Overrun = exceeding the worst-case execution time

Will only affect the current task and lower priority tasks

These will miss deadlines or, in the worst case, not get any execution time at all

Higher priority tasks will be unaffected.

47

- Introduction
- Execution Time Estimation
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - **Earliest Deadline Scheduling**
 - Scheduling of Aperiodic Tasks
- Alternative Scheduling Models
 - Reservation-Based Scheduling

48

Earliest Deadline First (EDF) Scheduling

- dynamic approach: all scheduling decisions are made on-line by the dispatcher
- the task with the smallest absolute deadline runs
- preemptive
- ready-queue sorted in deadline order
- "dynamic priorities"
- more intuitive to assign deadlines to tasks than to assign priorities
 - requires only local knowledge

49

Analysis:

- Simplest model:
 - periodic tasks
 - each task i has a period T_i ,
 - a worst-case computation time requirement C_i , and
 - a relative deadline D_i
 - $D_i = T_i$
 - independent task execution
 - ideal kernel

50

Result:

If the utilization U of the system is not more than 100% then all deadlines will be met.

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

Necessary and sufficient condition

Advantage: Processor can be fully used.

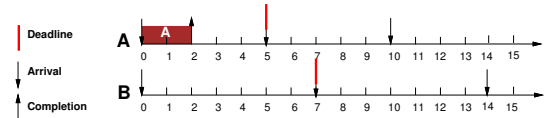
Less restrictive assumptions make the analysis harder (see RTCS for the analysis in the case $D_i \leq T_i$.)

51

Example

Task name	T	D	C
A	5	5	2
B	7	7	4

Utilization: $2/5 + 4/7 = 0.971$

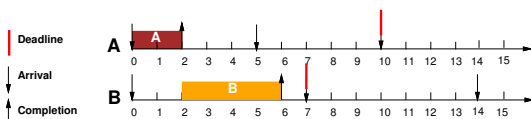


52

Example

Task name	T	D	C
A	5	5	2
B	7	7	4

Utilization: $2/5 + 4/7 = 0.971$

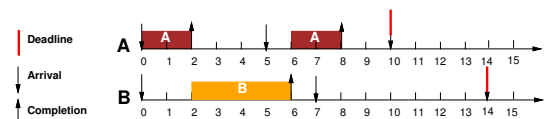


53

Example

Task name	T	D	C
A	5	5	2
B	7	7	4

Utilization: $2/5 + 4/7 = 0.971$

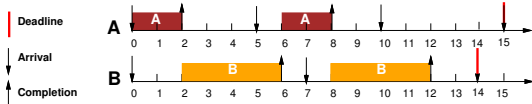


54

Example

Task name	T	D	C
A	5	5	2
B	7	7	4

Utilization: $2/5 + 4/7 = 0.971$

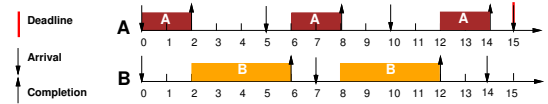


55

Example

Task name	T	D	C
A	5	5	2
B	7	7	4

Utilization: $2/5 + 4/7 = 0.971$



56

Overrun Behaviour

In the case of overrun all tasks will be affected, i.e., all tasks may miss deadlines.

The "Domino effect"

However, in general EDF is more fair than priority-based scheduling

- the available resources will be distributed among all the tasks

57

EDF: Summary

Also for EDF there exists a very well-developed schedulability theory

Resource access protocols similar to priority inheritance and ceiling.

58

Scheduling of Aperiodic Tasks 1

Soft aperiodic tasks combined with hard periodic tasks:

Server techniques:

- Main idea: use a special high-priority server task for scheduling the pending aperiodic work
- the server has time tickets (a budget) that can be used to schedule and execute the aperiodic tasks
- if there is aperiodic work pending and the server has unused tickets, the aperiodic tasks execute until they finish or the available tickets are exhausted.
- several servers proposed, e.g., priority exchange server, deferrable server, sporadic server, slack stealer, constant bandwidth server
- the main differences concern how the capacity of the server is replenished and the maximum capacity of the server
- schedulability results available

59

60

- Introduction
- Execution Time Estimation
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - Earliest Deadline Scheduling
 - **Scheduling of Aperiodic Tasks**
- Alternative Scheduling Models
 - Reservation-Based Scheduling

Scheduling of Aperiodic Tasks 2

Hard aperiodic tasks:

- on-line admission control
- if an aperiodic task is accepted we should be able to guarantee that it finishes before its deadline

Using deadline-monotonic scheduling all invocations of aperiodic tasks will meet their deadlines if

$$\forall t : U(t) \leq UB(n),$$

where $UB(n) = \frac{1}{2} + \frac{1}{2n}$ for $n < 3$ and

$$UB(n) = \frac{1}{1 + \sqrt{\frac{1}{2}(1 - \frac{1}{n-1})}}$$

for $n \geq 3$. ($\rightarrow 58.6\% n \rightarrow \infty$)

$U(t)$ = the synthetic utilization (utilization), $U(t) = \sum(C_i/D_i)$ where the sum extends over all tasks that have arrived and been accepted but whose deadlines have not yet expired.

61

- Introduction
- Execution Time Estimation
- Basic Scheduling Approaches
 - Static Cyclic Scheduling
 - Fixed Priority Scheduling
 - * Rate Monotonic Analysis
 - Earliest Deadline Scheduling
 - Scheduling of Aperiodic Tasks
- **Alternative Scheduling Models**
 - **Reservation-Based Scheduling**

62

Alternative Scheduling Models

The Multi-Frame Model:

- execution times and deadlines for tasks are allowed to vary from job to job according to a periodic pattern
- sufficient and necessary schedulability conditions
- Mok & Chen and Baruah

Sub-Task Scheduling:

- each task is divided into serially executed subtasks each characterized by an execution time, a priority, and a deadline
- in the recurrent task model is also possible to model branching in the execution of the subtasks (Baruah, Gonzalez Harbour)

63

Offset-Based Scheduling:

- the periodic task model is based on the assumptions that all tasks may arrive simultaneously
- if the is can be avoided the schedulability of a task set increases
- using task offsets task can be shifted in time
- schedulability analysis (Gutierrez and Harbour)

64

Reservation-Based Scheduling

If a tasks overruns (executes longer than anticipated) this will effect other tasks negatively

- In priority-based systems the priority decides which tasks that will be effected
- In deadline-based systems all tasks will be effected

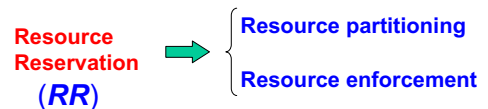
We want to provide temporal protection between tasks that guarantees that a certain task or group of tasks receives a certain amount of the CPU time.

Cp. Functional protection provided by the memory management unit in conventional OS (and in some RTOS)

65

Explicit Resource Reservation

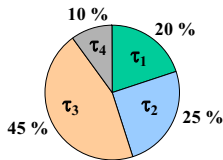
A much better approach is to use a programming model that enables the designer to [explicitly control the resources](#) assigned to a given activity.



6

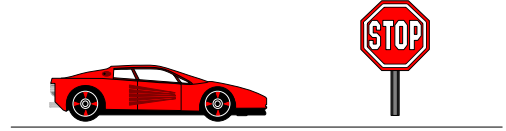
Resource Partitioning

- Partition each resource among system activities, based on performance and execution requirements:



Each task receives a fraction U_i of the processor and behaves as if it were executing alone on a slower processor of speed U_i

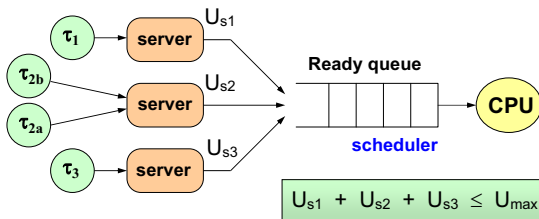
Resource Enforcement



- It is a mechanism that makes sure that each task does not consume more than its reserved amount.
- If a task attempts to execute more, it must be delayed, preserving the resource for the other tasks.

8

Implementation



	server	scheduler
Fixed priorities	Cap. Res.	RM
Dynamic priorities	CBS	EDF

9

Reservation-Based Scheduling: Static & Dynamic

Use static cyclic scheduling for some tasks and let the other tasks be priority-based (event-based) which only may execute during the idle periods of the static schedule

Used in Rubus from Arcticus

- Swedish RTOS used by Volvo
- red threads - statically scheduled
- blue threads - dynamically scheduled

66

Reservation-based Scheduling: Priority-based system

How can, conceptually, a reservation-based scheduling system be implemented on top of ordinary priority-based scheduling

Each task or task set receives a certain percentage of the CPU. (50% + 30% + 20%)

Can be viewed as if the tasks are executed on a correspondingly much slower CPU

Two variants:

- Each task set gets exactly its share of the CPU
- Each task set gets at least its share of the CPU

Question: Over which time horizon does the CPU reservation hold

67

Reservation-based Scheduling: Priority-based system 1

Each task set gets exactly its share of the CPU

The scheduler can be viewed as consisting of as many ready-queues as there are reservation sets.

An external timer is set up to generate interrupts when it is time to switch which ready-queue that is active

One idle process in each ready-queue

68

Reservation-based Scheduling: Priority-based system 2

Each task set gets at least its share of the CPU

One ready-queue.

Make sure that the tasks belonging to the currently serviced task set all have higher priority than the tasks in the tasks sets which are not serviced

An external timer is set up to generate interrupts when it is time to switch between the tasks sets.

Lower the priorities of the the tasks that have been serviced and raise the priorities of the tasks that should be serviced

A single idle task

69

Reservation-based Scheduling: Industrial Practice

Beginning to emerge in commercial RTOS

Integrity from Green Hills Software Inc

70

More Information

A nice introduction and overview of the state-of-the-art in uni-processor scheduling of real-time systems can be found in:

- "Real Time Systems by Fixed Priority Scheduling" by Ken Tindell and Hans Hansson, Dept. of Computer Systems, Uppsala University
- <http://www.docs.uu.se/hansh/fpsnotes-9710.ps>

71