

# Sleepy Sloth: Threads as Interrupts as Threads

**Wanja Hofer**, Daniel Lohmann, Wolfgang Schröder-Preikschat



32nd IEEE Real-Time Systems Symposium  
November 30, 2011



# My Personal Dilemma

## Motivation

Get a PhD in computer science (in operating systems group)

## Problem

Requires lots of work (build own operating system),



but I am a lazy sloth

## Solution

Let the hardware do the work!



# Control Flows in Embedded Systems

|                 | Activation Event | Sched./Disp. | Semantics       |
|-----------------|------------------|--------------|-----------------|
| ISRs            | HW               | by HW        | RTC             |
| Threads         | SW               | by OS        | Blocking        |
| Sloth [RTSS 09] | HW or SW         | by HW        | RTC             |
| Sleepy Sloth    | HW or SW         | by HW        | RTC or Blocking |

(RTC: Run-to-Completion)



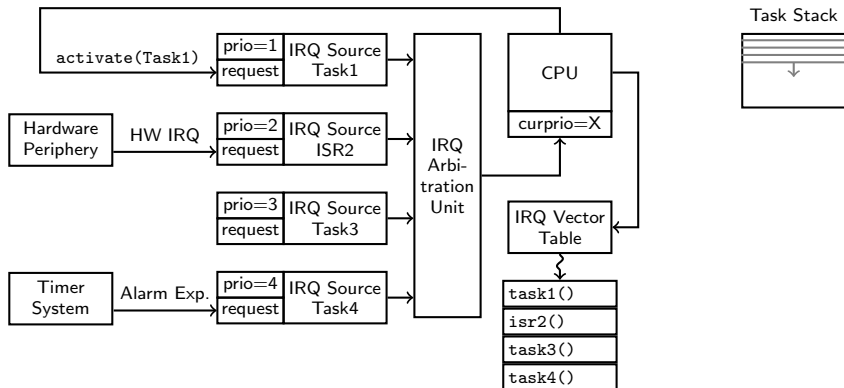
# Talk Outline

---

1. Sloth Revisited
2. Sleepy Sloth Design and Implementation
3. Evaluation
4. Conclusions and Future Work



# Sloth Revisited



- Platform must support IR priorities and software IR triggering



# Sleepy Sloth: Main Goal and Challenge

---

## Main Goal

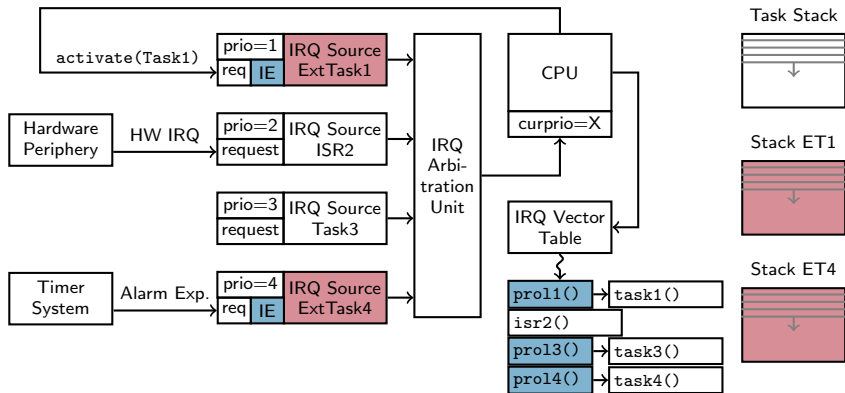
Support **extended blocking tasks** (with stacks of their own) while preserving Sloth's **latency benefits** by having threads run as ISRs

## Main Challenge

IRQ controllers do not support **suspension and re-activation** of ISRs



# Sleepy Sloth Design: Task Prologues and Stacks



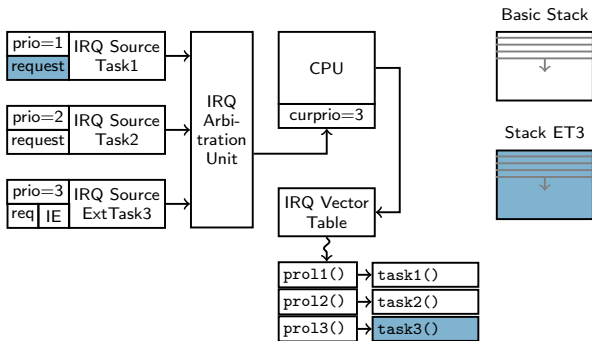
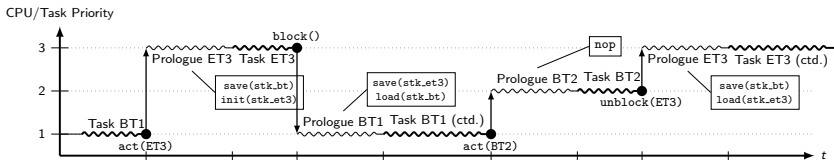
# Sleepy Sloth: Dispatching and Rescheduling

- Task prologue: switch stacks if necessary
  - Switch *basic task* → *basic task* omits stack switch
  - On job start: initialize stack
  - On job resume: restore stack
- Task termination: task with next-highest priority needs to run
  - Yield CPU by setting priority to zero
  - (Prologue of *next* task performs the stack switch)
- Task blocking: take task out of “ready list”
  - Disable task’s IRQ source
  - Yield CPU by setting priority to zero
- Task unblocking: put task back into “ready list”
  - Re-enable task’s IRQ source
  - Re-trigger task’s IRQ source by setting its pending bit





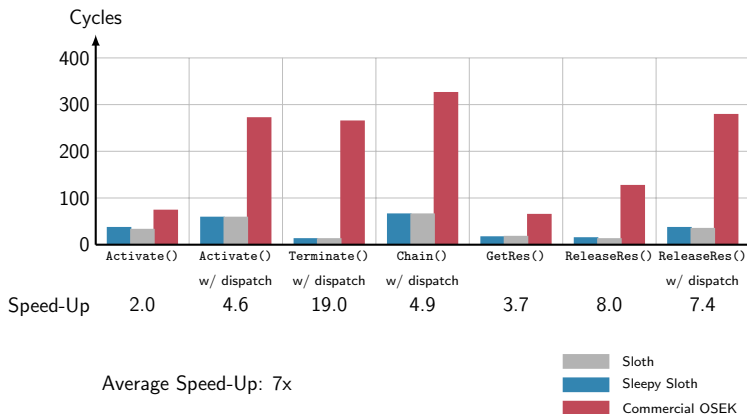
# Sleepy Sloth: Example Control Flow



- Reference implementation on Infineon TriCore microcontroller
- 32-bit load/store architecture
- Interrupt controller: 256 priority levels, about 200 IRQ sources with memory-mapped registers
- Measurements: system call latencies in 3 system configurations, compared to a leading commercial OSEK implementation
  1. Only basic run-to-completion tasks
  2. Only extended blocking tasks
  3. Both basic and extended tasks



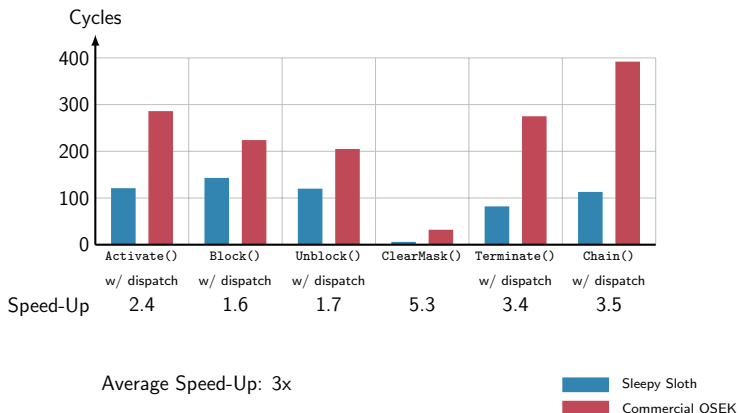
# Evaluation: Only Basic Tasks



- Sleepy Sloth outperforms commercial kernel with SW scheduler
- Sleepy Sloth as fast as original Sloth



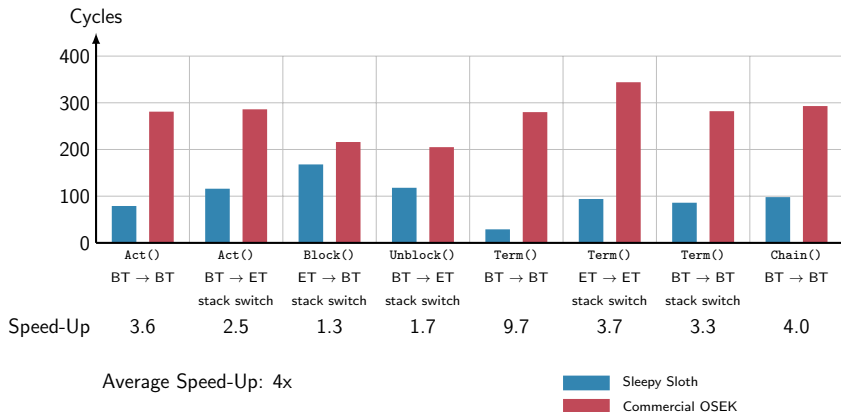
## Evaluation: Only Extended Tasks



- Still faster than commercial kernel with SW scheduler
- Sleepy Sloth: Extended switches slower than basic switches



# Evaluation: Extended *and* Basic Tasks



- Basic switches in a mixed system only slightly slower than in purely basic system



# Evaluation: Summary

## Main Goal

Support **extended blocking tasks** (with stacks of their own) while preserving Sloth's **latency benefits** by having threads run as ISRs



- Sleepy Sloth threads, which are implemented as ISRs, ...
  - ... can be run-to-completion or blocking  
→ flexible
  - ... can be activated by hardware or through a syscall  
→ flexible
  - ... are scheduled and dispatched by the IRQ controller  
→ fast (speed-up  $\approx 2-5$ )
  - ... run in a single priority space, the IRQ priority space  
→ no priority inversion
- Future work: investigate suitability of the Sloth concept for ...
  - ... integration with RTLinux/RTAI on x86-APIC
  - ... multi-core IRQ controllers: Pandaboard (dual-Cortex-A9), x86-APIC
  - ... time-triggered scheduling elements: AUTOSAR schedule tables

