

Effects of Clock Resolution on the Scheduling of Interactive and Soft Real-Time Processes

by Yoav Etsion, Dan Tsafir, Dror G. Feitelson

Presented by James Volpe

Key Points

- Modern interactive applications not supported in commodity operating systems
- Instead of specialized APIs, tune commodity operating systems
- Clock interrupt rates unchanged last 30 years
- Increasing clock interrupt rate can help
- Additional overhead is acceptable

2

Real-time Problem

- Xine MPEG player
- Short clip of 500 frames

3

Process Scheduling

4

Process Scheduling Cont.

- Commodity operating systems use priority-based scheduling
- Static and dynamic priorities
- High CPU usage, the lower the priority
- Conflict with modern multimedia applications

5

Real-time Problem Up Close

6

Related Work

- RT-Linux, one-shot timers, soft timers, firm timers and priority adjustments
- Require special APIs and/or non-trivial modifications to the system
- Why not increase the clock interrupt rate and see what happens?

7

Motivation

- 100 Hz clock interrupt rates
 - 10 MHz CPU 100,000 instructions/interrupt
 - 1 GHz CPU 10,000,000 instructions/interrupt
- Clock interrupt rate has become too coarse
 - Two orders of magnitude of additional instructions per interrupt

8

Benefits of High Tick Rate

- Take advantage of today's faster processors
- Better timing for meeting deadlines
- Accurate billing of processes


9

Test Platform

- Pentium 90 to Pentium-IV 2.4 GHz
- Measurements done on 664 MHz Pentium III
- Linux kernel 2.4.8 (RedHat 7.0)
- Klogger used for logging scheduling related events: context switching, recalculation of priorities, forks, execs
- Workload from Emacs, Xine, Quake 3, CPU-bound processes


10

Timing Advantage



11

Timing Advantage



12

Billing Advantage


Application	Billing ratio		Missed quanta	
	@100Hz	@1000Hz	@100Hz	@1000Hz
Emacs	1.0746	0.9468	95.96%	73.42%
Xine	1.2750	1.0249	89.46%	74.81%
Quake	1.0310	1.0337	54.17%	23.23%
X Server*	0.0202	0.9319	99.43%	64.05%
CPU-bound	1.0071	1.0043	7.86%	7.83%
CPU+Quake	1.0333	1.0390	26.71%	2.36%

* When running Xine

Table 3. Scheduler billing success rate

13

Billing Advantage Cont.



14

The Cost Is Overhead

- More clock interrupts increases overhead
- Time needed to process interrupts
- Higher number of context switches
- Operating systems become less efficient as clock rate increases

15

Clock Interrupt Overhead

- Do operating systems become as fast as hardware?

Processor	Default		Without 8253	
	Cycles	µs	Cycles	µs
P-90	814±180	9.02	498±466	5.53
PP-200	1654±553	8.31	462±762	2.32
PII-350	2342±303	6.71	306±311	0.88
PIII-664	3972±462	5.98	327±487	0.49
PIII-1.133	6377±602	5.64	426±914	0.38
PIV-2.4	14603±436	6.11	445±550	0.19
A1.467	10494±396	7.15	202±461	0.14

Table 1. Interrupt processing overheads on different processor generations

16

Clock Interrupt Overhead Cont.


- Do operating systems become as fast as hardware?

Processor	Context switch		Cache BW MB/s	Trap	
	Cycles	µs		Cycles	µs
P-90	1871±656	20.75	28±1	153±24	1.70
PP-200	1530±389	7.69	705±26	379±75	1.91
PII-350	1327±331	3.80	1314±29	343±68	0.98
PIII-664	1317±424	1.98	2512±32	348±163	0.52
PIII-1.133	1330±441	1.18	4286±82	364±278	0.32
PIV-2.4	3792±857	1.59	3016±47	1712±32	0.72
A1.467	1436±477	0.98	3962±63	274±20	0.19

Table 2. Other overheads on different processor generations

17


Clock Interrupt Overhead Cont.



- Overhead from sorting an array with 50ms quantum

18

Clock Interrupt Overhead Cont.



- Overhead from sorting an array with quantum equal to 5 ticks

19

Conclusion

- Increasing clock interrupt rate to 1000 Hz achieves timing and billing advantages
- Overhead is minimal
- Suggest making a settable parameter instead of compiled constant

20

Evaluation

- Results are promising
- Production level example would strengthen argument
- Changing clock interrupt rate in Linux not an end user job
- Testing is needed for implementation

21

Questions / Discussion



22