

Teachers: Robert Dick Peter Dinda
Office: L477 Tech 338, 1890 Maple Ave.
Email: dickrp@ece.northwestern.edu pdinda@cs.northwestern.edu
Phone: 467-2298 467-7859
Webpage: <http://ziyang.ece.northwestern.edu/EXTERNAL/realtime>

1

Goals for lecture

- Lab four
- Example scheduling algorithm design problem
 - Will initially focus on static scheduling
- Real-time operating systems
- Comparison of on-line and off-line scheduling code

3

Example problem: Static scheduling

- What is an FPGA?
- Why should real-time systems designers care about them?
- Multiprocessor static scheduling
- No preemption
- No overhead for subsequent execution of tasks of same type
- High cost to change task type
- Scheduling algorithm?

5

Rate monotonic scheduling

Main idea

- 1973, Liu and Layland derived optimal scheduling algorithm(s) for this problem
- Schedule the job with the smallest period (period = deadline) first
- Analyzed worst-case behavior on any task set of size n
- Found utilization bound: $U(n) = n \cdot (2^{1/n} - 1)$
- 0.828 at $n = 2$
- As $n \rightarrow \infty$, $U(n) \rightarrow \log 2 = 0.693$
- Result: For any problem instance, if a valid schedule is possible, the processor need never spend more than 71% of its time idle

7

| | |
|--------------------------------|----|
| 1 Reading assignment | 29 |
| 2 Lab six | 32 |

2

Lab four

- Talk with Promi SD101
- Sample sound at 3 kHz
- Multihop

4

Problem: Uniprocessor independent task scheduling

- Problem
 - Independent tasks
 - Each has a period = hard deadline
 - Zero-cost preemption
- How to solve?

6

Optimality and utilization for limited case

- Simply periodic: All task periods are integer multiples of all lesser task periods
- In this case, RMS/DMS optimal with utilization 1
- However, this case rare in practice
- Remains feasible, with decreased utilization bound, for in-phase tasks with arbitrary periods

8

Rate monotonic scheduling

- Constrained problem definition
- Over-allocation often results
- However, in practice utilization of 85%–90% common
 - Lose guarantee
- If phases known, can prove by generating instance

9

Proof sketch for RMS utilization bound

- Consider case in which no period exceeds twice the shortest period
- Find a pathological case
 - Utilization of 1 for some duration
 - Any decrease in period/deadline of longest-period task will cause deadline violations
 - Any increase in execution time will cause deadline violations

11

Proof sketch for RMS utilization bound

- See if there is a way to increase utilization while meeting all deadlines
- Increase execution time of high-priority task
 - $e'_i = p_{i+1} - p_i + \epsilon = e_i + \epsilon$
- Must compensate by decreasing another execution time
- This always results in decreased utilization
 - $e'_k = e_k - \epsilon$
 - $U' - U = \frac{e'_i}{p_i} + \frac{e'_k}{p_k} - \frac{e_i}{p_i} - \frac{e_k}{p_k} = \frac{\epsilon}{p_i} - \frac{\epsilon}{p_k}$
 - Note that $p_i < p_k \rightarrow U' > U$

13

Proof sketch for RMS utilization bound

- Get utilization as a function of adjacent task ratios
- Substitute execution times into $\sum_{k=1}^n \frac{e_k}{p_k}$
- Find minimum
- Extend to cases in which $p_n > 2 \cdot p_k$

15

Critical instants

Main idea:

A job's critical instant a time at which all possible concurrent higher-priority jobs are also simultaneously released

Useful because it implies latest finish time

10

RMS worst-case utilization

- In-phase
- $\forall k \text{ s.t. } 1 \leq k \leq n-1 : e_k = p_{k+1} - p_k$
- $e_n = p_n - 2 \cdot \sum_{k=1}^{n-1} e_k$

12

Proof sketch for RMS utilization bound

- Same true if execution time of high-priority task reduced
- $e'_i = p_{i+1} - p_i - \epsilon$
- In this case, must increase other e or leave idle for $2 \cdot \epsilon$
- $e'_k = e_k + 2\epsilon$
- $U'' - U = \frac{2\epsilon}{p_k} - \frac{\epsilon}{p_i}$
- Again, $p_k < 2 \rightarrow U'' > U$
- Sum over execution time/period ratios

14

Notes on RMS

- Other abbreviations exist (RMA)
- DMS better than or equal RMA when deadline \neq period
- Why not use slack-based?
- What happens if resources are under-allocated and a deadline is missed?

16

Essential features of RTOSs

- Provides real-time scheduling algorithms or primitives
- Bounded execution time for OS services
 - Usually implies preemptive kernel
 - E.g., linux can spend milliseconds handling interrupts, especially disk access

17

Threads vs. processes

- Threads: Low context switch overhead
- Threads: Sometimes the only real option, depending on hardware
- Processes: Safer, when hardware provides support
- Processes: Can have better performance when IPC limited

19

TinyOS

- Most behavior event-driven
- High rate → Livelock
- Research schedulers exist

21

μ C/OS-II

- Similar to BD threads
- More flexible
- Bigger footprint

23

Threads

- Threads vs. processes: Shared vs. unshared resources
- OS impact: Windows vs. Linux
- Hardware impact: MMU

18

Software implementation of schedulers

- TinyOS
- Light-weight threading executive
- μ C/OS-II
- Linux
- Static list scheduler

20

BD threads

- Brian Dean: Microcontroller hacker
- Simple priority-based thread scheduling executive
- Tiny footprint (fine for AVR)
- Low overhead
- No MMU requirements

22

Old linux scheduler

- Single run queue
- $\mathcal{O}(n)$ scheduling operation
- Allows dynamic goodness function

24

$\mathcal{O}(1)$ scheduler in Linux 2.6

- Written by Ingo Molnar
- Splits run queue into two queues prioritized by goodness
- Requires static goodness function
 - No reliance on running process
- Compatible with preemptible kernel

25

Real-time operating systems

- Embedded vs. real-time
- Dynamic memory allocation
- Schedulers: General-purpose vs. real-time
- Timers and clocks: Relationship with HW

27

Reading assignment

- Read Chapter 12 in J. W. S. Liu, *Real-Time Systems*. Prentice-Hall, Englewood Cliffs, NJ, 2000
- Read K. Ghosh, B. Mukherjee, and K. Schwan, "A survey of real-time operating systems," tech. rep., College of Computing, Georgia Institute of Technology, Feb. 1994

29

Lab four

- Please email or hand in the write-up for lab assignment four
- Problems? See me.
 - Will need everything from lab four working for lab six

31

Real-time linux

- Run linux as process under real-time executive
- Complicated programming model
- RTAI (Real-Time Application Interface) attempts to simplify
 - Colleagues still have problems at > 18 kHz control period

26

Summary

- Static scheduling
- Example of utilization bound proof
- Introduction to real-time operating systems

28

Goals for lecture

- Lab four?
- Lab six
- Simulation of real-time operating systems
- Impact of modern architectural features

30

Lab six

- Develop priority-based cooperative scheduler for TinyOS that keeps track of the percentage of idle time.
- Develop a tree routing algorithm for the sensor network.
- Send noise, light, and temperature data to a PPC, via the network root.
- Have motes respond to *send audio samples* and *buzz* commands.
- Play back or display this data on PPCs to verify the that the system functions.

32

Outline

- Introduction
- Role of real-time OS in embedded system
- Related work and contributions
- Examples of energy optimization
- Simulation infrastructure
- Results
- Conclusions

33

Introduction

- Real Time Operating Systems important part of embedded systems
 - Abstraction of HW
 - Resource management
 - Meet real-time constraints
- Used in several low-power embedded systems
- Need for RTOS power analysis
 - Significant power consumption
 - Impacts application software power
 - Re-used across several applications


35

Related work and contributions

- **Instruction level power analysis**
V. Tiwari, S. Malik, A. Wolfe, and T.C. Lee, Int. Conf. VLSI Design, 1996
- **System-level power simulation**
Y. Li and J. Henkel, Design Automation Conf., 1998
- **MicroC/OS-II**: J.J. Labrosse, R & D Books, Lawrence, KS, 1998
- **Our work**
 - First step towards detailed power analysis of RTOS
 - Applications: low-power RTOS, energy-efficient software architecture, incorporate RTOS effects in system design

37

Single task network interface



Checksum computation and output

Procuring Ethernet controller has high energy cost

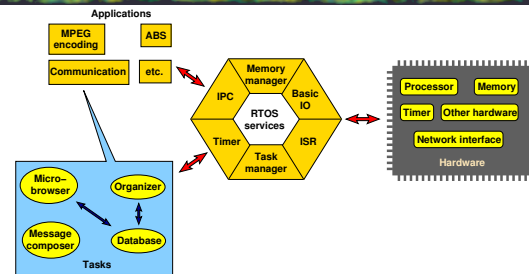
39

Introduction

- Real-Time Operating Systems are often used in embedded systems.
- They simplify use of hardware, ease management of multiple tasks, and adhere to real-time constraints.
- Power is important in many embedded systems with RTOSs.
- RTOSs can consume significant amount of power.
- They are re-used in many embedded systems.
- They impact power consumed by application software.
- RTOS power effects influence system-level design.

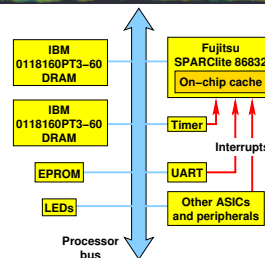
34

Role of RTOS in embedded system



36


Simulated embedded system



- Easy to add new devices
- Cycle-accurate model
- Fujitsu board support library used in model
- μ C/OS-II RTOS used

38

TCP example




Straight-forward implementation

Multi-task implementation

40


Multi-tasking network interface



RTOS power analysis used for process re-organization to reduce energy
 21% reduction in energy consumption. Similar power consumption.

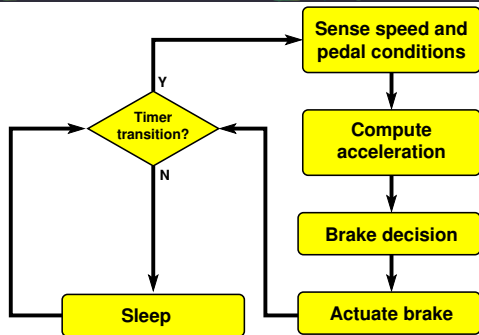
41

ABS example timing



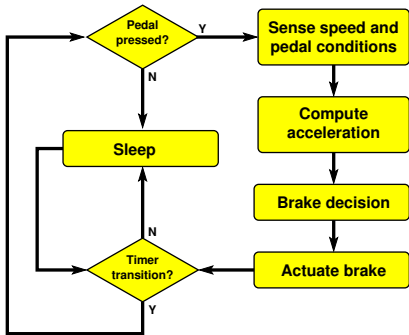
43

Periodically triggered ABS



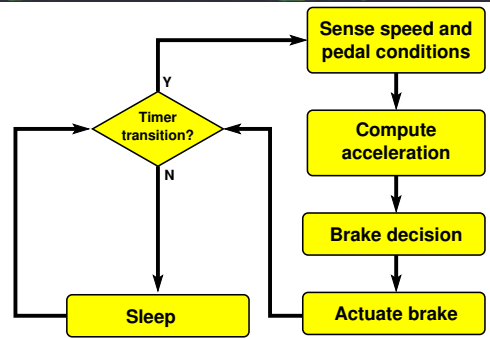
45

Selectively triggered ABS




47

ABS example




42

Straight-forward ABS implementation




44

Periodically triggered ABS timing



46


Selectively triggered ABS timing



63% reduction in energy and power consumption


48

Power-optimized ABS example



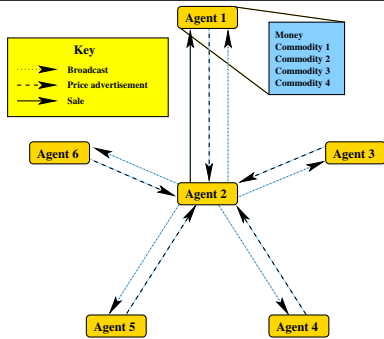
49

Experimental results



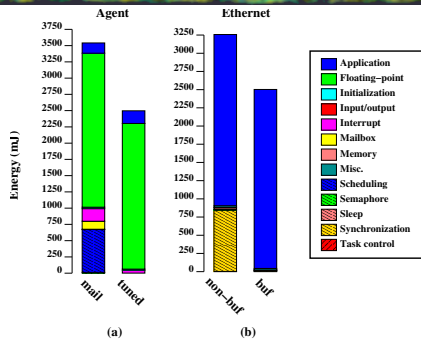
51

Agent example



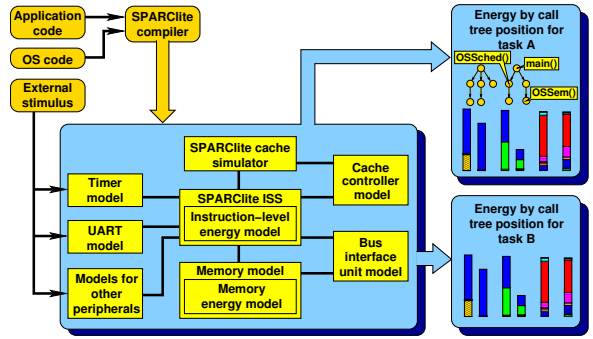
53

Experimental results




55

Infrastructure




50

Experimental results – time



52

Experimental results



54

Optimization effects

TCP example:

- 20.5% energy reduction
- 0.2% power reduction
- RTOS directly accounted for 1% of system energy

ABS example:

- 63% energy reduction
- 63% power reduction
- RTOS directly accounted for 50% of system energy

Mailbox example: RTOS directly accounted for 99% of system energy

Semaphore example: RTOS directly accounted for 98.7% of system energy

56

Partial semaphore hierarchical results

| | Function | Energy/invocation (μ J) | Energy (%) | Time (ms) | Calls |
|--------------------------------------|----------------------------|------------------------------|------------|-----------|-------|
| realstart 6.41 mJ total 2.02 % | init_dvcs | 0.41 | 0.00 | 0.00 | 1 |
| | init_timer | 1.31 | 0.00 | 0.00 | 1 |
| | 5.51 mJ total 1.74 % | | | | |
| | startup | 887.44 | 0.28 | 2.18 | 1 |
| | save_data | 1.56 | 0.00 | 0.00 | 1 |
| | init_data | 1.31 | 0.00 | 0.00 | 1 |
| | init_dss | 0.88 | 0.00 | 0.00 | 1 |
| | cache_on | 2.72 | 0.00 | 0.01 | 1 |
| | 155.18 mJ total 48.88 % | | | | |
| | Task1 | win_uartJrap | 1.90 | 1.20 | 9.73 |
| | _JOSDisableInt | 0.29 | 0.09 | 0.78 | 1000 |
| | _JOSDisableInt | 0.32 | 0.10 | 0.89 | 1000 |
| 31.18 mJ total 9.82 % | | | | | |
| spanscm_terminate | win_uartJrap | 0.75 | 0.00 | 0.00 | 1 |
| | _JOSDisableInt | 2.48 | 0.78 | 5.33 | 999 |
| | _JOSDisableInt | 0.29 | 0.18 | 1.59 | 1999 |
| | _JOSDisableInt | 0.29 | 0.18 | 1.59 | 1999 |
| | OSEventTaskWait | 3.75 | 1.18 | 9.22 | 999 |
| | OSSched | 19.07 | 6.00 | 47.97 | 999 |
| | OSSemPost | 0.29 | 0.09 | 0.78 | 1000 |
| | _JOSDisableInt | 0.29 | 0.09 | 0.78 | 1000 |
| | OSTimeGet | 0.27 | 0.08 | 0.70 | 1000 |
| | _JOSDisableInt | 0.29 | 0.09 | 0.78 | 1000 |
| | CPUInt | 1.09 | 0.00 | 0.00 | 1 |
| | exceptionHandler | 4.77 | 0.02 | 0.17 | 15 |
| | print | 2.05 | 0.65 | 5.06 | 1000 |
| | vprint | 108.89 | 34.30 | 290.53 | 1000 |

57

Conclusions

- RTOS can significantly impact power
- RTOS power analysis can improve application software design
- Applications
 - Low-power RTOS design
 - Energy-efficient software architecture
 - Consider RTOS effects during system design

59

Summary

- Labs
- Simulation of real-time operating systems
- Impact of modern architectural features

61

Energy per invocation for μ C/OS-II services

| Service | Minimum energy (μ J) | Maximum energy (μ J) |
|---------------------|---------------------------|---------------------------|
| OSEventTaskRdy | 18.02 | 20.03 |
| OSEventTaskWait | 7.98 | 9.05 |
| OSEventWaitListInit | 20.43 | 21.16 |
| OSInit | 1727.70 | 1823.26 |
| OSMboxCreate | 27.51 | 28.82 |
| OSMboxPend | 7.07 | 82.91 |
| OSMboxPost | 5.82 | 84.55 |
| OSMemCreate | 19.40 | 19.75 |
| OSMemGet | 6.64 | 8.22 |
| OSMemInit | 27.41 | 27.47 |
| OSMemPut | 6.38 | 7.91 |
| OSQInit | 20.10 | 20.93 |
| OS Sched | 6.96 | 52.34 |
| OSSemCreate | 27.87 | 29.04 |
| OSSemPend | 6.54 | 73.64 |
| etc. | etc. | etc. |

58

Impact of modern architectural features

- Memory hierarchy
- Bus protocols ISA vs. PCI
- Pipelining
- Superscalar execution
- SIMD
- VLIW

60