Benchmarking AMD64 and EMT64

Computing in High Energy and Nuclear Physics

13-17 February 2006, T.I.F.R. Mumbai, India

Computing Facilities and Networking: CFN-6

02/15/06

Hans Wenzel, Oliver Gutsche (USCMS / Fermilab)
Mako Furukawa (University of Nebraska)

Introduction

- Comparison of benchmarks of single-core 32bit and single- and dual-core 64bit processors
- Architecture has influence on important parameters for high-performance large size computing installations
  - Performance / computing power
  - Cooling requirements
  - Electricity consumption
- Results help in decision process for future procurements

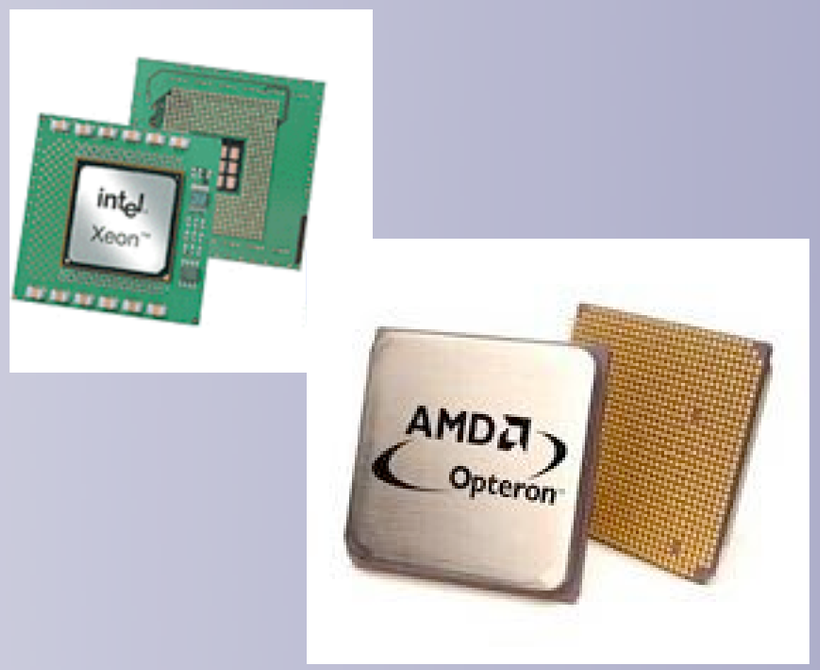

Outline

- Single- vs. dual-core
- Compared processors
- Operation modes: 32bit, compatibility and 64bit
- Test environment and list of used applications
- Benchmark comparisons for the different applications
- Power consumption benchmark
- Summary & Outlook

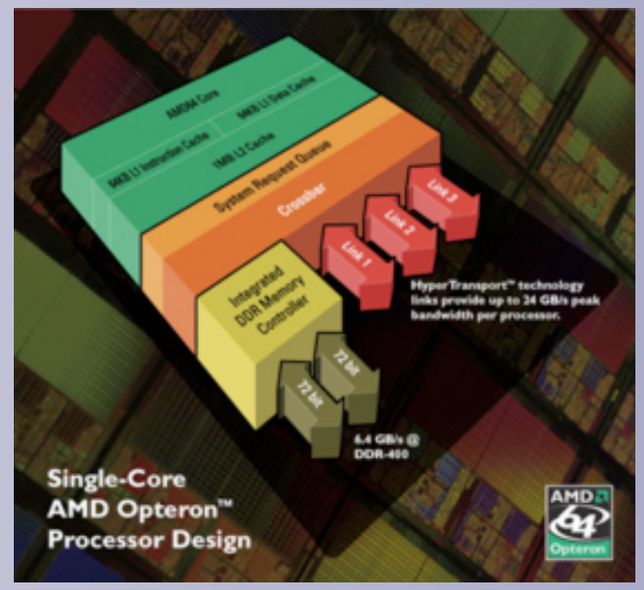

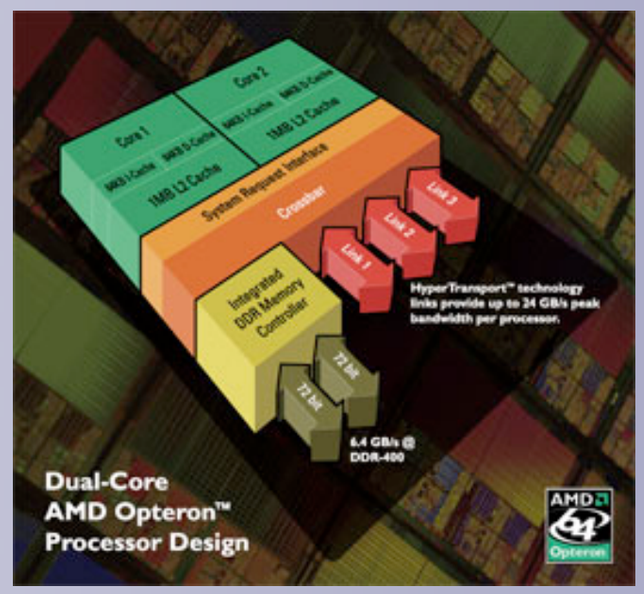

Single- vs. dual-core processors

- Dual-core processors promise to nearly double computing power
  - Without increasing infrastructure costs for cooling and electricity
- Dual-core processors provide:
  - Lower space requirements and need for fewer racks
  - Fewer console connections
  - Fewer network connections
    - Might adversely affect performance
    - Same network interface is shared by 4 instead of 2 cpu-cores
    - more processors competing for the same bandwidth (same is true for disk IO)

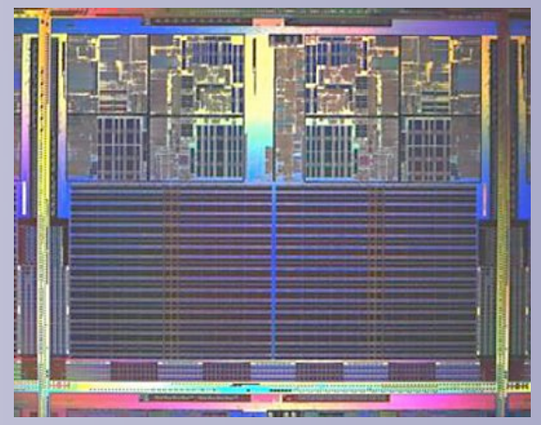

➡ **Half the cooling and electricity cost for the same CPU power**

Processors

- AMD Athlon 64 and Xeon test have been made on single processor machines, all other tests have been made on dual processor machines

➡ up to 4 processes in parallel possible on dual-core machines

Processor	Speed	Architecture	Cores
Opteron 244	1.8 GHz	AMD-64	Single
Opteron 265	1.8 GHz	AMD-64	Dual
Opteron 246	2.0 GHz	AMD-64	Single
Opteron 246HE	2.0 GHz	AMD-64	Single
Opteron 270	2.0 GHz	AMD-64	Dual
Opteron 248	2.2 GHz	AMD-64	Single
Opteron 275	2.2 GHz	AMD-64	Dual
Opteron 250	2.4 GHz	AMD-64	Single
AMD Athlon 64 3500+	2.2 GHz	AMD-64	Single
AMD Athlon 64 4200+	2.2 GHz	AMD-64	Dual
Xeon	3.4 GHz	EMT64	Single
Xeon	3.6 GHz	EMT64	Single
Xeon	2.4 GHz	IA32	Single
Xeon	2.8 GHz	IA32	Single

Operation modes for 64bit processors

- 32-bit legacy mode (32/32):
  - processors will act like any other IA32 compatible processor
    - 32-bit OS and 32-bit applications
    - not possible to use flat memory addressing above 4GB and additional General Purpose Registers (GPRs)
- Compatibility mode (64/32):
  - intermediate mode using 64-bit OS and 64-bit drivers
    - **enabled to support both 32-bit and 64-bit applications**
    - each 32-bit application limited to 4GB memory (imposed on per-process level, huge improvement to 32/32)
- Full 64-bit mode (64/64):
  - 64-bit OS and 64-bit applications recompiled to take full advantage of the various enhancements of the 64-bit addressing architecture
    - single application can have a virtual address space up to 40-bits (1TB of addressable memory)
    - access to full physical memory range and access to new GPRs as well as expanded GPRs

Test environment

- Scientific Linux 3.03/3.04
  - available in 32bit and 64bit versions

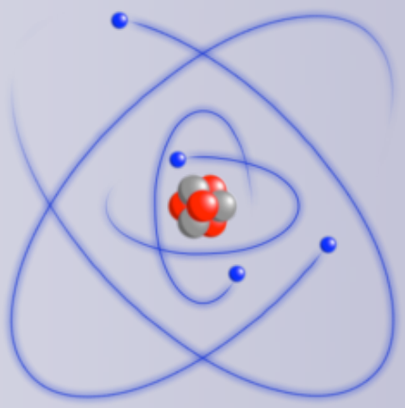

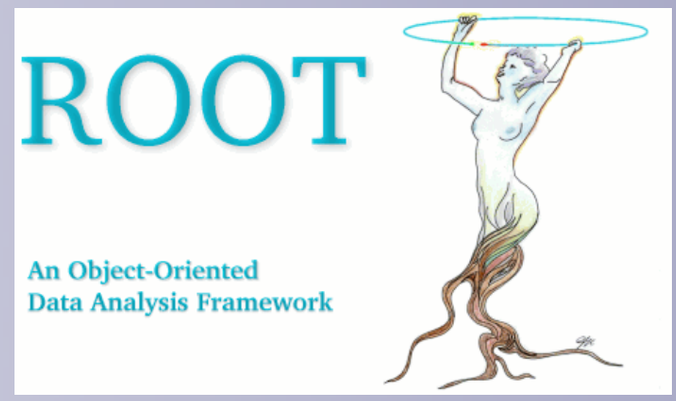

- CMS specific OO applications:
  - OSCAR: GEANT 4 based CMS detector simulation (OSCAR_3_7_0)
  - ORCA: CMS OO reconstruction program (ORCA_8_7_5)
    - 32bit applications ➡ port to 64bit not available yet

➡ only running modes: (32/32) and 64bit compatibility mode (64/32)

- General HEP applications:
  - ROOT 4.02/00 for 32bit and 64bit, recompiled for 64bit using gcc 3.2.3 from SL 3.04 (”configure linuxx8664gcc”)
  - PYTHIA 6.227 for 32bit and 64bit, recompiled for 64 bit using g77 based on gcc 3.2.3 from SL 3.04

Test applications

- ROOT: stress benchmark
- PYTHIA: generate 100,000 SUSY events at $\sqrt{s}$=14TeV
  - PYTHIA main65.f example: compiled with O2 option (g77 -O2 -o main65 main65.f pythia6227.o)
- OSCAR: simulate 300 single pion events of p_T=50GeV
- ORCA: digitize OSCAR output (writeAllDigis) and reconstruct them (writeDST)

Application	Operating mode		
	(32/32)	(64/32)	(64/64)
PYTHIA	Yes	Yes	Yes
ROOT	Yes	Yes	Yes
OSCAR	Yes	Yes	N/A
ORCA (Digi)	Yes	Yes	N/A
ORCA(DST)	Yes	Yes	N/A

Pythia

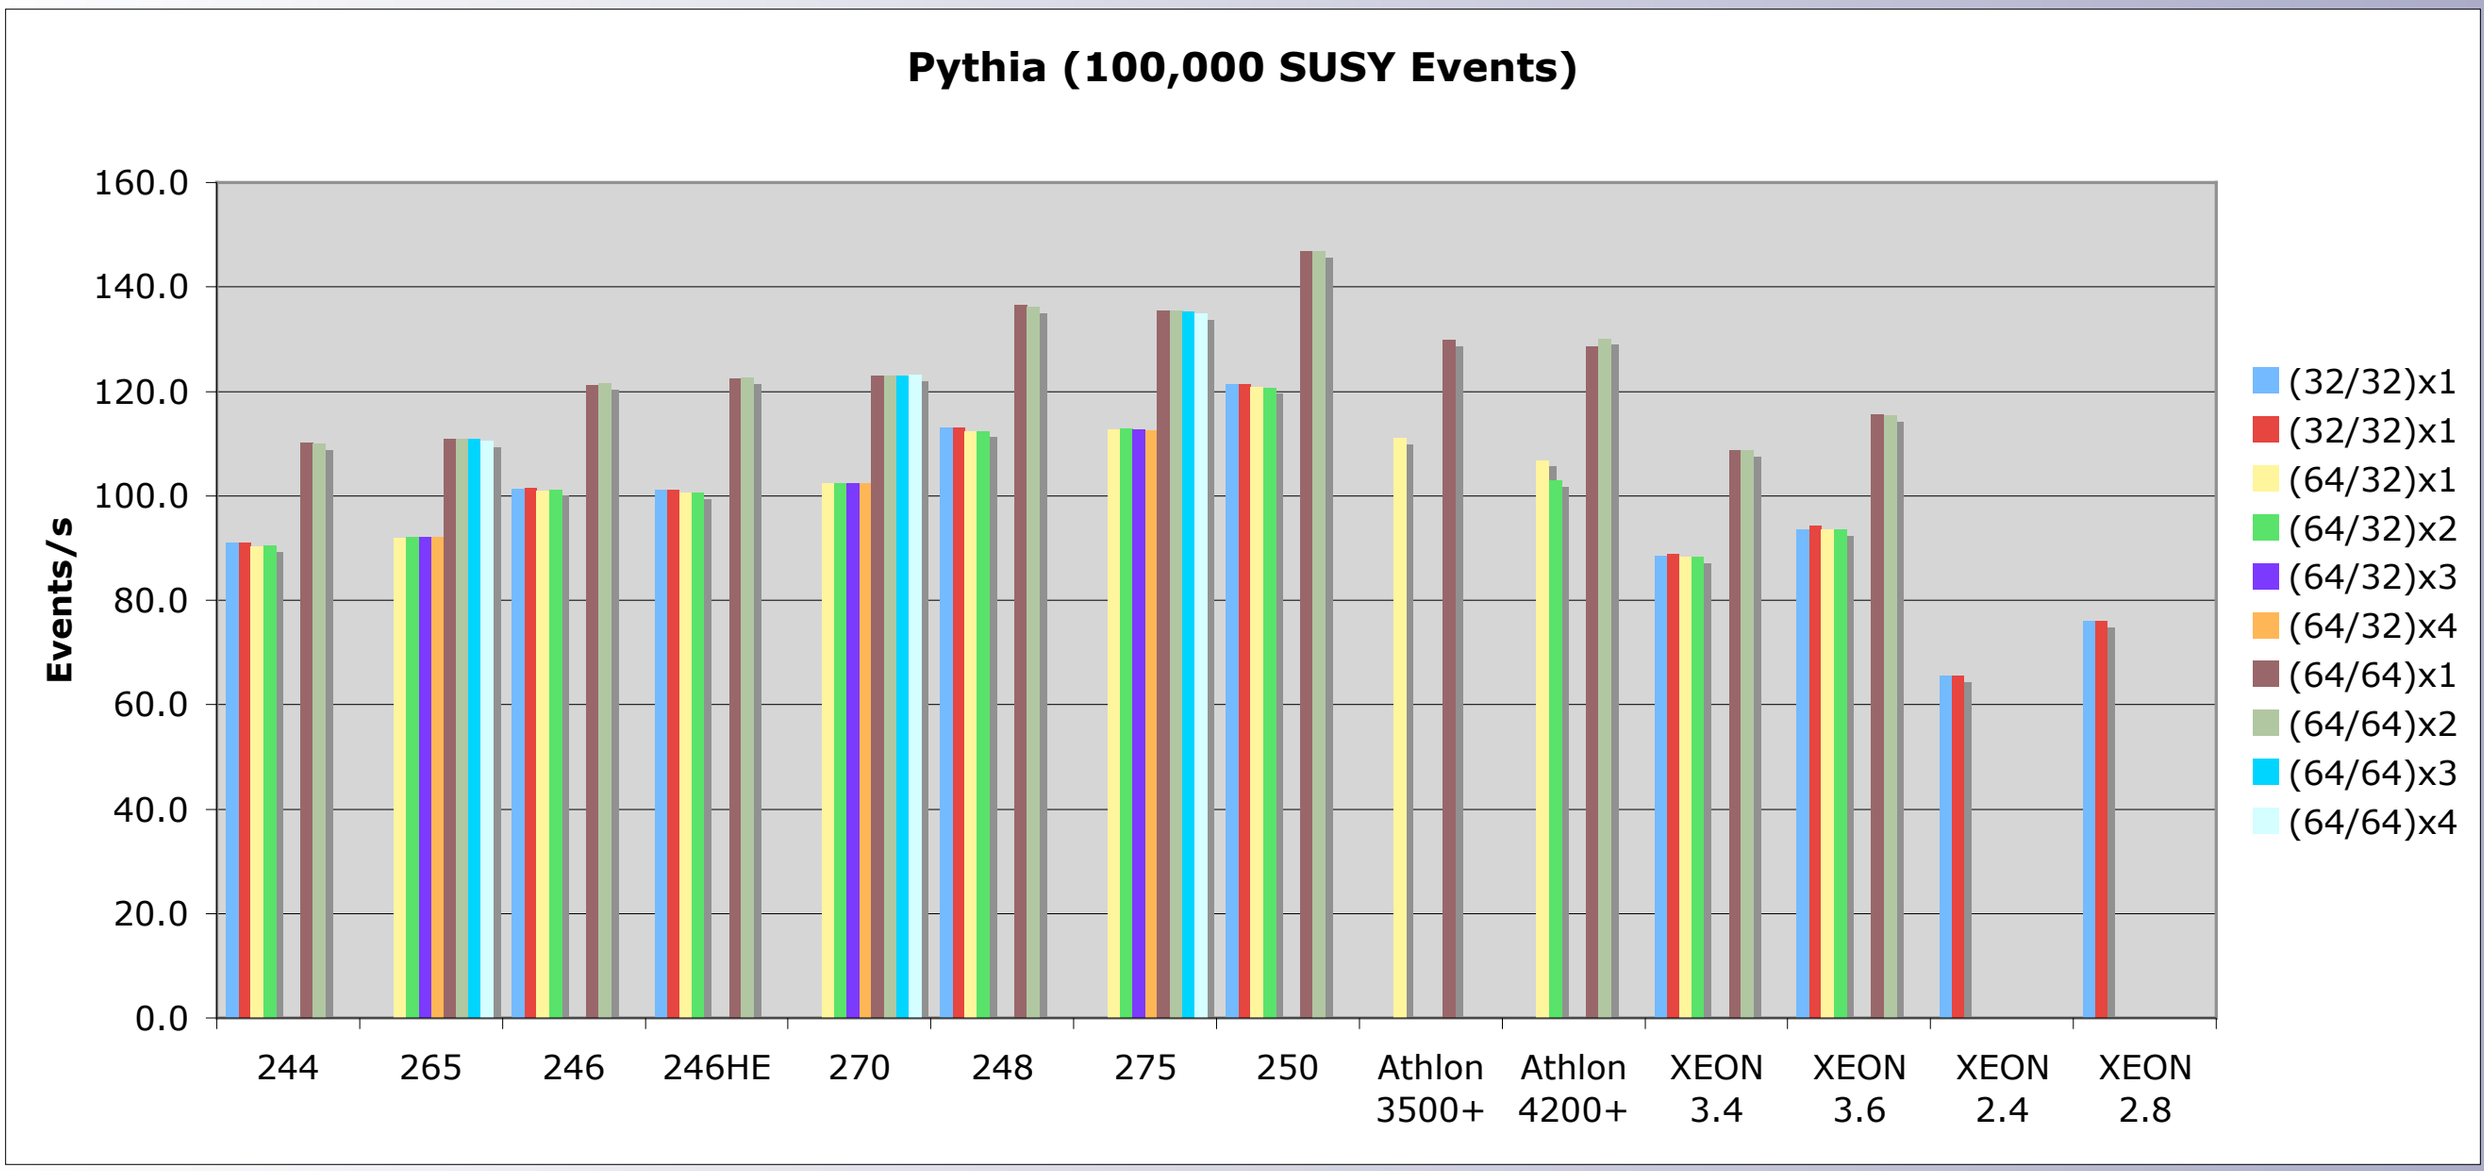

- 20% performance boost from (64/32) to (64/64)

Pythia

[events/s]	(32/32)x1	(32/32)x1	(64/32)x1	(64/32)x2	(64/32)x3	(64/32)x4	(64/64)x1	(64/64)x2	(64/64)x3	(64/64)x4
244	91.0	91.0	90.3	90.4			110.1	110.0		
265			92.0	92.1	92.1	92.0	110.9	110.9	110.8	110.4
246	101.3	101.4	100.9	101.1			121.2	121.5		
246HE	101.1	101.2	100.5	100.6			122.4	122.5		
270			102.4	102.4	102.4	102.3	123.0	123.0	123.0	123.1
248	113.0	112.9	112.3	112.3			136.4	136.1		
275			112.7	112.8	112.6	112.5	135.3	135.3	135.2	134.9
250	121.3	121.3	120.8	120.7			146.9	146.8		
Athlon 3500+			111.1				129.8			
Athlon 4200+			106.7	102.9			128.6	130.0		
XEON 3.4	88.5	88.7	88.1	88.3			108.6	108.6		
XEON 3.6	93.5	94.2	93.4	93.4			115.6	115.4		
XEON 2.4	65.5	65.5								
XEON 2.8	75.9	75.9								

- No performance drop when running 4 processes in parallel

ROOT Stress

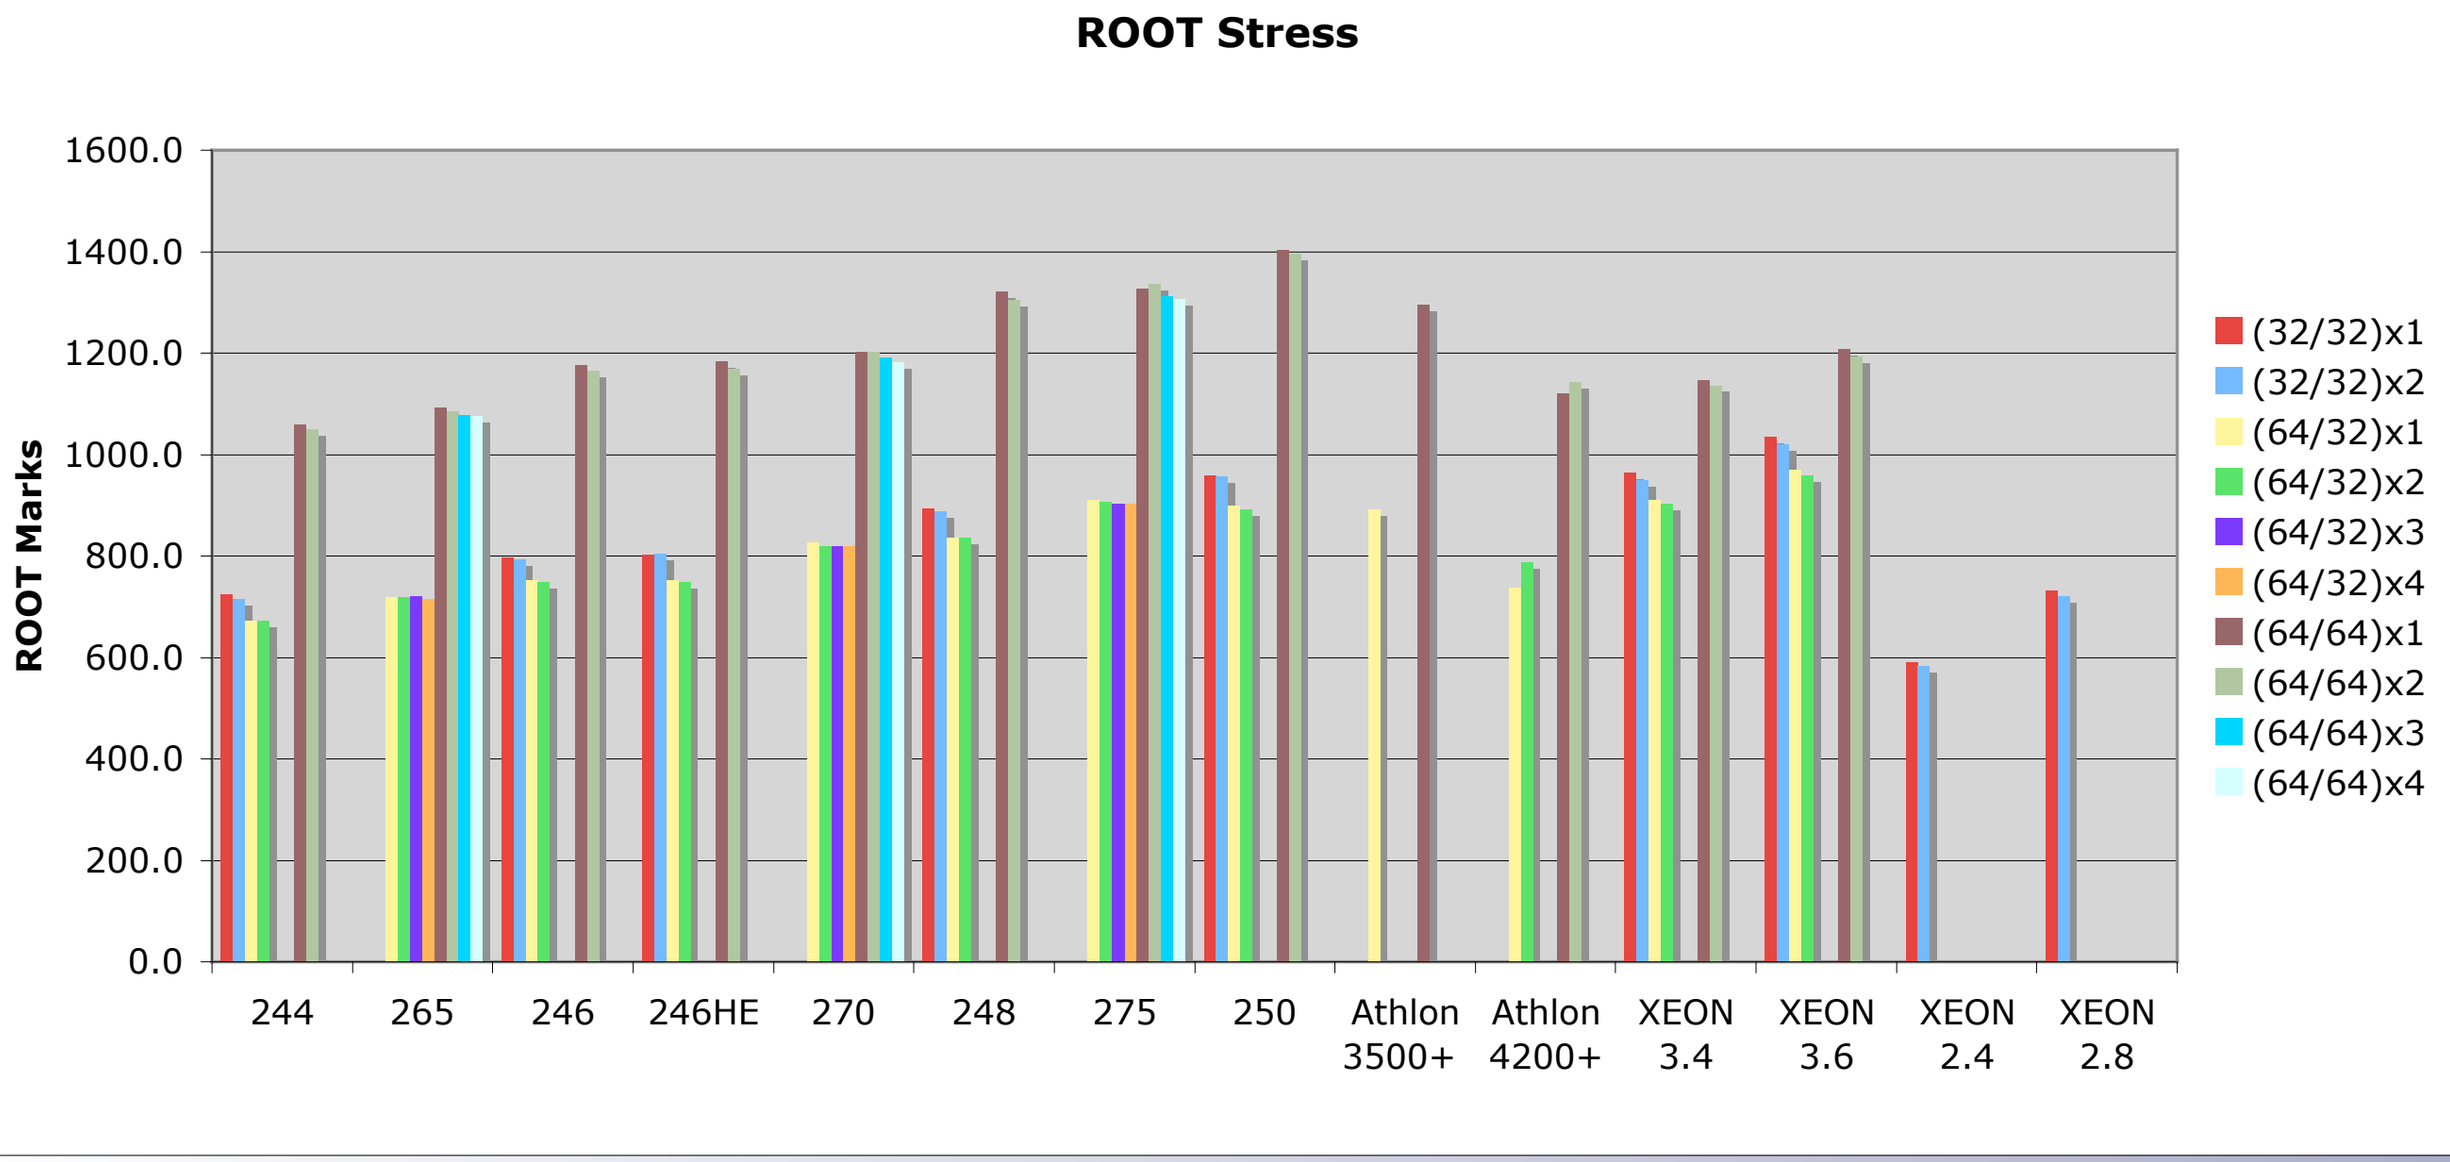

- Very small performance drop of 1.6% when running 4 processes in full 64bit mode in parallel
- 6% slower in compatibility mode (64/32) compared to (32/32)

ROOT Stress

[ROOT marks]	(32/32)x1	(32/32)x1	(64/32)x1	(64/32)x2	(64/32)x3	(64/32)x4	(64/64)x1	(64/64)x2	(64/64)x3	(64/64)x4
244	723.6	715.4	672.5	673.3			1059.9	1050.8		
265			719.3	718.7	719.6	715.7	1093.3	1085.4	1077.1	1075.9
246	797.1	793.5	751.6	748.3			1176.1	1165.7		
246HE	801.0	805.1	752.2	749.1			1183.9	1170.4		
270			827.3	819.3	819.8	819.3	1201.6	1201.8	1190.9	1182.8
248	894.2	888.4	834.7	835.8			1320.8	1306.0		
275			910.2	907.1	904.2	903.0	1328.0	1335.1	1311.1	1306.0
250	957.6	956.7	898.3	891.4			1402.6	1395.6		
Athlon 3500+			891.4				1295.9			
Athlon 4200+			738.6	788.0			1120.0	1142.0		
XEON 3.4	964.6	950.1	909.3	903.8			1145.9	1135.8		
XEON 3.6	1034.8	1020.6	969.7	958.7			1208.6	1194.4		
XEON 2.4	590.7	583.8								
XEON 2.8	731.0	720.1								

- Opterons run 32% faster in (64/64) mode compared to (32/32)
- Xeons only gain 16%
- Opteron 250 runs in (64/64) 2.4 times faster than our current worker nodes (2.4 xeons in (32/32))

OSCAR

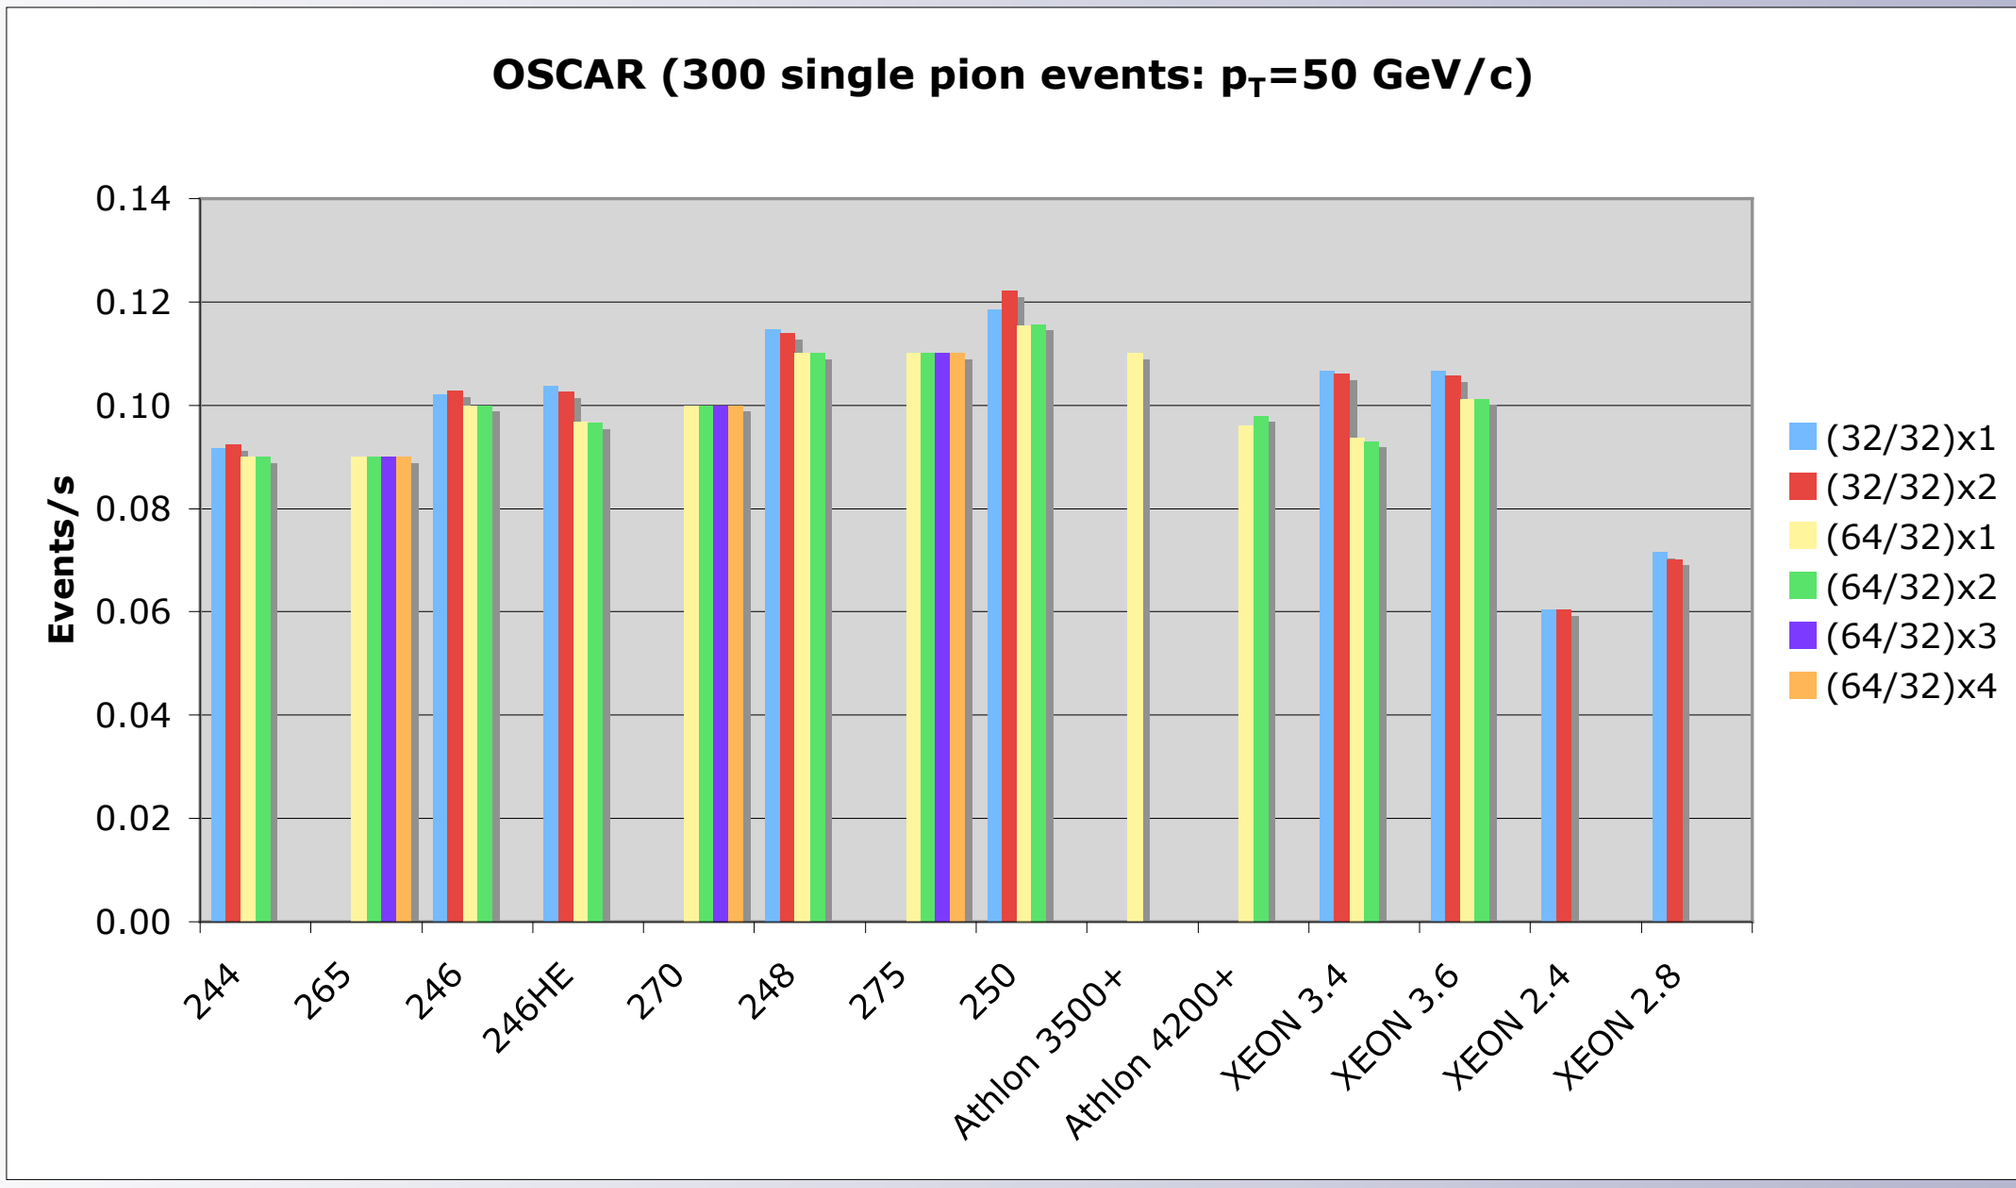

- Scales well with increasing number of parallel processes➡ CPU intensive and not I/O intensive
- Performance drop only 0.9% / 2.2% / 1.6% for 265 / 270 / 275 for 4 processes in parallel

OSCAR

	(32/32)x1	(32/32)x1	(64/32)x1	(64/32)x2	(64/32)x3	(64/32)x4
244	0.09	0.09	0.09	0.09		
265			0.09	0.09	0.09	0.09
246	0.10	0.10	0.10	0.10		
246HE	0.10	0.10	0.10	0.10		
270			0.10	0.10	0.10	0.10
248	0.11	0.11	0.11	0.11		
275			0.11	0.11	0.11	0.11
250	0.12	0.12	0.12	0.12		
Athlon 3500+			0.11			
Athlon 4200+			0.10	0.10		
XEON 3.4	0.11	0.11	0.09	0.09		
XEON 3.6	0.11	0.11	0.10	0.10		
XEON 2.4	0.06	0.06				
XEON 2.8	0.07	0.07				

- 4% slower in compatibility mode (64/32) compared to (32/32)

ORCA Digitization

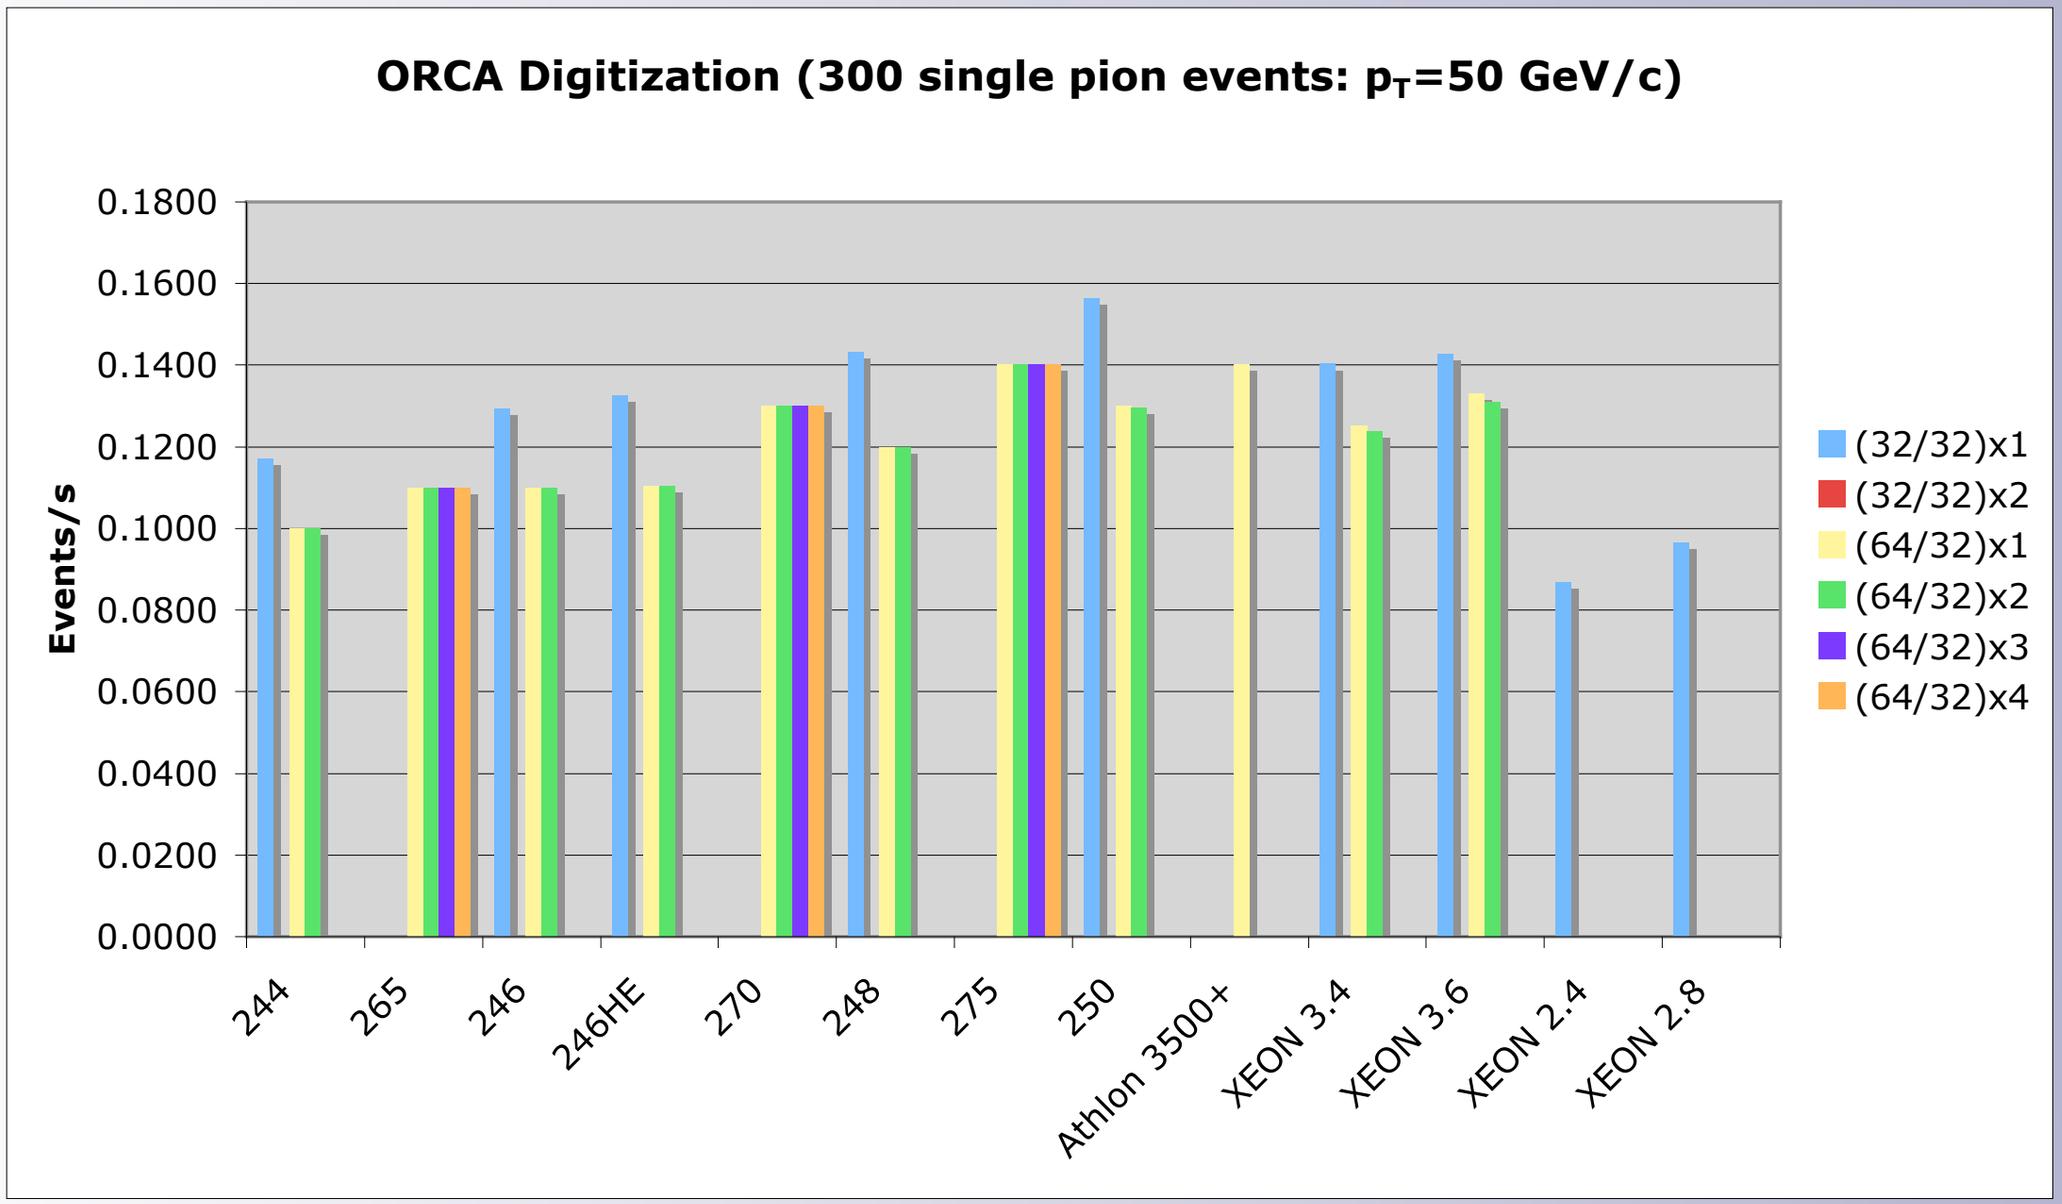

Performance drop 0.6% / 3.3% / 2.8% for 265 / 270 / 275 for 4 processes in parallel

➡ Increased I/O activity ?

ORCA Digitization

	(32/32)x1	(32/32)x1	(64/32)x1	(64/32)x2	(64/32)x3	(64/32)x4
244	0.12		0.10	0.10		
265			0.11	0.11	0.11	0.11
246	0.13		0.11	0.11		
246HE	0.13		0.11	0.11		
270			0.13	0.13	0.13	0.13
248	0.14		0.12	0.12		
275			0.14	0.14	0.14	0.14
250	0.16		0.13	0.13		
Athlon 3500+			0.14			
XEON 3.4	0.14		0.13	0.12		
XEON 3.6	0.14		0.13	0.13		
XEON 2.4	0.09					
XEON 2.8	0.10					

- 15% slower in compatibility mode (64/32) compared to (32/32)

ORCA (DST)

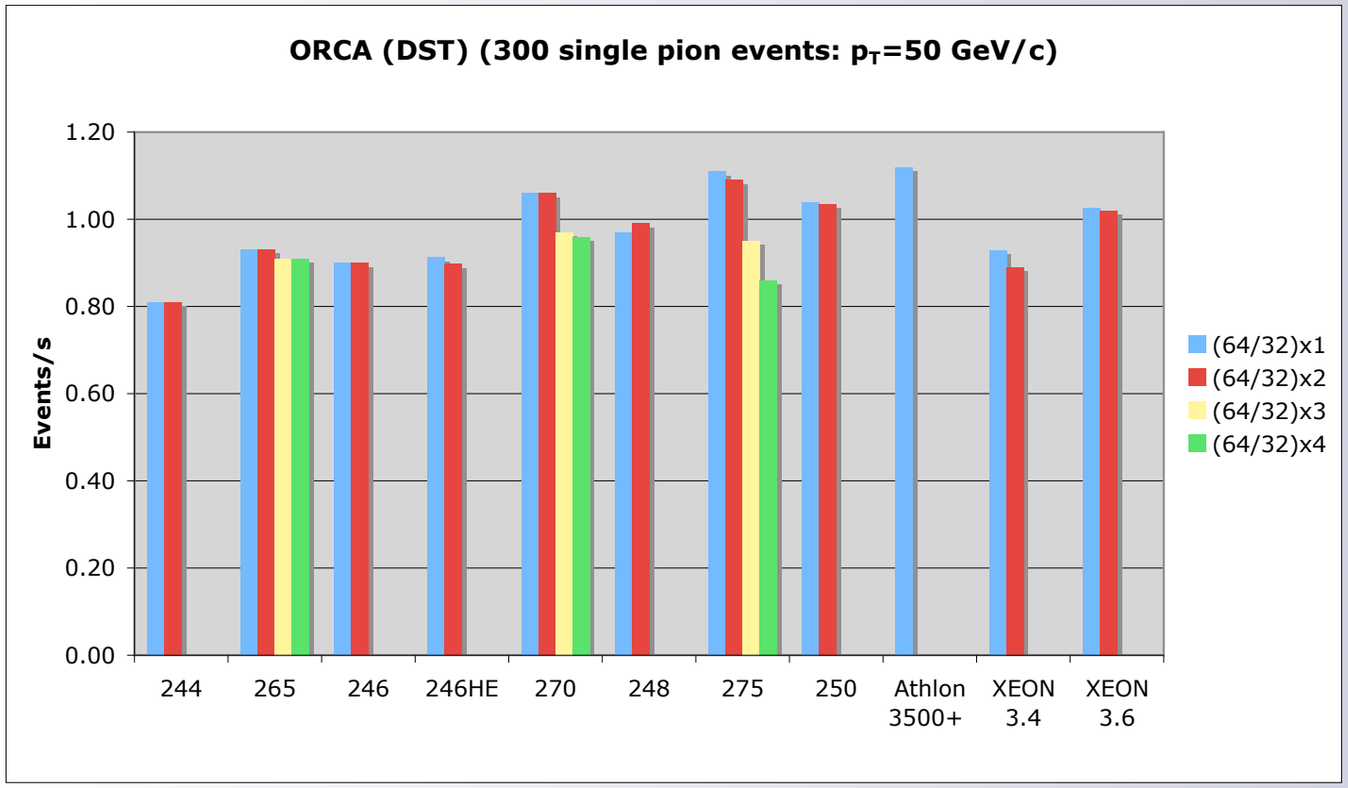

	(64/32)x1	(64/32)x2	(64/32)x3	(64/32)x4
244	0.81	0.81		
265	0.93	0.93	0.91	0.91
246	0.90	0.90		
246HE	0.91	0.90		
270	1.06	1.06	0.97	0.96
248	0.97	0.99		
275	1.11	1.09	0.95	0.86
250	1.04	1.04		
Athlon 3500+	1.12			
XEON 3.4	0.93	0.89		
XEON 3.6	1.03	1.02		
XEON 2.4				
XEON 2.8				

- Performance drop 2.2% / 9.4% / 22.5% for 265 / 270 / 275 for 4 processes in parallel

➡ very I/O dependent

Power consumption benchmark

- Complicated to achieve meaningful power consumption benchmark, requires equal measurement conditions
  - power connection, measurement device, etc.
  - measurement done using endless loop of π calculation and "dd"

Processor	Core	Idle	Loaded
244	Single	1.20	1.40
248	Single	1.40	1.60
270	Dual	1.20	1.70

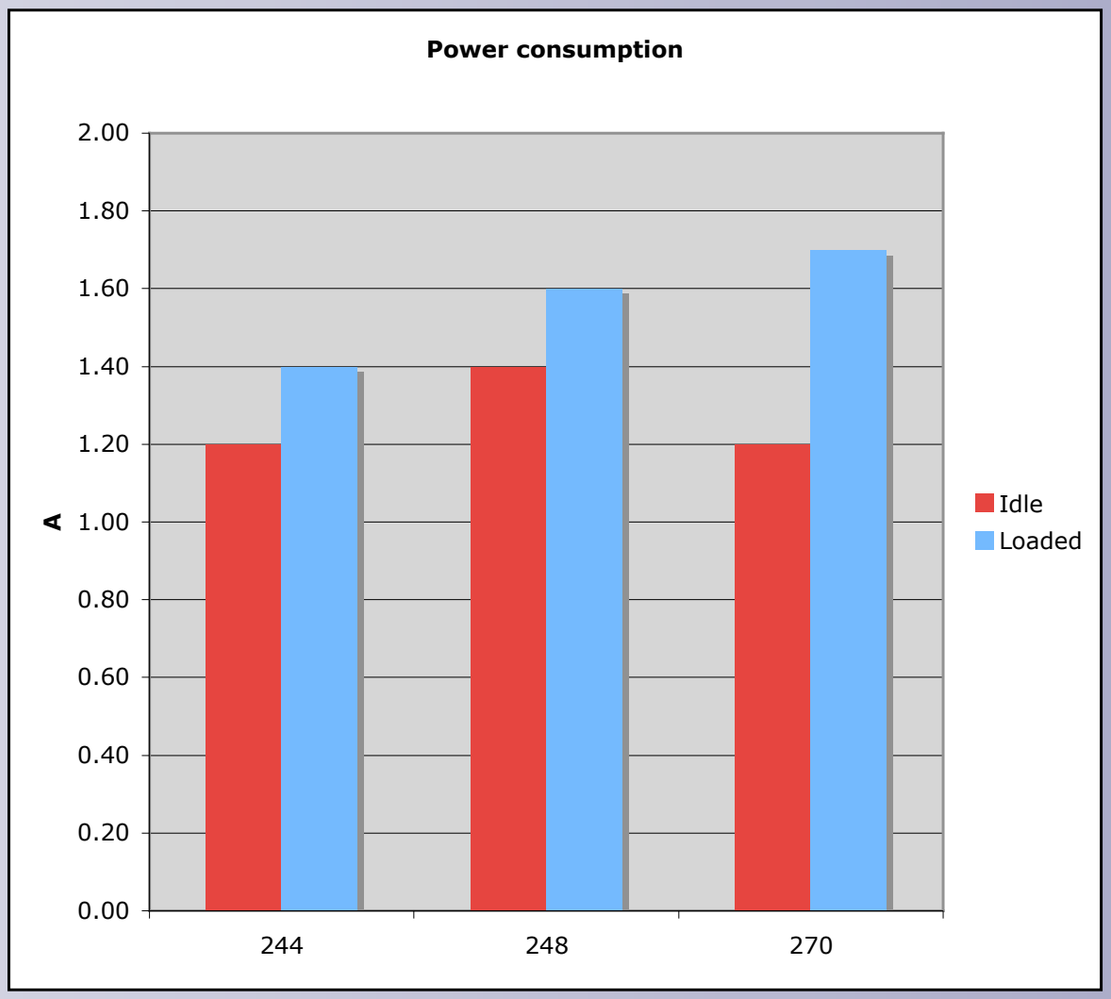

- Dual-core consumes almost the same electricity for twice the computing power

Summary & Outlook

- Todays 32bit applications can be executed by the tested 64bit machines without porting, either in 32bit mode or in 64bit compatibility mode. This allows for a painless transition to this platform without any risk of incompatibilities.

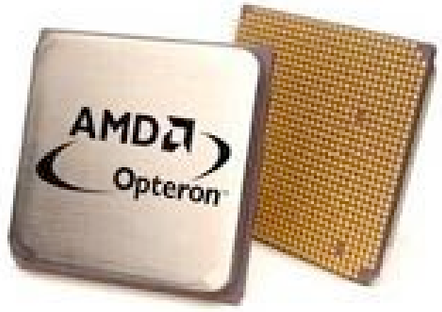

- The 64bit architecture promises a big increase in computing power for ported applications

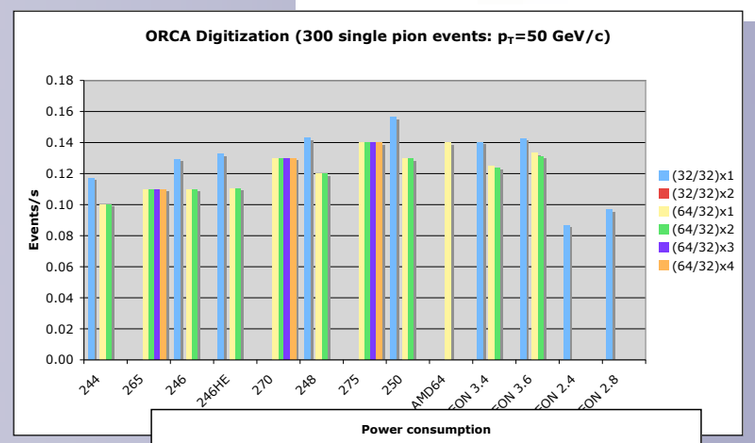

- Dual-core processors provide almost twice the computing power compared to single-core processors at the same electricity consumption and cooling needs. Also multicore processors are expected to be available soon.

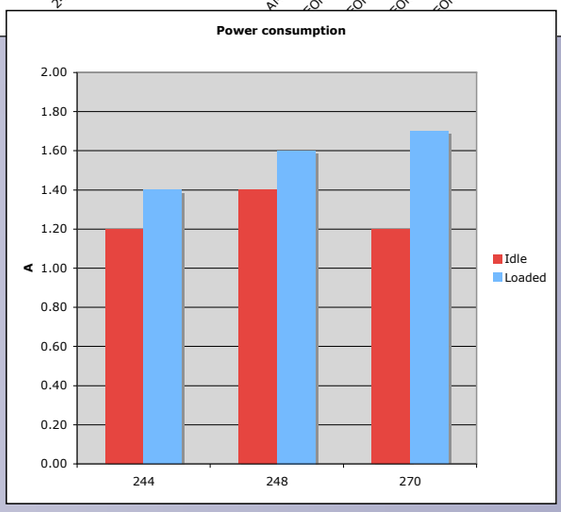

- The ongoing benchmarks show a significant advantage of dual-core 64bit processors for large scale high-performance computing facilities for high energy physics