

Integrated Schedulers for a Predictable Interrupt Management on Real-Time Kernels

S. Sáez A. Crespo

Instituto de Automática e Informática Industrial
Universidad Politécnica de Valencia

19th International Conference on Reliable Software Technologies
June 2014, Paris

Index

- 1 Introduction
- 2 Integrated Interrupt Model
- 3 Fully Integrated Interrupt Model
- 4 Scheduling Cartesian Tree
- 5 Conclusions

Index

- 1 Introduction
- 2 Integrated Interrupt Model
- 3 Fully Integrated Interrupt Model
- 4 Scheduling Cartesian Tree
- 5 Conclusions

Event-driven real-time schedulers

- ▶ Activate tasks when certain events occur:
 - ▶ external interrupts, timer interrupts, software mechanisms, ...
- ▶ Depending on the activation event, tasks can be classified in:

Hardware Activated Tasks

Tasks are woken up by *Interrupt Service Routines*.

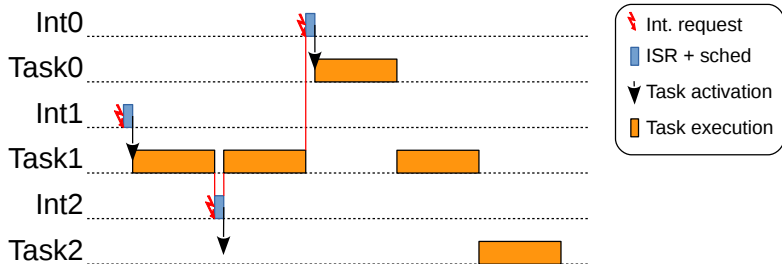
Software Activated Tasks

Tasks are woken up by software mechanisms: timing events, delays, semaphores, barriers, ...

- ▶ Tasks have priorities that establish the execution order.
- ▶ These priorities can be used to map task's criticality level.
- ▶ It is desirable that a high priority task did not suffer unnecessary interference from lower priority tasks.

Classical Interrupt Model

- ▶ Since interrupts have higher priorities than any task, their ISRs are always executed despite of:
 - ▶ the priority of the task currently in execution
 - ▶ the task that is going to be activated by a given interrupt.



- ▶ High priority tasks suffer interference from interrupts used to activate lower priority tasks that are obviously not executed at activation instant.

Classical Interrupt Model (cont.)

- ▶ Blocking Time at priority i is:

$$B(i) = \underbrace{|L(i)|}_{\text{Nr of lower priority tasks}} \times \underbrace{(\delta^{\text{isr}} + \delta^{\text{Sched-A}})}_{\text{Task activation overhead}}$$

- ▶ δ^{isr} → mostly ISR execution time

- ▶ $\delta^{\text{Sched-A}}$ → scheduler has to ...

Q_{insert}^R → update the *ready queue*,

$Q_{\text{find-min}}^R$ → determine the highest priority task

... each time a task is activated.

Question:

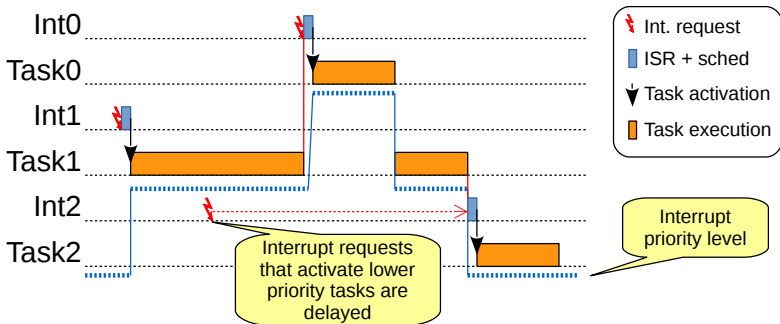
If the activated task is not going to be immediately executed, why does the ISR have to interfere with high priority tasks?

Index

- 1 Introduction
- 2 Integrated Interrupt Model**
- 3 Fully Integrated Interrupt Model
- 4 Scheduling Cartesian Tree
- 5 Conclusions

Integrated Interrupt Model

- ▶ A unique priority space is used for tasks and ISRs.
- ▶ Each time a task is activated, the interrupt priority level is changed to avoid lower priority interrupts to be attended.

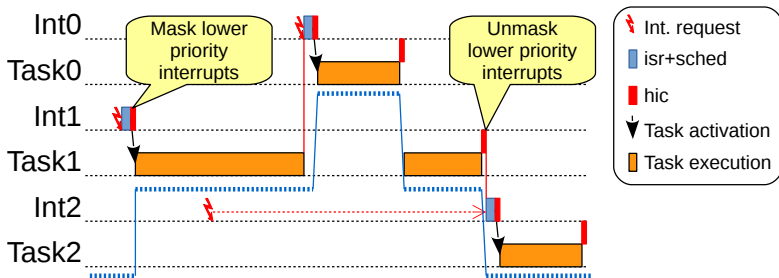


- ▶ The *Hardware Interrupt Controller* receives lower priority interrupt requests but the CPU is not notified.

Integrated Interrupts Model: Hardware Activated Tasks

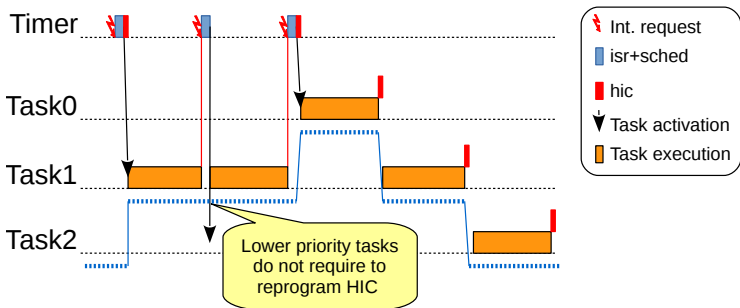
- ▶ Requires a Hardware Interrupt Controller with multiple interrupt priorities.
- ▶ If this HW support is not present, an additional overhead is introduced:

δ^{hic} time to mask unnecessary interrupts according to the new priority level.



Integrated Interrupts Model: Software Activated Tasks

- ▶ To fully support SAT multiple HW timers with different priorities are required.
 - ⇒ This HW is not commonly available.
- ▶ SATs have to share the Timer interrupt
 - ⇒ Timer interrupt is always enabled.



Integrated Interrupts Model: Software Activated Tasks

- ▶ Only SATs activations can produce interference, but the interference is higher.
- \exists Blocking Time at priority i due to unnecessary activations:

$$B^{\text{SAT}}(i) = |L_{\text{SAT}}(i)| \times (\delta^{\text{isr}} + \delta^{\text{Sched-A}} + \delta^{\text{timer}})$$

$$\delta^{\text{Sched-A}} = \underbrace{Q_{\text{delete-min}}^W}_{\text{Remove } \tau \text{ from timer queue}} + \underbrace{Q_{\text{find-min}}^W}_{\text{Find out next activation time}} + \underbrace{Q_{\text{insert}}^R}_{\text{Insert } \tau \text{ into ready queue}} + \underbrace{Q_{\text{find-min}}^R}_{\text{Schedule}}$$

δ^{timer} → Overhead of reprogramming the HW timer.

Index

- 1 Introduction
- 2 Integrated Interrupt Model
- 3 Fully Integrated Interrupt Model**
- 4 Scheduling Cartesian Tree
- 5 Conclusions

A new proposal: Virtual Integrated Interrupt Model

GOAL: To program timer interrupts corresponding **only** to higher priority tasks.

- ▶ Each time a task starts its execution, it has to:
 1. Find the closer waiting task with a priority higher than the current one.
⇒ **The next preemptor.**
 2. Program the timer interrupt for activating only the next preemptor.
- ▶ When the next preemptor wakes up, previously ignored lower priority tasks are also awakened.

Drawbacks

- This approach gives rise to additional scheduling overheads.
- Commonly used data structures for ready and waiting queues are not adequate for these new scheduling operations.

Scheduling overheads using ready/waiting queues

- ▶ Waiting queue is sorted by release time
→ to find the next preemptor could have a $O(N^W)$ cost.
- ▶ When the next preemptor is activated
→ N_p lower priority tasks have to be moved from the waiting queue into the ready queue.
→ in the worst case $N_p = N^W$.

$$\begin{aligned}
 \delta_I^{\text{Sched-A}} &= N_p \times \overbrace{(Q_{\text{delete-min}}^W + Q_{\text{insert}}^R)}^{\text{wake up a lower priority task}} \\
 &+ \underbrace{Q_{\text{find-min}}^R}_{\text{select the next task}} + \underbrace{Q_{\text{find-preemptor}}^W}_{\text{determine its next preemptor}}
 \end{aligned}$$

Index

- 1 Introduction
- 2 Integrated Interrupt Model
- 3 Fully Integrated Interrupt Model
- 4 Scheduling Cartesian Tree**
- 5 Conclusions

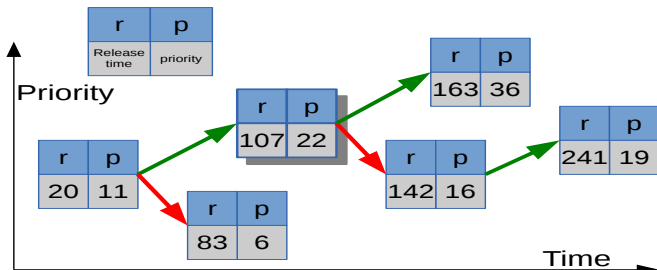
Scheduling Cartesian Tree

Goals

- ▶ To avoid massive task movements from waiting queue.
- ▶ To efficiently determine the next preemptor.

Cartesian Tree

- ▶ A binary tree sorted by two keys:
priority (top-bottom subtrees) and release time (node depth).



Scheduling Cartesian Tree (cont.)

- ▶ Only one tree that represents the ready and waiting queue.
- ▶ SC-Tree is sorted by absolute release time:
 ⇒ a task activation does not modify the structure.

$$\delta_l^{\text{Sched-A}} = N_p \times (\cancel{Q_{\text{delete-min}}^W + Q_{\text{insert}}^R}) + Q_{\text{find-min}}^R + Q_{\text{find-preemptor}}^W$$

- ▶ Next preemptor is the *top* child node:
 ⇒ activation time overhead is constant. **No release jitter!!**

$$\delta_l^{\text{Sched-A}} = C_{\text{find-min}} + C_{\text{find-preemptor}}$$

- ▶ Main scheduling overhead occurs during task suspension.
 - ▶ It can be accounted as part of the WCET.
 - ▶ A careful implementation could allow preemptive SC-Tree operations
 → It produces no/low blocking times.

Priority Inheritance

- ▶ When a task locks/unlocks a shared resource and an Inheritance Protocol is used, priority changes are produced.
- ▶ These temporal priority changes have an additional cost in an Integrated Interrupt Model:

$$Q^{PO} = \delta^{hic} + Q_{insert}^R + Q_{delete-min}^R$$
$$+ 2 \times Q_{find-min}^R \xrightarrow{C_f} + 2 \times Q_{find-preemptor}^W \xrightarrow{C_p}$$

- ▶ When the shared resource is freed, the cost of activate pending medium priority tasks is lower than if they had been activated in their release instants:
 - **no** ISR has been executed
 - **no** interrupt priority level has been changed.

Index

- 1 Introduction
- 2 Integrated Interrupt Model
- 3 Fully Integrated Interrupt Model
- 4 Scheduling Cartesian Tree
- 5 Conclusions**

Summarising

- ▶ A new approach has been presented to completely avoid activation interference from lower priority tasks.
- ▶ Suitability of dual-queue schedulers to implement this approach has been evaluated.
- ▶ A new data structure has been proposed and its overhead compared against the classical model.
- ▶ The paper provides the necessary tools to check if this approach is suitable for a given system taking into account the real system overheads.

Pending issues

- ▶ SC-Tree behaviour during priority inheritance can be improved.
- ▶ To study the applicability of the Integrated Interrupt Model to dynamic priorities.

Integrated Schedulers for a Predictable Interrupt Management on Real-Time Kernels

S. Sáez A. Crespo

Instituto de Automática e Informática Industrial
Universidad Politécnica de Valencia

19th International Conference on Reliable Software Technologies
June 2014, Paris

Thank you! Any question?