Process Cruise Control

on Intel Atom Processors

Martin Farkas-Schmid Sebastian Klaas

Department of Computer Science
Karlsruhe Institute Of Technologie (KIT)

Power Management Lab Winter Term 2009/2010

Outline

1. Introduction
   - Process Cruise Control
2. Our Implementation
3. Results
   - Benchmarking Results
   - Power Consumption
4. Possible Extensions / Improvements

General Idea

- power management policy based on frequency scaling
- using embedded hardware monitors
- individual clock frequency for each task
- performance should only suffer slightly

Trade-offs / Restrictions

- power consumption vs. QoS
- does only work on architectures providing performance counters

The Policy

- find set of countable events, characterizing tasks
- make partitions (frequency domains) - each according to an optimal speed
- determine the optimal frequency at each task-switch

Settings

- performance counters count:
  - retired instructions
  - memory bus activity
- maximum performance loss: 15%

Performance Measurement

- evaluate performance counter data at each task-switch
- calculate ratio: (retired instructions) div (bus activity)
- store ratio in task_struct
- store new frequency value in task_struct
- values are stored in task_struct

Frequency Switch

- immediately before a task is scheduled
- set frequency according to value in task_struct
- switch functions provided by a new governor

Frequency Domains

Instr. per Mem.act.	According Frequency
0 - 50	800 Mhz
50 - 100	1067 Mhz
100 - 200	1333 Mhz
200 - ...	1600 Mhz

Some Examples

Program	Avg. Instr. per Mem.act.	Best Frequency
cksum	68-86	1333 MHz
gzip	108-160	1600 MHz

- (values are examples - recorded with an old buggy kernel version ...)

Linux 2.6.30

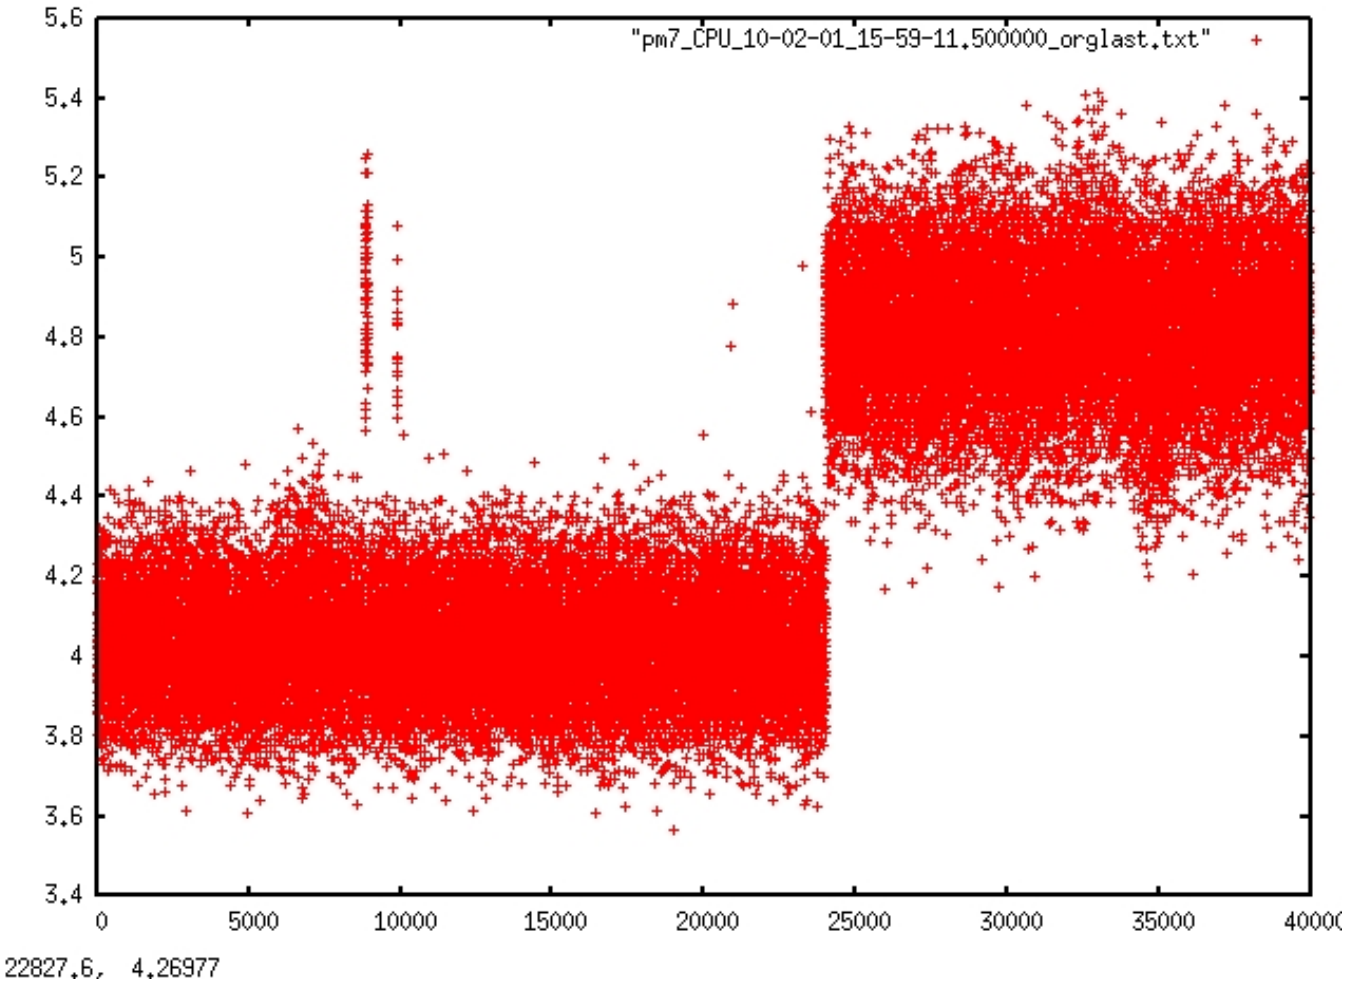
"pm7_CPU_10-02-01_15-59-11.500000_orglast.txt"
5.6
5.4
5.2
5
4.8
4.6
4.4
4.2
4
3.8
3.6
3.4
0
5000
10000
15000
20000
25000
30000
35000
40000
22827.6, 4.26977

Linux 2.6.30 with PCC

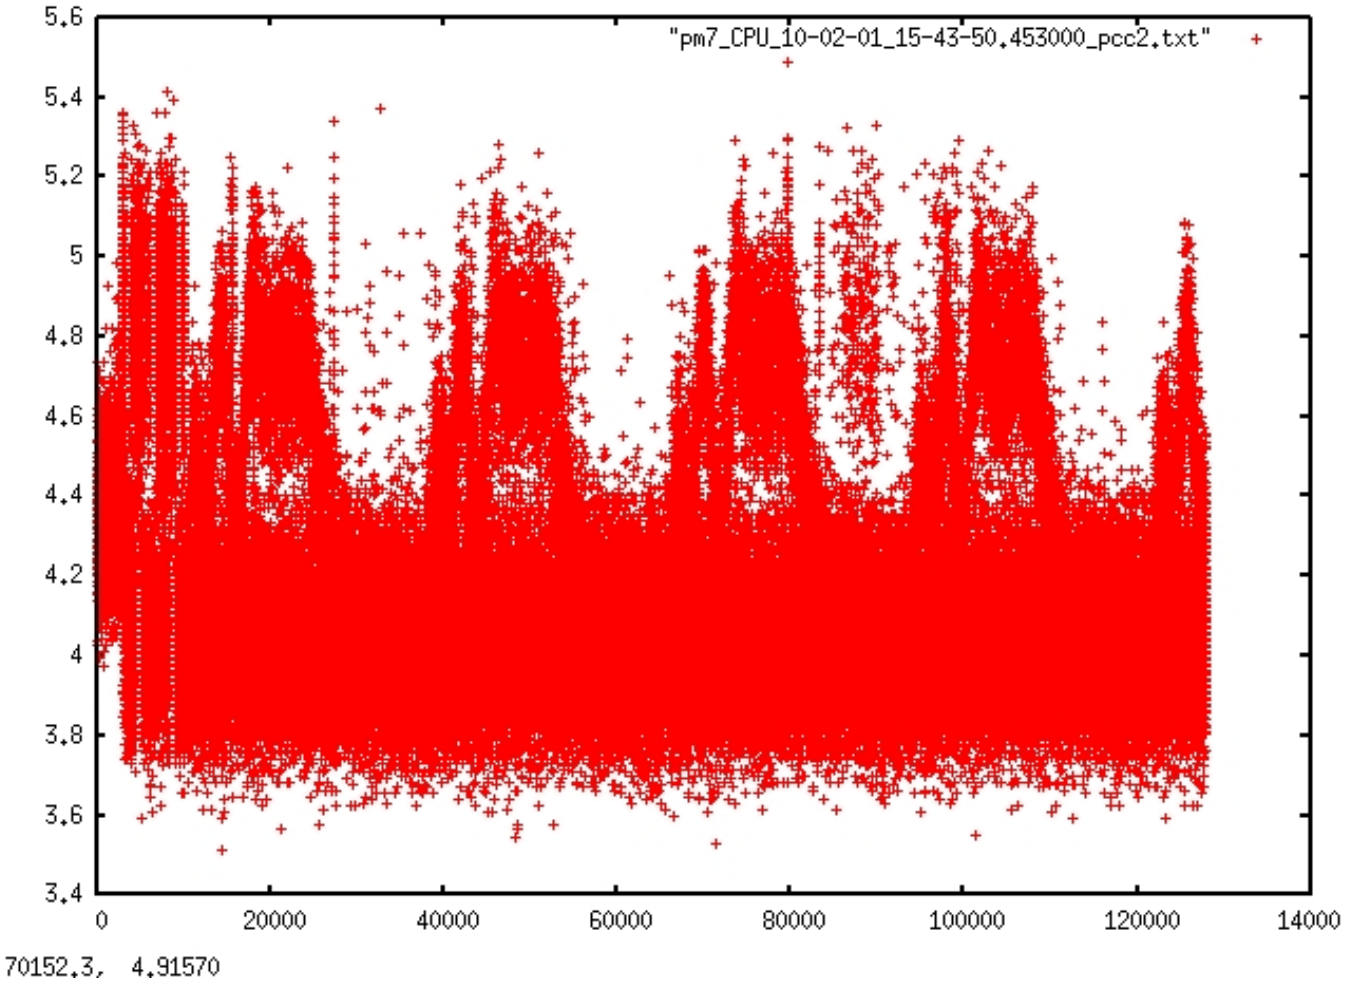

Linux 2.6.30 with PCC

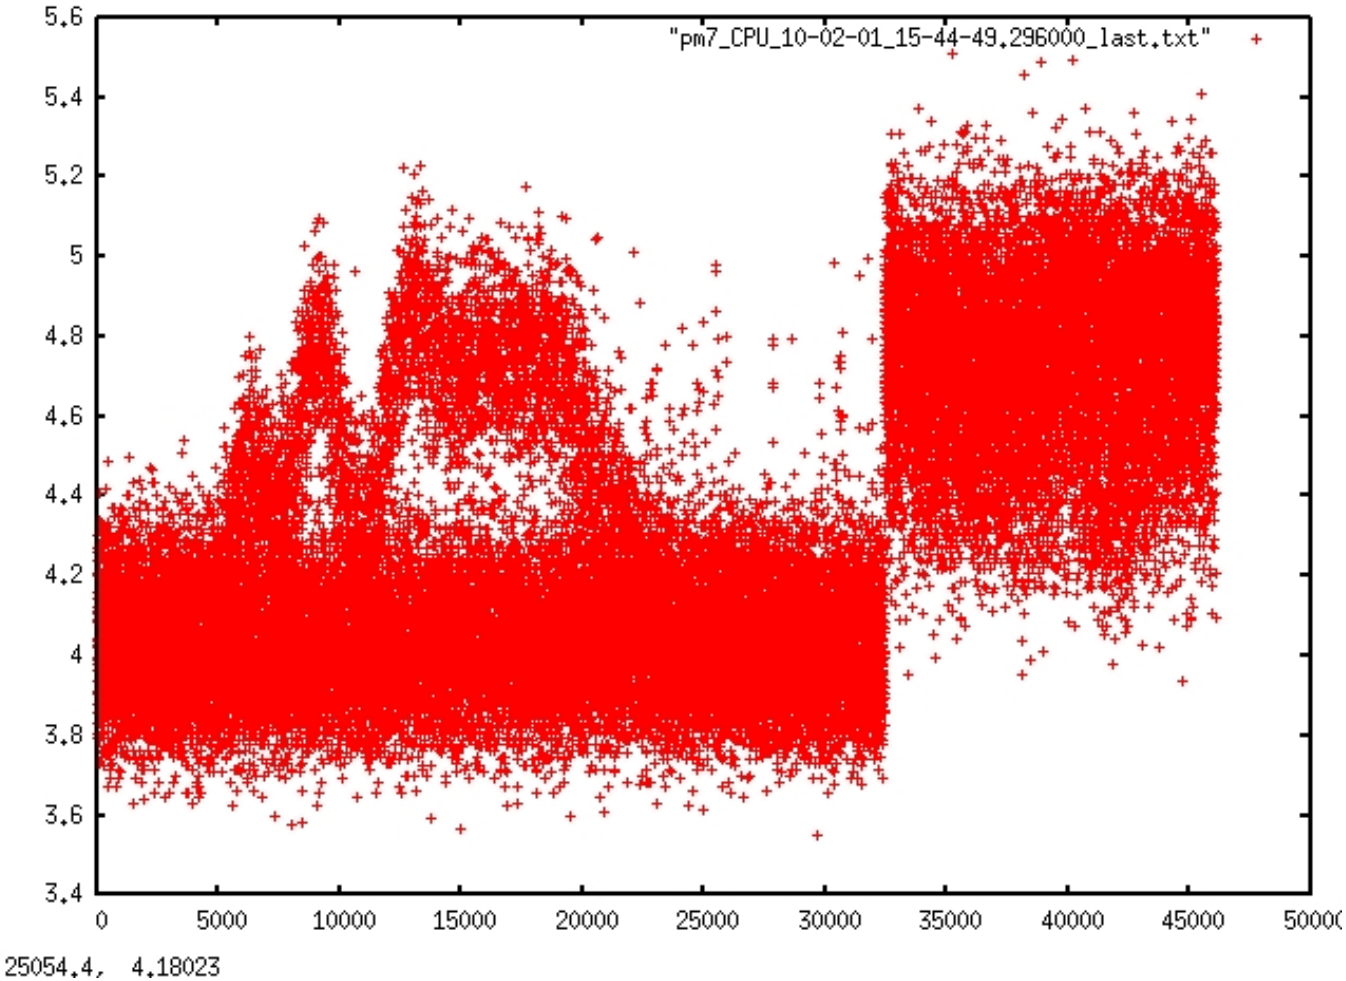

Possible Extensions / Improvements

- let the user decide about the acceptable loss of performance
- frequency decision based on more than one timeslice
- improve frequency domains (more dimensions, more detailed)

Literature

Andreas Weissel, Frank Bellosa:
Process Cruise Control: Event-driven clock scaling for dynamic power managment
Cases 2002, October 8-11, 2002, Grenoble France
http://doi.acm.org/10.1145/581630.581668