Energy-Efficient Static Priority and Speed Assignment for Real-Time Tasks with Non-Deterministic Release Times

Simon Perathoner¹, Jian-Jia Chen², Lothar Thiele¹

¹ ETH Zurich (Switzerland) ² KIT Karlsruhe (Germany)

RTCSA, August 24th, 2010, Macau SAR, China



Outline

- Motivation
- Problem Definition
- Analysis Framework
- Proposed Algorithms
- Experimental Evaluation
- Conclusion

Embedded Systems

Computer systems physically embedded into larger device



Requirements

- Performance
- Energy
- Fault tolerance
- Size
- Weight
- Cost

Methods for CPU Power Management

- Dynamic Voltage Scaling (DVS)
 - For CMOS circuits $E \propto V_{dd}^2$ (ignoring leakage)
 - Save energy by reducing supply voltage (clock frequency)

- Dynamic Power Management (DPM)
 - Switch between different power states
 - Helps to reduce power dissipation due to leakage



4

Real-Time System



Goal: Minimize energy consumption while guaranteeing deadlines



Related Work

- Offline Approaches
 - Take scheduling decisions statically (at design time) based on expected worst-case workload
 - E.g. [Yao et al. 1995], [Quan et al. 2007], [Yun et al. 2003]
- Online Approaches
 - Take scheduling decisions dynamically (at run time) by adapting to actual workload
 - E.g. [Yao et al. 1995], [Aydin et al. 2001], [Mochocki et al. 2005]

Limitations of existing approaches:

- 1) Very restrictive assumptions on task release times
 - Examples: Periodic / sporadic event streams
 - Full a priori knowledge of release times

But in real systems task release times are less deterministic !

- Tasks triggered by physical events
- Tasks triggered by other tasks in complex distributed architectures
- 2) Task priorities are not subject to optimization





Outline

- Motivation
- **Problem Definition**
- Analysis Framework
- Proposed Algorithms
- Experimental Evaluation
- Conclusion

Problem Definition



9



Example

$$\Gamma = \{\tau_{1}, \tau_{2}, \tau_{3}\}$$
Task release: Periodic with non-deterministic
but bounded jitter
$$p_{1} = 10 \quad j_{1} = 3 \quad C_{1} = 1 \quad D_{1} = 1 \quad [ms]$$

$$p_{2} = 5 \quad j_{2} = 3 \quad C_{2} = 1 \quad D_{2} = 9$$

$$p_{3} = 8 \quad j_{3} = 1 \quad C_{3} = 1 \quad D_{3} = 10$$

$$P(s) = \hbar(0.08 + 1.52s^{3}) Watt$$
(a)
$$\Pi(\tau_{1}) = 1 \qquad \Pi(\tau_{2}) = 2 \qquad \Pi(\tau_{3}) = 3$$

$$\Sigma(\tau_{1}) = 1 \qquad \Sigma(\tau_{2}) = 0.6 \qquad \Sigma(\tau_{3}) = 0.\overline{3}$$

$$E_{a}(10s) = 3475mJ$$



Example

$$\Gamma = \{\tau_{1}, \tau_{2}, \tau_{3}\}$$
Task release: Periodic with non-deterministic
but bounded jitter
$$\begin{array}{c} p_{1} = 10 \quad j_{1} = 3 \quad C_{1} = 1 \quad D_{1} = 1 \quad [ms] \\ p_{2} = 5 \quad j_{2} = 3 \quad C_{2} = 1 \quad D_{2} = 9 \\ p_{3} = 8 \quad j_{3} = 1 \quad C_{3} = 1 \quad D_{3} = 10 \end{array}$$

$$P(s) = \hbar(0.08 + 1.52s^{3}) Watt$$
(b)
$$\Pi(\tau_{1}) = 1 \qquad \Pi(\tau_{2}) = 2 \qquad \Pi(\tau_{3}) = 3 \\ \Sigma(\tau_{1}) = 1 \qquad \Sigma(\tau_{2}) = 0.5 \qquad \Sigma(\tau_{3}) = 0.5 \end{array}$$

$$E_{b}(10s) = 3357mJ$$



Institute of Technology Simon Perathoner In Networks Laboratory

Example

$$\Gamma = \{ \tau_1, \tau_2, \tau_3 \}$$
 Task release: Periodic with non-deterministic
but bounded jitter
$$p_1 = 10 \quad j_1 = 3 \quad C_1 = 1 \quad D_1 = 1 \quad [ms] \\ p_2 = 5 \quad j_2 = 3 \quad C_2 = 1 \quad D_2 = 9 \\ p_3 = 8 \quad j_3 = 1 \quad C_3 = 1 \quad D_3 = 10$$
 $P(s) = \hbar(0.08 + 1.52s^3) Watt$

(c)
$$\Pi(\tau_1) = 1$$
 $\Pi(\tau_2) = 3$ $\Pi(\tau_3) = 2$
 $\Sigma(\tau_1) = 1$ $\Sigma(\tau_2) = 0.445$ $\Sigma(\tau_3) = 0.445$ $E_c(10s) = 3163mJ$



|--|

Outline

- Motivation
- Problem Definition
- Analysis Framework
- Proposed Algorithms
- Experimental Evaluation
- Conclusion



RTC: Greedy Processing Component



RTC: Scheduling Analysis



Task schedulable if

 $WCRT \leq D$

RTC: Preemptive Fixed Priority Scheduling





Outline

- Motivation
- Problem Definition
- Analysis Framework
- **Proposed Algorithms**
- Experimental Evaluation
- Conclusion

Contributions

- Algorithm for static assignment of priorities
- Algorithm for static assignment of speeds
- Algorithm for combined optimization of priorities and speeds

Priority Assignment

- 1) Find Π for given Σ
 - Energy consumption is fixed
 - Π affects schedulability
 - \rightarrow Find $~\Pi~$ under which $~\Gamma~$ is schedulable, if it exists

Observation:

Order of high-priority tasks does not affect low-priority task



Priority Assignment Algorithm



Minimum Global Speed

2) Find efficient (Π, Σ) under the restriction of equal speeds for all tasks

Binary search procedure

- Choose a global speed s_{qlob}
- If there is no schedulable priority order increase s_{qlob}
- Otherwise decrease s_{glob}
- Repeat until minimum speed is approximated with precision ϵ



Speed Assignment

- 3) Find Σ for given Π
 - A global speed is usually not energy-efficient
 - One should reduce individual speeds to obtain a *tight* assignment
 - But how to choose the individual speeds?



- \rightarrow Intuitions:
 - Never run a low-priority task with a higher speed than a high-priority task
 - Use as close speeds as possible

Speed Assignment

3) Find Σ for given Π

Proposal: Use speed assignments with the following properties:

$$s_{pr}^1 \ge s_{pr}^2 \ge \dots \ge s_{pr}^N$$

$$s^{i-1}_{pr}>s^{i}_{pr}$$
 only if WCRT $au^{i-1}_{pr}=D^{i-1}_{pr}$

Bottleneck Algorithm

3) Find Σ for given Π



- Assign minimum global speed
- Find bottleneck
- Fix speeds until bottleneck
- Repeat procedure for lower tasks
- Stop when last task becomes bottleneck

Time complexity:

$$O(N^2 \log \frac{1}{\epsilon})$$

Potential Non-Optimality

Example:

Speed assignment according to proposed strategy: x = 0.5, y = 0.5



 \rightarrow As close speeds as possible are not always beneficial

Combined Priority and Speed Assignment

4) Find Π and Σ (general case)

Combination of both algorithms:



At each iteration of the bottleneck algorithm do priority reordering

Time complexity:



Simon Perathoner

Computer Engineering and

Networks Laboratory

Adaptations for Discrete Speeds

In the discrete case there is no notion of ,bottleneck' task (speed changes by arbitrary small ϵ are not possible)

 \rightarrow Repeat speed assignment at each priority level of the chain

Speed assignment:
$$O(N^2 \log M)$$
Brute-force algorithms
 $O(M^N)$ Combined assignment: $O(N^4 \log M)$ $O(N! \cdot M^N)$

1

Outline

- Motivation
- Problem Definition
- Analysis Framework
- Proposed Algorithms
- Experimental Evaluation
- Conclusion

Experimental setup

9 test cases

10 tasks per test case

Periodic event streams with jitter

Intel XScale 1GHz

 $S_{cont} = [0, 1]$

$$S_{discr} = \{0.15, 0.4, 0.6, 0.8, 1\}$$

$$P(s) = \hbar (0.08 + 1.52s^3) Watt$$

Parameters for test case No. 6 [ms]:

	$ au_1$	$ au_2$	$ au_3$	$ au_4$	$ au_5$	$ au_6$	$ au_7$	$ au_8$	$ au_9$	$ au_{10}$
р	14	28	15	30	28	7	25	25	30	22
j	20	35	3	19	35	5	30	20	0	15
d	2	2	0	0	4	0	5	0	0	0
С	1	1	1	1	1	1	2	1	1	2
D	21	1	53	406	276	41	6	342	4	181

 $\epsilon = 10^{-4}$

RTC Toolbox for Matlab www.mpa.ethz.ch

Computer Engineering and Networks Laboratory

30

Results

Comparison of 3 different assignment policies for $\Pi\,$ and $\Sigma\,$:

- a) **Priorities for minimum global speed**
- b) Fixed priorities from a), Bottleneck algorithm for speeds
- c) Combined optimization of priorities and speeds

Results for test case No. 6:

Policy \Priority	1	2	3	4	5	6	7	8	9	10
(a) Alg. 2	$ au_2$	$ au_9$	$ au_7$	$ au_{10}$	$ au_8$	$ au_6$	$ au_1$	$ au_5$	$ au_4$	$ au_3$
	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
(b) Alg. 2 + 3	$ au_2$	$ au_9$	$ au_7$	$ au_{10}$	$ au_8$	$ au_6$	$ au_1$	$ au_5$	$ au_4$	$ au_3$
	1.000	.9445	.9445	.9445	.9445	.9445	.9445	.5776	.5776	.5776
(c) Combined	$ au_2$	$ au_9$	$ au_7$	$ au_8$	$ au_1$	$ au_6$	$ au_3$	$ au_5$	$ au_4$	$ au_{10}$
	1.000	.7500	.7500	.5914	.5914	.5914	.5914	.5914	.5914	.5914

Simulated worst-case energy consumption for an execution of 10⁴ ms:



Computer Engineering and Networks Laboratory

Outline

- Motivation
- Problem Definition
- Analysis Framework
- Proposed Algorithms
- Experimental Evaluation
- Conclusion

Conclusion

- New methods for the design of energy-efficient real-time systems
- Static assignment of priorities and speeds to tasks with arbitrary nondeterministic release-times described by arrival curves
- Algorithms are based on two observations:
 - 1) Rearranging high-priority tasks does not affect low-priority tasks
 - 2) Use of priority-monotonic speeds is beneficial
- Combined optimization of priorities and speeds gives best results



Thank you!

Simon Perathoner

perathoner@tik.ee.ethz.ch

