

0

NICTA

Lecture Content

- · Definition of Real-Time Systems (RTS)
- Scheduling in RTS
- Schedulability Analysis
- Worst Case Execution Time Analysis

No: 2

No: 3

- · Time and Distributed RTS
- · Rate Based Scheduling

© NICTA 2007/2008

© NICTA 2007/2008







© NICTA 2007/2008

No: 6

1

NICTA



- Fast context switches?
 should be fast anyway
- Small size?
- should be small anyway
- Quick response to external triggers?
 not necessarily quick but predictable
- Multitasking?
- often used, but not necessarily
- "Low Level" programming interfaces?
- might be needed as with other embedded systemsHigh processor utilisation?
- desirable in any system (avoid oversized system)

- The computation is obsolete if the job is not finished on time.
- Cost may be interpreted as loss of revenue.
- Typical example are forecast systems.







0

NICTA

| ade | 140 | dal | | |
|-----|-----|-----|--|--|

- Periodic tasks
 - Time-driven. Characteristics are known a priori
 - Task t_i is characterized by (T_i, C_i)
 - E.g.: Task monitoring temperature of a patient in an ICU.
- Aperiodic tasks
 - Event-driven. Characteristics are not known a priori
 - Task $\pmb{\tau}_i$ is characterized by (C_i, D_i) and some probabilistic profile
 - for arrival patterns (e.g. Poisson model) - E.g.: Task activated upon detecting change in patient's condition.
- · Sporadic Tasks
 - Aperiodic tasks with known minimum inter-arrival time $(\mathsf{T}_i,\,\mathsf{C}_i)$

No: 21

© NICTA 2007/

Preemption NICTA Why preemptive scheduling is good:

- It allows for shorter response time of high priority tasks
- As a result it is likely to allow for a higher utilisation of the processor before the system starts missing deadlines
- · Why preemptive scheduling is bad:
 - It leads to more task switches then necessary
 - The overheads of task switches are non-trivial
 - The system becomes harder to analyse whether it is able to meet all its deadlines
 - Preemption delay (cache refill etc.) becomes more expensive with modern processors

© NICTA 200 No: 24

Preemption Fixed Priority

- · Cooperative preemption?
 - Applications allow preemption at given points
 - Reduction of preemptions
 - Increase of latency for high priority tasks

 Fixed Priority Scheduling(FPS)
 Image: Constraint of the second scheduling s

| | FIIOIIty | 0 | 1 | U |
|---------------|----------|----|----|----|
| Task τ_1 | 1 | 5 | 20 | 20 |
| Task τ_2 | 2 | 8 | 30 | 20 |
| Task τ_3 | 3 | 15 | 50 | 50 |

No: 28

0

NICTA

| NICTA 2007/2008 | No: 25 |
|-----------------|--------|

© NICTA 2007/2008

Event Triggered Systems

"... The asynchronous design of the [AFTI-F16] DFCS introduced a random, unpredictable characteristic into the system. The system became untestable in that testing for each of the possible time relationships between the computers was impossible. This random time relationship was a major contributor to the flight test anomalies. Adversely affecting testability and having only postulated benefits, asynchronous operation of the DFCS demonstrated the need to avoid random, unpredictable, and uncompensated design characteristics."

D. Mackall, flight-test engineer AFTI-F16 AFTI-F16 flight tests

No: 26

| NUCT | |
|------|---|
| NICI | • |

0.

NICTA

Earliest Deadline First (EDF)

- · Dynamic priorities
- · Scheduler picks task, whose deadline is due next
- Advantages:

٠

- Optimality
- Reduces number of task switches
- Optimal if system is not overloaded
- · Drawbacks:
 - Deteriorates badly under overload
 - Needs smarter scheduler
 - Scheduling is more expensive

© NICTA 2007/2008

No: 29



Fixed Priority Scheduling

- · Priorities may assigned by
 - Deadline: shortest deadline \Rightarrow highest priority
 - Period: shortest period \Rightarrow highest priority
 - "Importance"
- Scheduler picks from all ready tasks the one with the highest priority to be dispatched.

No: 27

Benefits:

© NICTA 2007/2008

- Simple to implement
- Not much overhead
- Minimal latency for high priority tasks
- Drawbacks
- Inflexible
- Suboptimal (from analysis point of view)

© NICTA 2007/2008



No: 31







Time Triggered/Driven Scheduling

· Mostly static scheduling

© NICTA 2007/2008

- Time triggered scheduling allows easier reasoning and monitoring of response times
- · Can be used to avoid preemption
- Can be used in event triggered systems, but increases greatly the latency
- · Most often build around a base rate
- Can be implemented in big executive, using simple function calls

No: 33



- Caused by faulty components of the system

 Babbeling idiot or
 - A receiver part erroneously "receiving input"
 - EMI

© NICTA 2007/2008

0.

NICTA

- Or caused by wrong assumptions regarding the embedding environment
 - Basically wrong event rates or event correlation





No: 37







No: 41

- · Similar case for server tasks.
- · Pathfinder example

© NICTA 2007/2008





Non-Preemptable Critical Section

• GOOD

- Simple
- No deadlock
- No unbounded priority inversion
- No prior knowledge about resources.
- Each task blocked by at most 1 task of lower priority
- Works with fixed and dynamic priorities. (especially good for short critical sections with high contention)

No: 43

0.

NICTA

0.

NICTA

• BAD

© NICTA 2007/2008

- Tasks blocked even when no contention exists.

Basic Priority Ceiling Protocol 0 NICTA Task τ_1 Task r_2 Task $\mathbf{\tau}_3$ Task 🗗 Task τ_5 © NICTA 2007/2008 No: 46



Basic Priority Ceiling Protocol

- Lower priority task inherits priority of blocked task.
- Task may be denied resource even when available.
- Also known as Original Priority Ceiling Protocoll (OPCP) •

NICTA

0.

 \sum^{n} usage (k,i)C(k)

Priority Inheritance

- GOOD .
- No deadlock.
 - No unbounded priority inversion.
 - Blocking time reduced.
- BAD
 - Task may be denied resource even when available.

No: 47

- Need a priori knowledge of use of resources.

$$B_i = \max_{k=1}^{K} usage(k,i)C(k) \qquad B_i =$$

Basic Priority Ceiling

© NICTA 2007/2008

Priority Inheritance

- When lower priority job blocks, it inherits priority of • blocked job.
- GOOD
 - No unbounded priority inversion
 - Simple
 - No prior knowledge required
 - Works with fixed and dynamic priorities.
- BAD
 - Possible Deadlock.
 - Blocking of jobs not in resource contention.

No: 45

- Blocking time could be better
- Indirection a pain in the neck

© NICTA 2007/20



Immediate Priority Ceiling Protocol

- Lower priority task inherits priority of potentially blocked task. Task may be denied resource even when available.
 - GOOD
 - Simple.Shared run-time stack.
 - Shared full-time stack.
 Reduced Context-Switching
 - No deadlock
 - No unbounded priority inversion.
- BAD
 - Task may be denied resource even when available
 - Task may be affected by blocking effect without using any resources

No: 49

- Need a priori knowledge of use of resources.
- No self suspension while holding a resource

© NICTA 2007/2008

| | | NICTA |
|-------------------------|--------------------|--------------------------------------|
| | | THE UNIVERSITY OF NEW SOUTH WALES |
| Schedulability Analysis | of Real-Time Syste | ems |
| | | |
| o NICTA 0007/0008 | | |
| UNCTA2007/2008 | No: 52 | |

Implementation Comparison

- Non-preemptable critical sections
 - Easy to implement. Either blocking interrupts or syscall to have that implemented on behalf of task
- Priority Inheritance
 - Fairly straightforward, however requires various references (e.g. which thread is holding a resource)
- Basic Priority Ceiling
 - Requires application designer to explicitly identify which resources will be requested later (when first resource request of nested requests is made) on top of references
- · Immediate priority ceiling
 - Very easy to implement: Only requires ceilings associated with each resource mutex (that's something which may be automated if all tasks known
- Alternatively server task encapsulating the critical section
 NICTA 2007/2008
 No: 50

- - In the classical sense this is, whether all the deadlines are met under all circumstances;
 - Recent move to satisfaction of Quality-of-Service constraints;
 - Relies on availability of computation time of tasks – WCET;

No: 53

- Execution time profiles.

© NICTA 2007/2008

© NICTA 200

Reflective/Feedback-based Scheduling Critical Inst

0

NICTA

0

NICTA

Adaptive systems

© NICTA 2007/2008

- · By definition soft real time
- Adjusts scheduling based on information about change

No: 51

- Capable of better coping with "the unknown"
- Connects quite well with adaptive applications

- Critical Instant
- Trivial for independent tasks
 - All events happen at the same time;
 - However, implicitly consider all possible phases (take nothing for granted).
- However, get's more tricky (but tighter) having dependencies
 - What phasing of other activities produces the biggest load.
 - An activity is a string of tasks triggered by a single event.

Response Time Analysis

· Does not directly consider deadlines

0.

NICTA

- Makes the assumption of jobs being executed in order
- · Usually used in fixed priority systems







| Rate Monotonic Analysis | NICTA |
|---|-------|
| Looks at <i>utilisation</i> do determine whether a task is schedulable Initial work had following requirements: All tasks with deadlines are periodic All tasks are independent of each other (there exists no precedence relation, nor mutual exclusion) T_i= D_i C_i is known and constant Time required for context switching is known | |

Formal RTA

 Assumptions j<i ⇒priority j is lower than priority i

No: 57

· Critical instant

© NICTA 2007/2008

• Iterative process $w_i^0 = C_i$

$$w_i^{n+1} = C_i + \sum_{\forall j < i} \left\lceil \frac{w_i^n}{T_j} \right\rceil * C_j$$

Rate Monotonic Analysis contd

No: 59

Bound is given by:

© NICTA 2007/2008

© NICTA 2007/2008

0.

NICTA

$$\mu = \sum_{\forall i} \left\lceil \frac{C_i}{T_i} \right\rceil \le n * (2^{\frac{1}{n}} - 1)$$

- Has been relaxed in various ways, but still it is only an approximate technique.
- Further info can be found here: http://www.heidmann.com/paul/rma/PAPER.htm











| Is it a Problem? | 18 | |
|------------------|----|-------|
| | | NICIA |

- Safety critical computer systems exist and are deployed
- Yes, but ...
 - Safety critical systems have been
 - · highly restrictive in terms of HW/SW used
 - · highly restrictive in terms of complexity
 - used a lot of manual inspection and pen and paper work

| © NICTA 2007/2008 | No: 63 | © NICTA 2007/2008 | No: 66 |
|-------------------|--------|-------------------|--------|
| | | | |

| ls it a Prob | lem? contd | | Structural / | Analysis contd | |
|--|--|------------------------------|--|--|------------------------|
| – The stuff ir | n the last slide doesn't scale | ! | Object Cod | e | NICIA |
| industry no | ot in the safety critical arena | have | – Pros: | | |
| been using | measurements with safety | factors. | – All compile | er optimisations are done | |
| Worked fi work goo Excessive | ne with simple architectures, but d with more advanced computers e overestimation and underestim | t doesn't S ation with | – This is wh macros, p | at really is running, no trouble reprocessors | e with |
| same fac doesn't h | tor for different programs, paramelep too much | eterisation | – Cons: – Needs to s | second guess what all the var | riables |
| Large body | of academic work, but lit | ttle | meant | | |
| industrial up | | | | e original information is lost | alle |
| | | | object m | ember functions | ano, |
| | | | | | |
| Generic Pro | blem Partitioning | NICTA | Structural / | Analysis contd | NICTA |
| Some analysis but the general Program Str. Ana Constr. Const. Constr. Constr. Constr. Constr. Const | ctural alysis aint/Flow alysis Computation | s of these, | Assembly (– Pros: – All compile – Cons: – Same as (– Potentially | Code er optimisations done Object Code + v still some macros | |
| © NICTA 2007/2008 | No: 68 | | © NICTA 2007/2008 | No: 71 | |
| Structural / • Can work o – Source co – Object coo – Assembly | Analysis In: de and/or de or Code | NICTA | • Source Coo – Pros: – All informa – Structure i continue o indirect ca – Cons: – Trouble w | Analysis contd de ation the user has put there is in pure form (e.g. multiple loo conditions, object member fun ills) | there p ictions, |
| | | | - Needs to s | second guess what the comp | iler <i>will</i> |
| © NICTA 2007/2008 | No: 80 | | O | No: 72 | |
| | 140.00 | | | 1907.7.2 | |





- This information can be gained:
 - By static analysis (most importantly abstract interpretation)

No: 74

- By observation (worst case?)
- By user annotations

© NICTA 2007/2008

Flow information

- hardware features of processors
- Interaction of software and hardware

Flow Info Characteristics 0 NICTA Structurally possible 10 flows (infinite) // A Basic finite if(...) // B (D,G // C Statically al // D stual for T if(...) // E // F else // G } while(...) // H else WCET fou } while(...) // J Relation between possible Basic block graph Example program executions and flow info © NICTA 200 No: 75

0 Static Analysis NICTA · Looking at basic blocks in isolation (tree based, **IPET** based) - Problem of caching effects · Path based analysis: popular but very expensive

No: 77

- · Problem of conservative assumptions
- · Hardware analysis is very expensive
 - Data caches and modern branch prediction are very hard to model right.
 - Call for simpler hardware, e.g. scratchpad memory instead of caches

© NICTA 2007/2008 No: 78

Measurement Based Analysis

- End-to-end measurements + safety factor used for industrial soft-real time system development
 - Failed for modern processors as WC could hardly be expressed as function of AC
- · Measurement on basic block level
 - Safer than end-to-end measurements but potentially very pessimistic

No: 79

- · What is the worst-case on HW?
- · Can it be reliably produced?

© NICTA 2007/2008

© NICTA 2007/2008

· What about preemption delay?

 Tree Representation

 Image: constraint of the second seco

Path Based Computation

0

NICTA

- · Follows each individual paths
- · Becomes quickly intractable for large applications
- Altenbernd and Co have tried a simplified approach:
 - Starting out from the entry point of a function follow a path in the CFG and annotate each node with the execution time up to this node
 - Do so with any other path, but whenever a join node is reached compare new value up to this point with annotated value







- Throughput
 - Managing continuous load
- · Fault tolerance

external events

- Managing bugs, HW faults
- · Reliability

© NICTA 2007/2008

- Ensuring uptime, HW/SW upgrades ...

No: 86

| | Configurable number of hardware interrupt lines | | Private Fast Interrupts (FIQ) | (Can be use as NMI) |
|-----------------------------------|--|--|---|-----------------------------------|
| | Inte | rrupt Distributor | | |
| Per-CPU aliased peripherals | | | • Trans CPU. Widog Interface | |
| | , , | 8 8 | | |
| Configurable between | CPU/VFP CPU/VF | P CPU/VFP | CPU/VFP | |
| 1 and 4 Symmetric | L1 Memory L1 Mem | ory L1 Memory | L1 Memory | |
| | | ; <u> </u> | <u>†‡_</u> | |
| Priv Periph | ate sral | | I & D Coherence 64bit bus Control bus | |
| | | ↑ ↑ | | |
| | Pr AX 6 | mary Optional 2*4 I R/W AXI R/W 4-bit 64-bit bus | Vecto (VF | r Floating Poir P) is optional |

Image taken from http://linuxdevices.com/files/misc/arm-mpcore-architecture-big.jpg
© NICTA 2007/2008 No: 69



0.

NICTA

CPU

Memory

NIC

CPU

Caches Memory

NIC

| Issues | Non-Preemptive |
|---|--|
| | • But!!! |
| Resource contention | Less efficient processor use |
| Execution units | Anomalies: response time can increase with |
| – Caches | Changing the priority list |
| – Memory | Increasing number of CPUs |
| – Network | Reducing execution times |
| Adding a CPU does not help | Weakening the precedence constraints |
| - Example double the load 2 instead of 1 CPU | - Bill packing problem for hard |
| | but practically useless, as preemption and task migration overhead outweigh gains of Multiprocessors. |
| NICTA 2007/2008 No: 91 | o NICTA 2007/2008 No: 94 |
| Solutions?? | And now? |
| Partitioning | No global solution. |
| - Resource contention still there! | Partitioning and reducing it to single CPU problem |
| Assignment using heuristics | good, but still contention of resources. |
| mostly theoretical so far | Next step: After figuring out how to do the |
| Assumptions: | scheduling, what about preemption delay? |
| Zero preemption cost | Industry works with SMP/SMT, but most often on a yory ad bac basis |
| Zero migration cost | a very au noc basis. |
| Infinite time slicing | Active and unsolved research area |
| Don't translate into reality Acceptance test and no task migration a way to make it work | • Why does it work on non-RT? |
| | Running the "wrong" task is not critical. |
| | |
| Solutions?? | |
| NICTA | NICTA |
| Quite often non-preemptive | Integrating Deal Time and |
| Fewer context switches | General-Purpose |
| Reasoning is easy | Computing |
| IEEE Computer reference to insanity | |
| – Testing is easier?? | Many thanks to: Scott A. Brandt |
| Reduce need for blocking | University of California. Santa Cruz |

No: 96

No: 93

• But!

© NICTA 2007/2008

Real-Time vs. General-Purpose OS

- Real-time and general-purpose operating systems implement many of the same basic operations
 - Process mgmt., memory mgmt, I/O mgmt, etc.
- · They aim for fundamentally different goals
 - Real-time: Guaranteed performance, timeliness, reliability
 - General-purpose: Responsiveness, fairness, flexibility, graceful degradation, rich feature set
- They have largely evolved separately
 - Real-time system design lags general-purpose system design by decades

No: 97

They need to merge

© NICTA 2007/2008



We want to run processes with different timeliness requirements in the same system

NICTA

- HRT, RB, SRT, and BE Existing schedulers largely provide point solutions:
- HRT or RB or one flavor of SRT or BE
- Hierarchical scheduling is a partial solution
- Allows apps with a variety of timeliness requirements, BUT - Static, inflexible hierarchies
- Goal: Uniform, fully dynamic integrated real-time scheduling Same scheduler for all types of

© NICTA 2007/2008

NICTA

NICTA

applications No: 100

Why?

- We want both flexible general-purpose processing and robust realtime processing
 - Multimedia is ubiquitous in general-purpose systems
- Real-time systems are growing in size and complexity · Such systems are possible
- Look at the popularity of RTLinux
 - GP hardware has grown powerful enough to support traditional hard real-time tasks (multimedia, soft modems, etc.) - Windows, MacOS, etc., are already headed in this direction
- Existing solutions are ad hoc
- RTLinux, MacOS, Windows?
- The world is already headed that way

- CPU, storage, memory, network, ... · They must be hard real-time at their core

- Microsoft, HP, Intel, Dell all want to develop integrated home systems
- Complex distributed real-time systems do more than hard real-time

· We need integrated solutions for each type of resource

- This is the only way to guarantee the hardest constraints

- SRT and BE support cannot be added as an afterthought

We need an overall model for managing the separate

- Defaults should be reasonable, and helpful

- Each process must be able to specify it's per-resource constraints

No: 99

They must provide native hard real-time, soft real-time,

· We need to get out in front and lead the way

How?

0 Separate Resource Allocation and Dispatching NICTA Observation: Scheduling consists of two distinct auestions: Deadline Resource allocation SRT Allocation How much resources to allocate to each process Resourc Dispatching llocati Resource Soft When to give each process the resources it has been allocated SRT Real Tim Existing schedulers integrate their management

No: 101

Real-time schedulers implicitly separate them somewhat via job admission

© NICTA 2007/2008

Dispatching

nstrained

0

NICTA

The (RAD) Scheduling Model

0. NICTA

constrained

- Separate management of Resource Allocation and Dispatching
 - and separate policy and mechanism



and best-effort support

- Neither can HRT

resources









Increasing Rate (= increasing WCET) Theorem 2: The resource usage of any task can be increased at any time, within the available resources Civer a feasible EDE ashedula at any time task T may

- Given a feasible EDF schedule, at any time task T_i may increase utilization by any amount up to 1-U without causing any task to miss deadlines in the resulting EDF schedule



Increasing Rate (= increasing WCET)

Theorem 2: The resource usage of any task can be increased at any time, within the available resources

- Given a feasible EDF schedule, at any time task T_i may increase utilization by any amount up to 1-U without causing any task to miss deadlines in the resulting EDF schedule



© NICTA 2007/2008

© NICTA 2007/2

RBED EDF Mode Change Theory

- Theorem 1: EDF is optimal under this task model
- Corollary: A new task may enter at any time, within available resources
- Theorem 2: The rate of any task can be increased at any time, within available resources
- Theorem 3: The period of any task can be increased at any time
- Theorem 4: The rate of any task can be lowered at any time, down to what it has already used in the current period
- Theorem 5: The period of any task can be reduced at any time, down to the time corresponding to the current period's resource usage
- Corollary: The period of any task can be increased at any time (without changing WCET)
- Corollary: The period of a job which is ahead of its target allocation can be reduced at any time, down to the time corresponding to its current resource usage (without changing WCET) as long as the resources are available for the rate change

No: 112

© NICTA 2007/2008



NICTA

0

ΝΙCTΔ

Increasing Rate (= increasing WCET)

 Theorem 2: The resource usage of any task can be increased at any time, within the available resources

 Given a feasible EDF schedule, at any time task *Ti* may increase utilization by any amount up to *1–U* without causing any task to miss deadlines in the resulting EDF schedule



Better Slack Management: BACKSLASH

NICTA

NICTA

- Existing algorithms tend to ignore the needs of "background" tasks
 - Slack provided when everything else is idle
 - Aim for "fair" allocation and 100% utilization
- Slack reclamation is critical in an integrated real-time system
 - Utilization is important for best-effort systems
 - Soft real-time and best effort performance depends on the effective use of slack
- BACKSLASH improves performance via slack scheduling
 - Focuses on when slack is allocated, and to which process

© NICTA 2007/2008 No: 114











| | BACKSLASH Conclusions | O • NICTA |
|---|--|---------------------|
| • | In an integrated system supporting HRT, SRT and BE, the performance of SRT (and BE) depends on the effective reclamation and distribution of slack | |
| • | Four principles for effective slack reclamation and distribution: | |

- 1. Distribute slack as early as possible
- 2. Give slack to the ready task with the highest priority
- 3. Allow tasks to borrow against future reservations
- 4. Retroactively give slack to tasks that needed it
- 5. SMASH: Conserve slack across idle times!

© NICTA 200

 Our results show that these principles are effective: BACKSLASH significantly outperforms the other algorithms and improves SRT (and/or BE) performance

No: 120















0 Effects of Switching Times NICTA **Ideal World** f₁ f_2 **Real World** t₁ f₂ © NICTA 200

No: 128



- Allocation
 - If resource locked, block
 - If (potentially modified) priority of task higher than the ceiling of any resource not used by that task but used at the time, allocate - Else, block
- Priorities:

.

- If task T₁ blocked at resource held by task T_{2:}
 - + $\boldsymbol{\tau}_2$ is lifted in priority to task $\boldsymbol{\tau}_1$

· revert to original priority once all resources are released O NICTA 2003



Basic Priority Ceiling and Deadlock



0

NICTA

- At any time the priority of task τ_i > ceiling priority of resource currently in use THEN
 - 1. task τ_{l} will not require any resource currently in use
 - 2. Any task τ_k with priority greater than task τ_l will not require any resource currently in use

i.e.:

 No task currently holding a resource can inherit a higher priority and preempt task τ_{I} w