CPU Benchmarking at GridKa

Manfred Alef

Grid Computing Centre Karlsruhe (GridKa)

Forschungszentrum Karlsruhe
Institute for Scientific Computing
Hermann-von-Helmholtz-Platz 1
D-76344 Eggenstein-Leopoldshafen

http://www.fzk.de, http://www.gridka.de

firstname.lastname@iwr.fzk.de

Oct 10-14 HEPiX 2005

Benchmarks – Why?

- ➔ What's the computing speed of a particular system?
- ➔ What's the total speed my PC cluster?
- ➔ What's the best price per compute speed?
- ➔ Compute speed vs. heat dissipation?
- ➔ Compute speed vs. electric power consumption?

Benchmark Used at GridKa:

SPEC CINT2000 (part of SPEC CPU2000) *

measures integer performance of the computer's processor, memory architecture and compiler.

* SPEC® and the benchmark name SPECint® are registered trademarks of the Standard Performance Evaluation Corporation (www.spec.org).

Published results at www.spec.org – Problem 1:

➔ Influence of system environment:

Example: search "Processor=246"
→ 1226 ... 1323 SPECint2000,
1115 ... 1248 SPECint_base2000!

Does the difference come from
- differently hardware components (mainboard, ...),
- operating system (Windows, Linux, ...; 32/64 bit),
- compiler (GCC, Intel, PathScale, ...), - options,
- optimized libs (MicroQuill Smartheap Library, ...)?
- ...

Types of Benchmarking:

➔ Look for best compiler or best optimization options:
Decide for a particular hardware and OS, and compile the benchmark with different compilers.

Types of Benchmarking:

➔ Look for best compiler or best optimization options:
Decide for a particular hardware and OS, and compile the benchmark with different compilers.

➔ Find out most efficient operating system:
Decide for a particular hardware and compiler, install the operating systems, and run the benchmark on all of the systems.

Types of Benchmarking:

➔ Look for best compiler or best optimization options:
Decide for a particular hardware and OS, and compile the benchmark with different compilers.

➔ Find out most efficient operating system:
Decide for a particular hardware and compiler, install the operating systems, and run the benchmark on all of the systems.

➔ Compare hardware platforms:
Decide for an operating system, a compiler, and run the benchmark on different hardware platforms.
→ Benchmarking at GridKa

Benchmark Usage at GridKa:

- ➔ Benchmark metrics:
  SPECint_base2000
- ➔ Improved sources (src.alt):
  252.eon.stdcpp, 255.vortex.closed_files
  (176.gcc.64bitgcc40)
- ➔ Operating system:
  Scientific Linux
- ➔ Compiler, optimizing options:
  gcc-3.4.3 `-O3 -funroll_loops -march=`*CPU_architecture*

Published results at www.spec.org – Problem 2:

➔ Results refer to only 1 benchmark run per system!

How does the arithmetic performance change, if all CPUs are stressed?

Is a two-processor system real such?

Benchmark Usage at GridKa:

➔ Number of simultaneously running benchmarks:
= number of CPUs/cores

```
#!/bin/sh
#  GridKa benchmark script (simpified)
...
#  2 benchmark runs at the same time
#  (on dual processor box):
( runspec ; runspec ) &
( runspec ; runspec ) &
```

Worker Nodes at GridKa:

Worker Nodes at GridKa:

Intel CPUs:

- 2x Intel Pentium 3, 1.266 GHz, 1 GB (SDRAM133), TYAN Thunder LE-T S2518
- 2x Intel Xeon, 2.2 GHz, 1 GB (DDR200), TYAN Tiger i7500 S2720GN
- 2x Intel Xeon, 2.66 ... 3.06 GHz, 1 GB (DDR266), TYAN Tiger i7501 S2723GN

\+ 1 Disk (IDE, 7200 RPM)

Worker Nodes at GridKa:

AMD CPUs:

- ➔ 2x AMD Opteron 246 (2.0 GHz), 2 GB (DDR333),
  TYAN Thunder K8S PRO S2882, Barebone TYAN Transport GX28 B2882
  (abbrev.: "O-246 T")
- ➔ 2x AMD Opteron 246 (2.0 GHz), 2 GB (DDR400),
  MSI MS-9145, Barebone MSI MS-9245
  (abbrev.: "O-246 M")

\+ 1 Disk (IDE, 7200 RPM)

Worker Nodes at GridKa:

Test nodes:

- ➔ 2x AMD Opteron 270 (2.0 GHz dual core),
  2 GB (DDR400),
  MSI MS-9145, Barebone MSI MS-9245
- ➔ Intel Pentium M 760 (2.0 GHz), 1 GB (DDR2-533),
  AOpen i915GMm-HFS or Commell LV-673

\+ 1 Disk (IDE, 7200 RPM)

Worker Nodes at GridKa:

Vendors:

- ➔ "Next corner's computer manufacturers", 2001-2005
- ➔ MEGWARE, Chemnitz, 2002
- ➔ DALCO, Switzerland, 2005

Benchmark Results:

CPU Speed (SPECint_base2000)

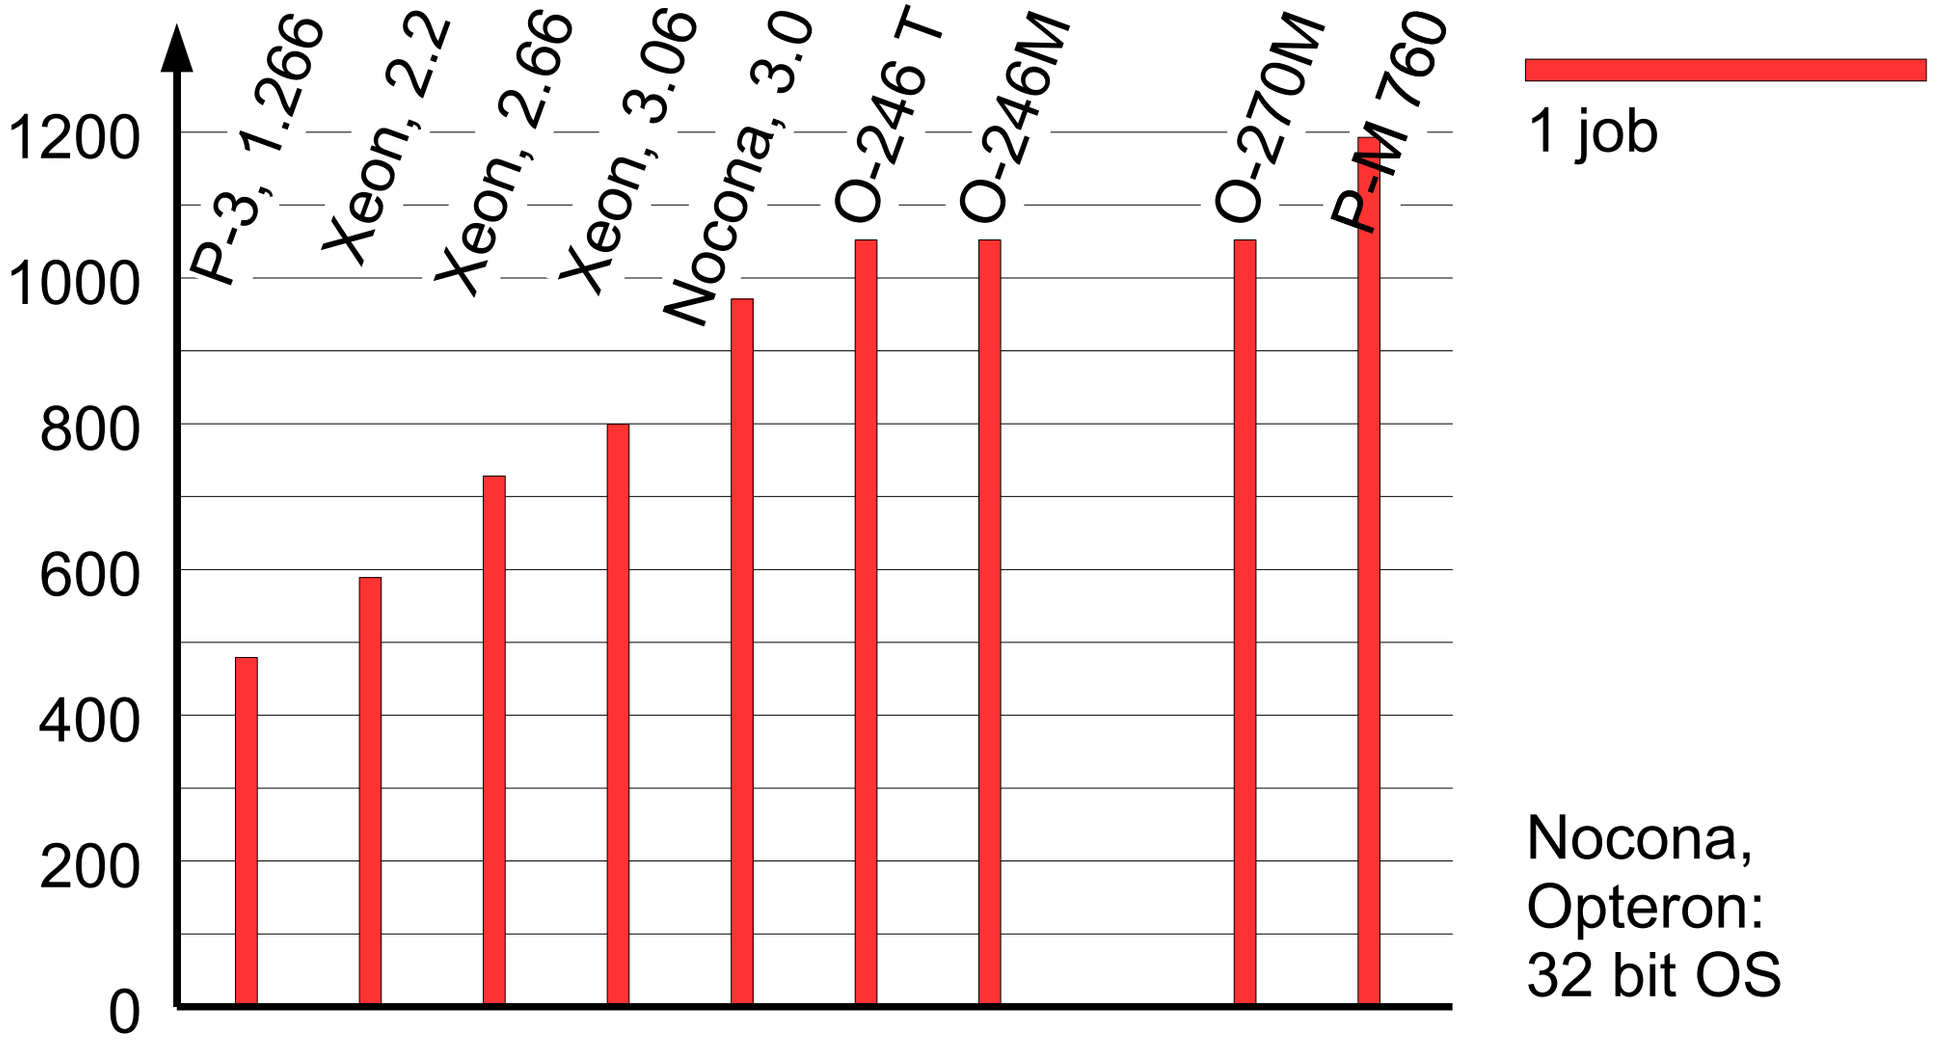

CPU Speed (SPECint_base2000 *per Job*)

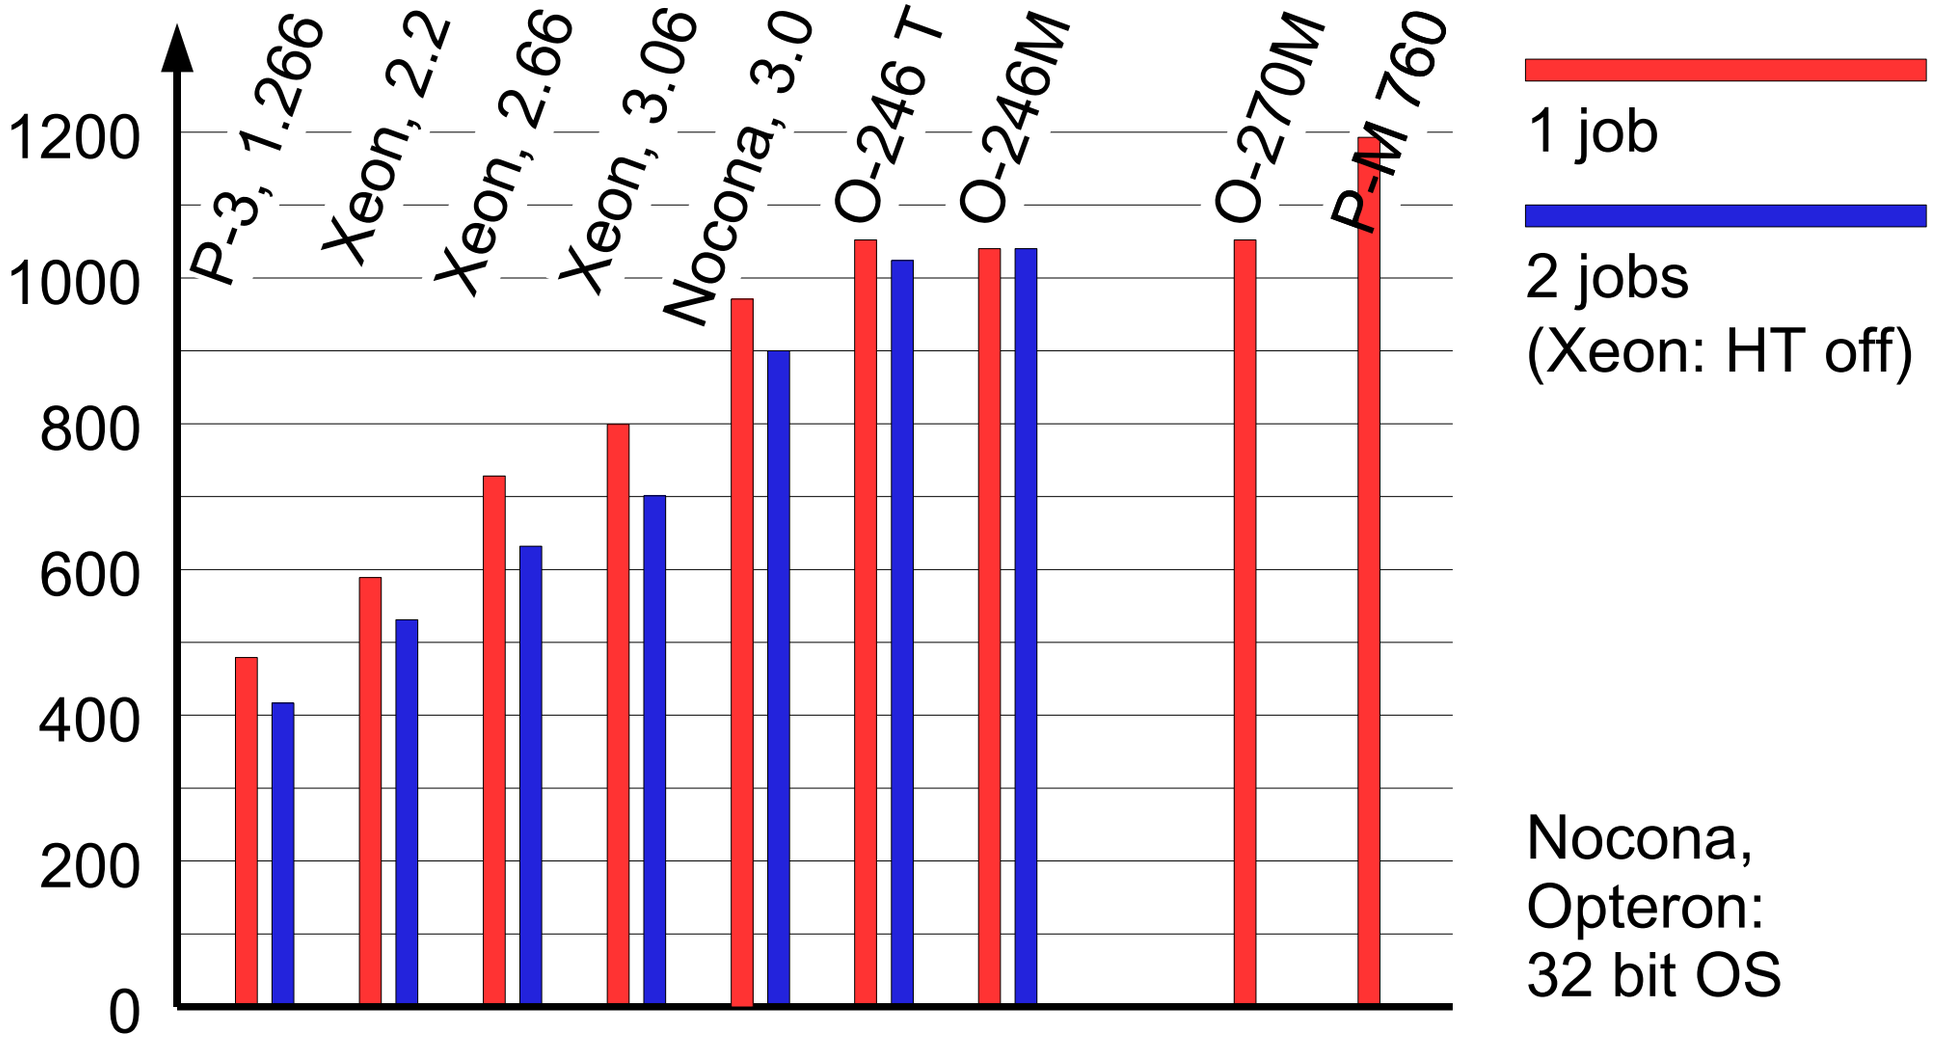

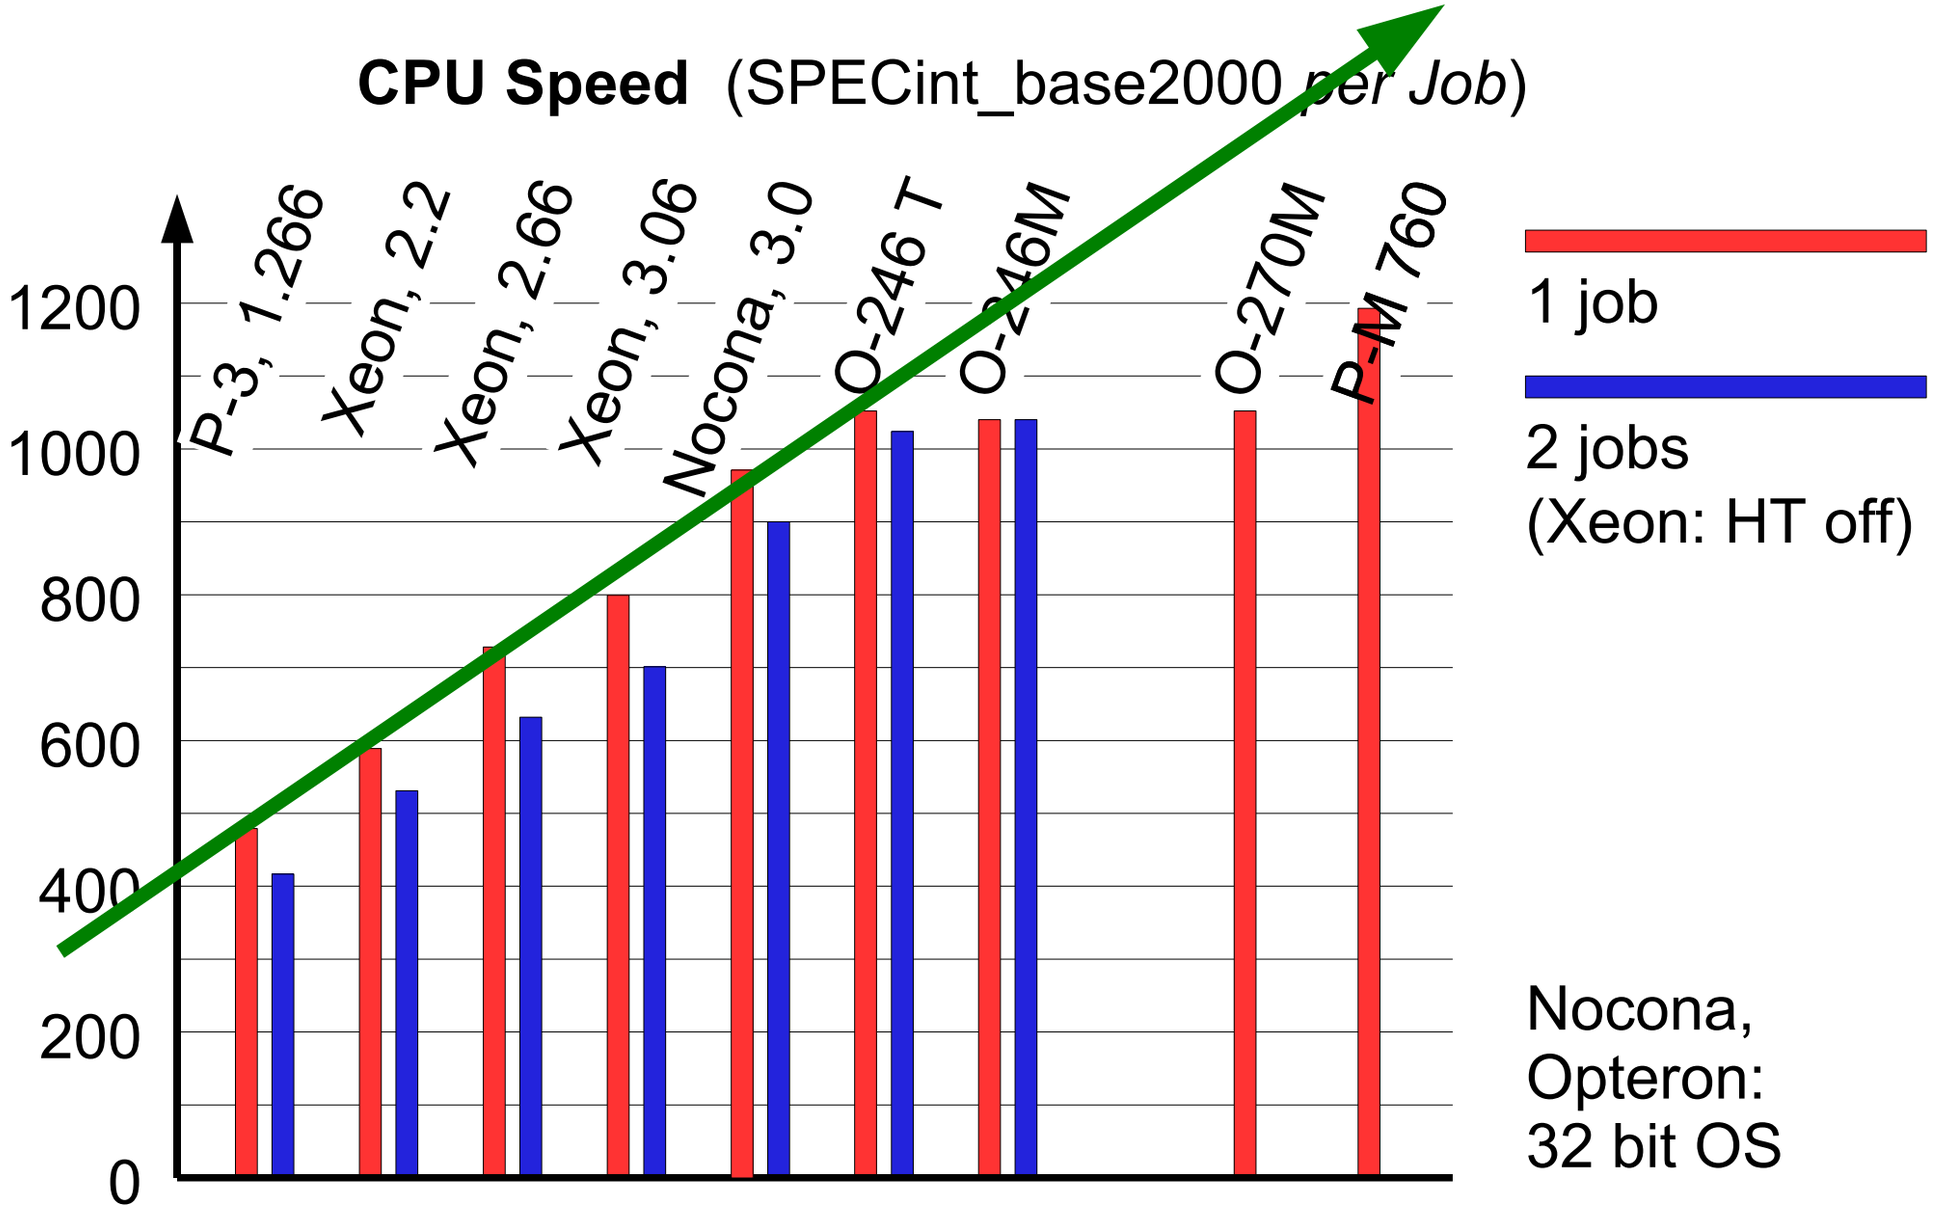
CPU Speed (SPECint_base2000 per Job)
P-3, 1.266
Xeon, 2.2
Xeon, 2.66
Xeon, 3.06
Nocona, 3.0
O-246 T
O-246M
O-270M
P-M 760
1200
1000
800
600
400
200
0
1 job
2 jobs
(Xeon: HT off)
Nocona,
Opteron:
32 bit OS

What about Moore's Law?

- ➔ Moore's Law:
  "Computer speed doubles every 24 months."
- ➔ Benchmark results:
  Computer speed-up less than 2 within the last 24 months!

What about Moore's Law?

of transistors per square inch ...

- Moore's Law:
  "~~Computer speed~~ doubles every 24 months."
- Benchmark results:
  Computer speed-up less than 2 within the last 24 months!
- Principal reasons:
  * HyperThreading,
  * 64 bit technology,
  * bigger caches.

CPU Speed (SPECint_base2000 *per Job*)

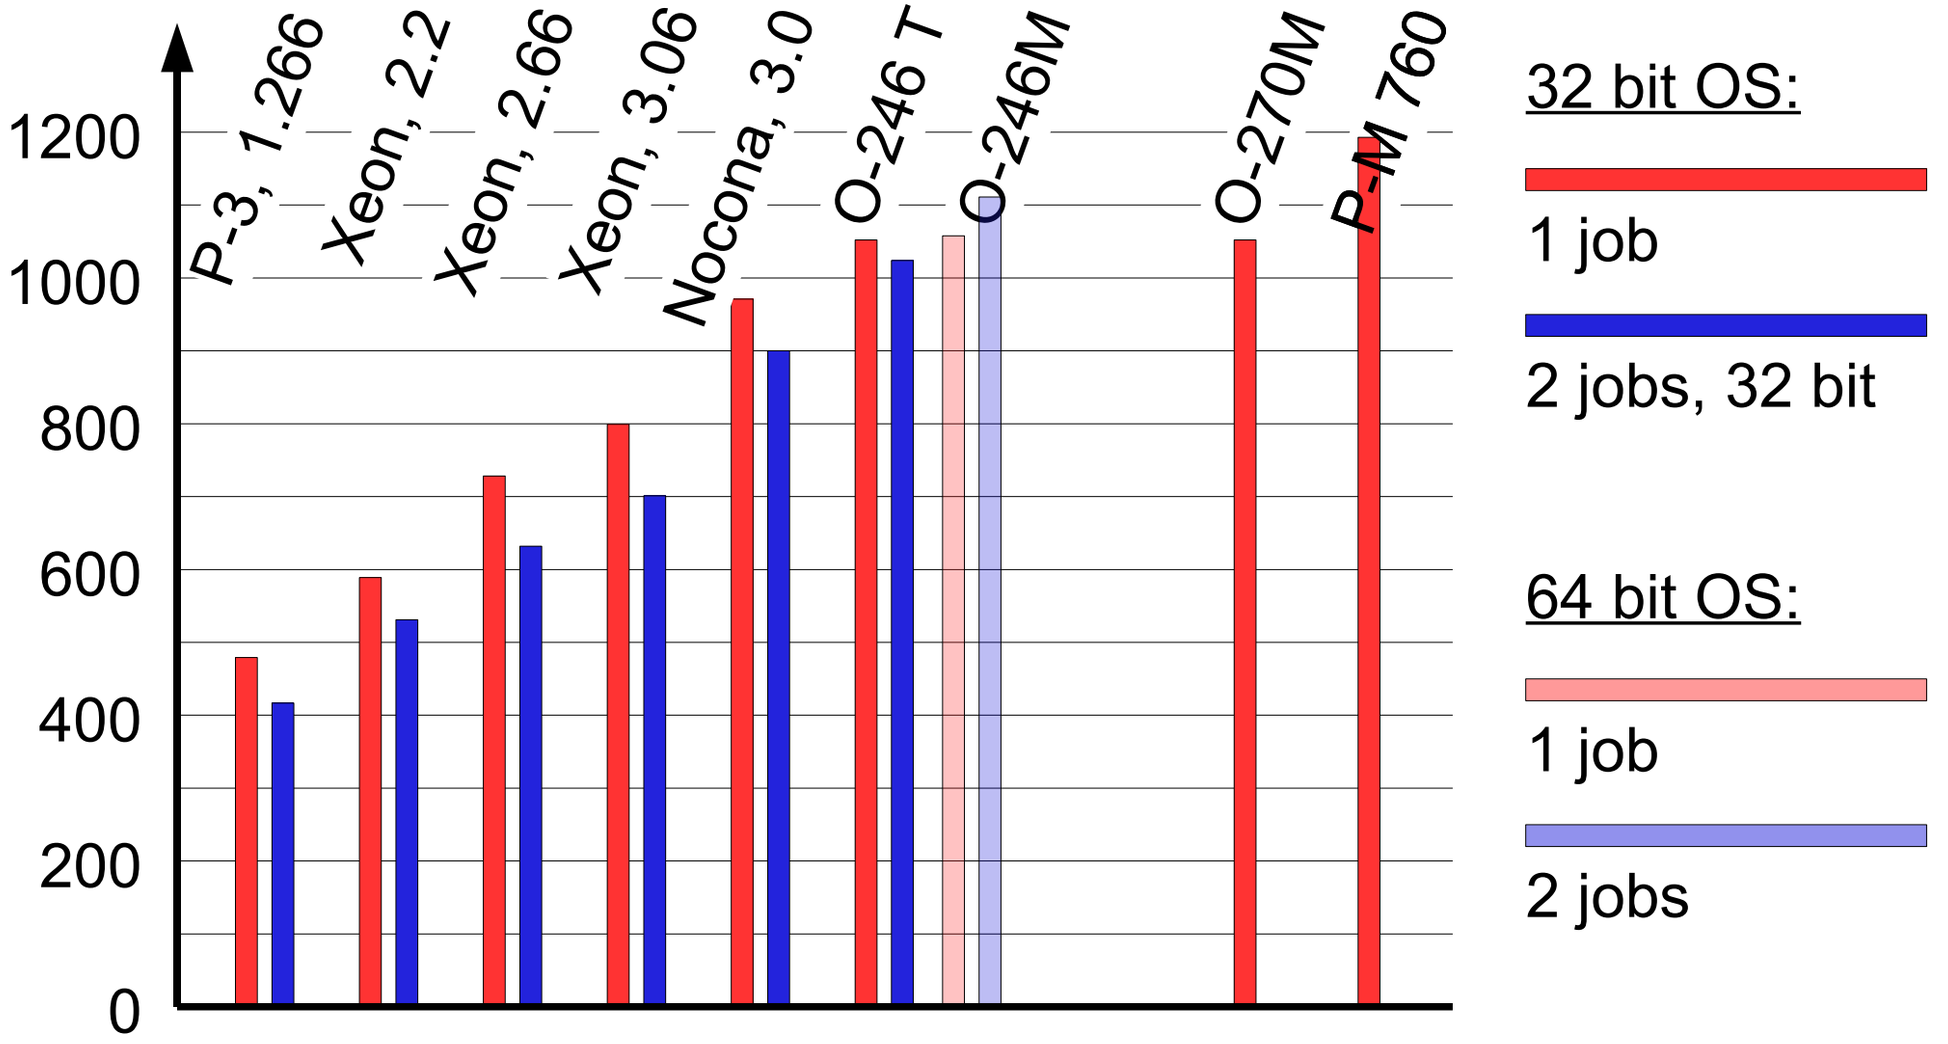

CPU Speed (SPECint_base2000 *per Job*)

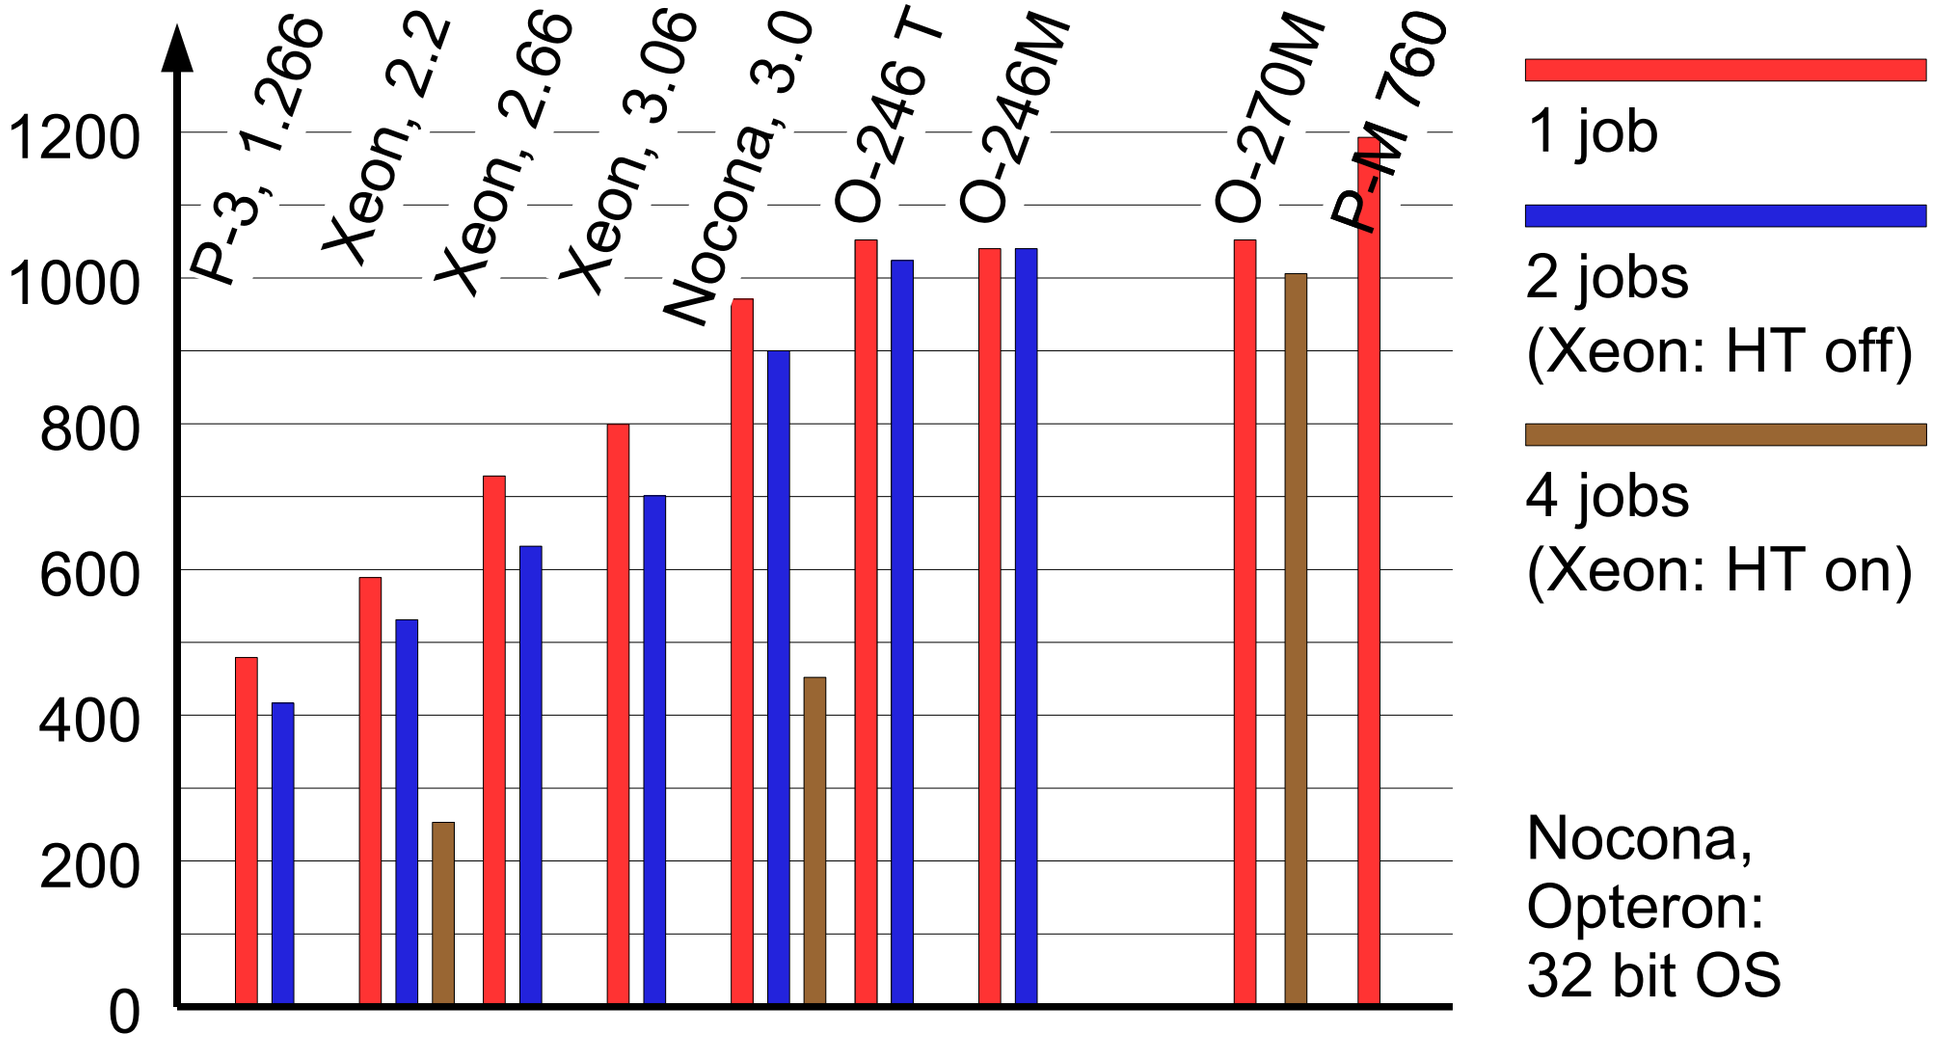

Interpretation:

➔ Multi core processors need sufficient memory throughput:

Memory Throughput (SPECint_rate_base2000)

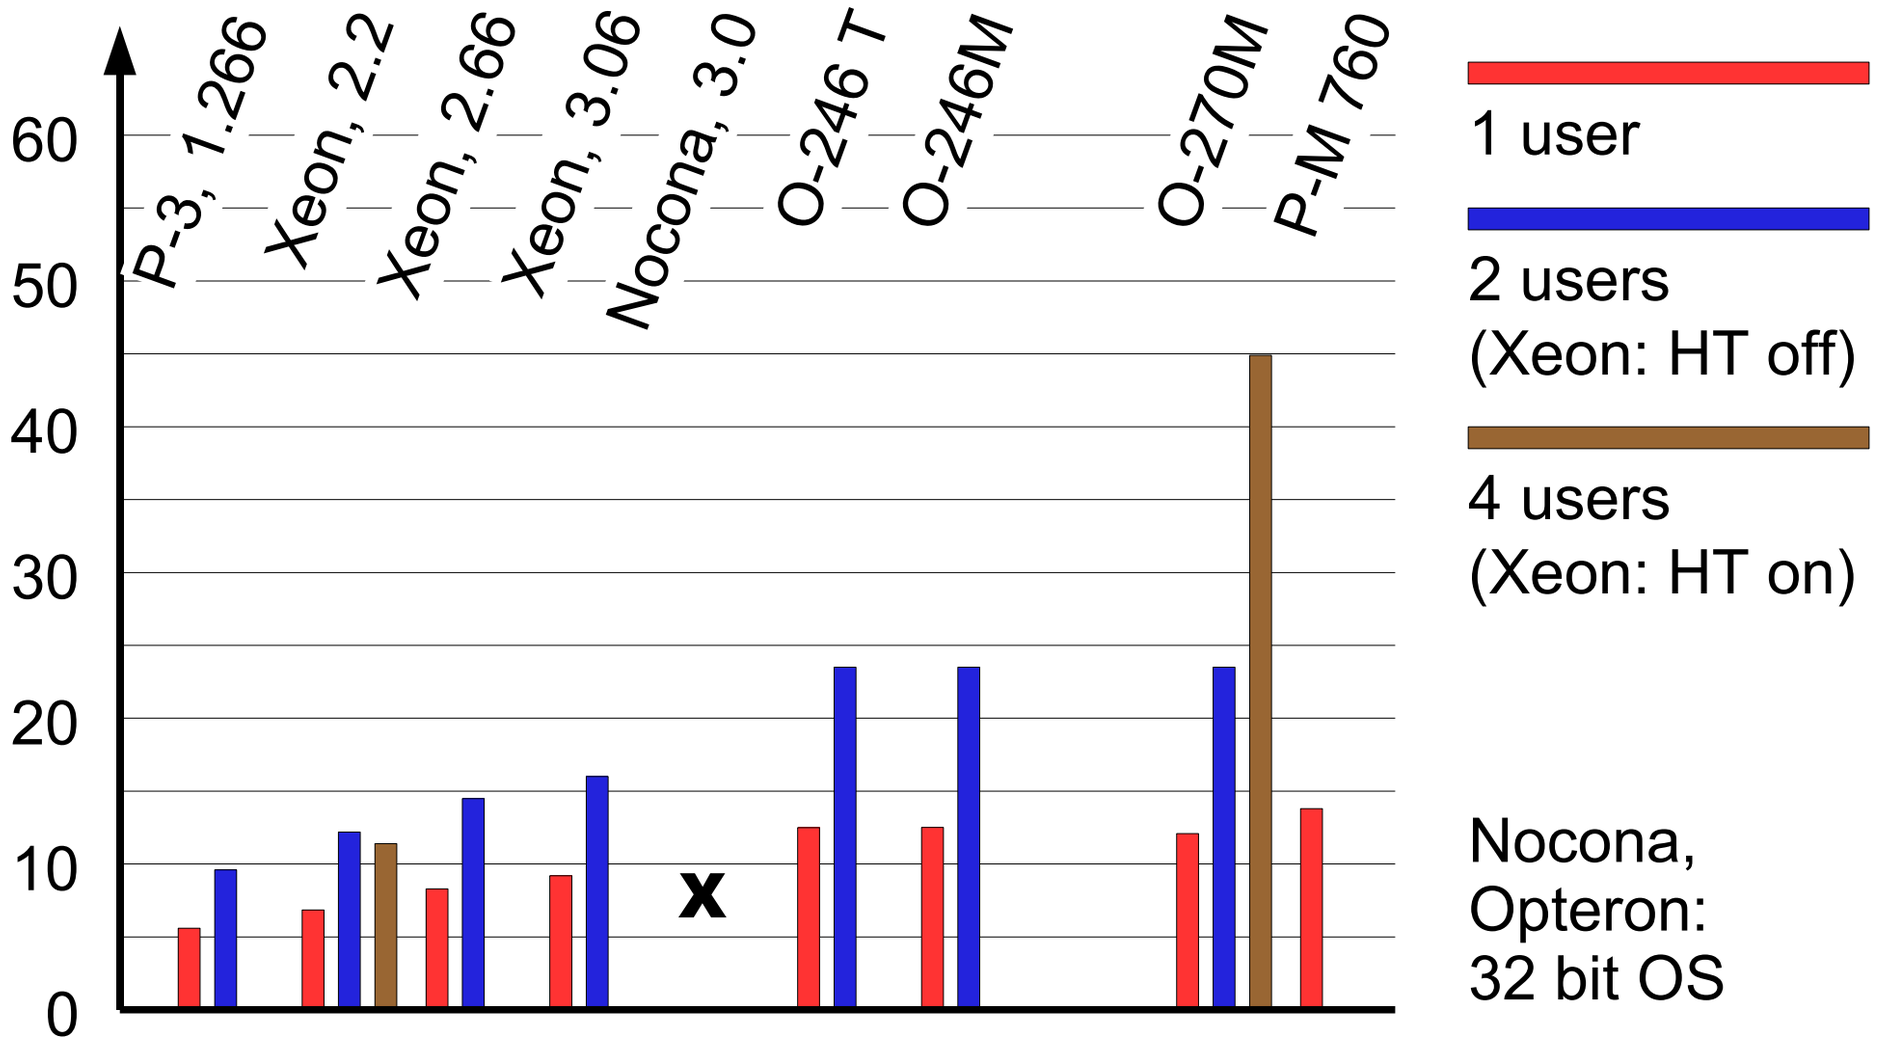

Computer Speed ~ Clock Cycles?

CPU Speed (SPECint_base2000 per Clock Cycle [GHz])

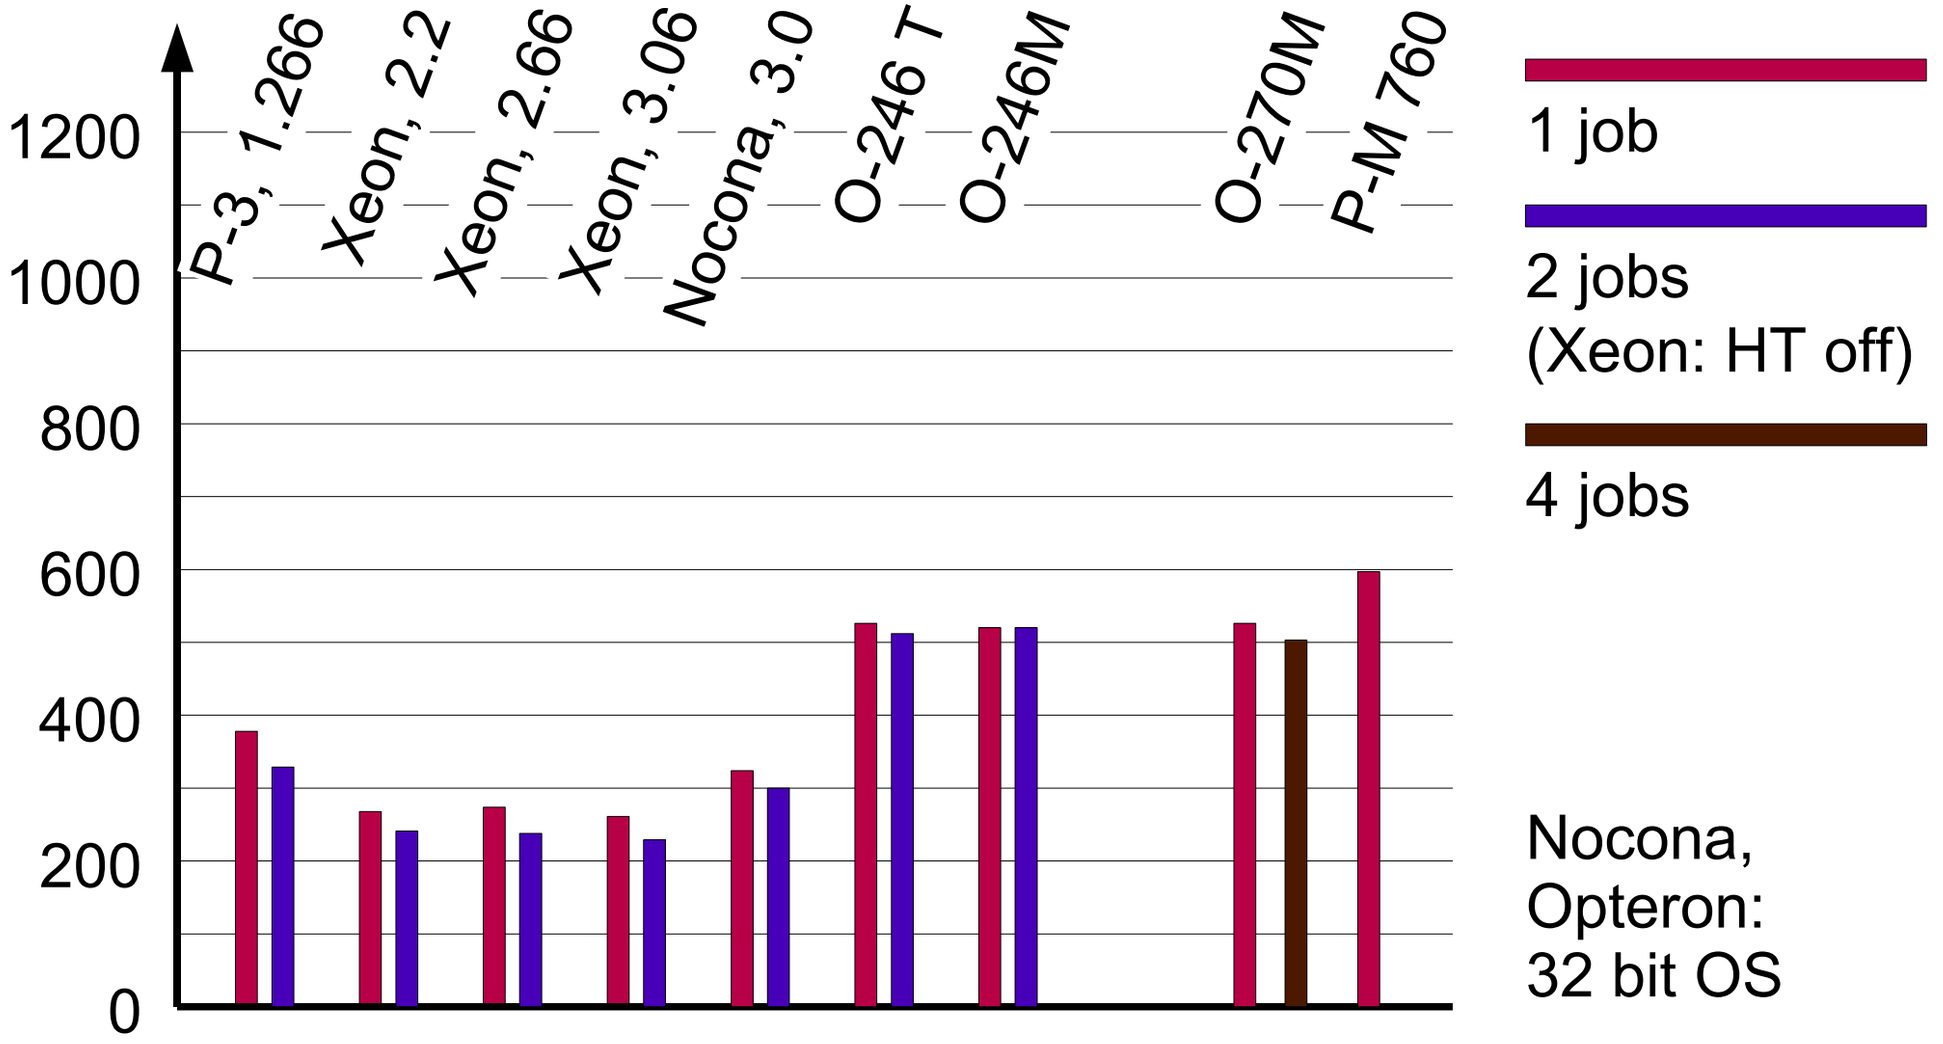

Other Benchmark Results (Contribution of GridKa Users)

CPU (GHz)	Xeon (3.06)	Opteron 246 (2.0)	Pentium-M 760 (2.0)
Benchmark	**Runtime (scaled)**		
CDF – track and and vertex fitting (Kurt Rinnert)			
	1.00		0.76
Compass – Monte Carlo simulation and data reconstruction (Sonja Hedicke)			
Geant	1.00	0.96	0.96
Coral	1.00	0.71	0.76
BaBar (Alexei Volk)			
	1.00	0.65	0.61
D0 (Ay Cano)			
pythia	1.00		0.80
d0gstar	1.00		0.91
d0sim	1.00		0.87
GridKa – SPEC CPU2000, GCC 3.4.3 -O3 -funroll_loops -march=xxx (Manfred Alef)			
SPECint_base2000	1.00	0.66...0.78	0.59...0.67
Electric power consumption of the whole chassis, divided by number of CPUs			
system idle	50 W	104 W	41 W
system busy	140 W	122 W	64 W

Presented at GridKa TAB 2005-08-11

Power Consumption and Heat Dissipation:

Power Consumption and Heat Dissipation:

- ➔ Electric installations may be insufficient.
- ➔ Removing the heat dissipation may be problematic.
- ➔ Cost of electricity within 3 years is about 50% of the initial costs (1 € ≈ 1.20 $ per W and year)!
- ➔ CPUs are the most power consuming components (at least in worker nodes).
- ➔ What about mobile CPUs? Dual core CPUs?

Power Consumption and Heat Dissipation:

➔ One of our first procurements ...

Power Consumption and Heat Dissipation:

➔ One of our first procurements ...

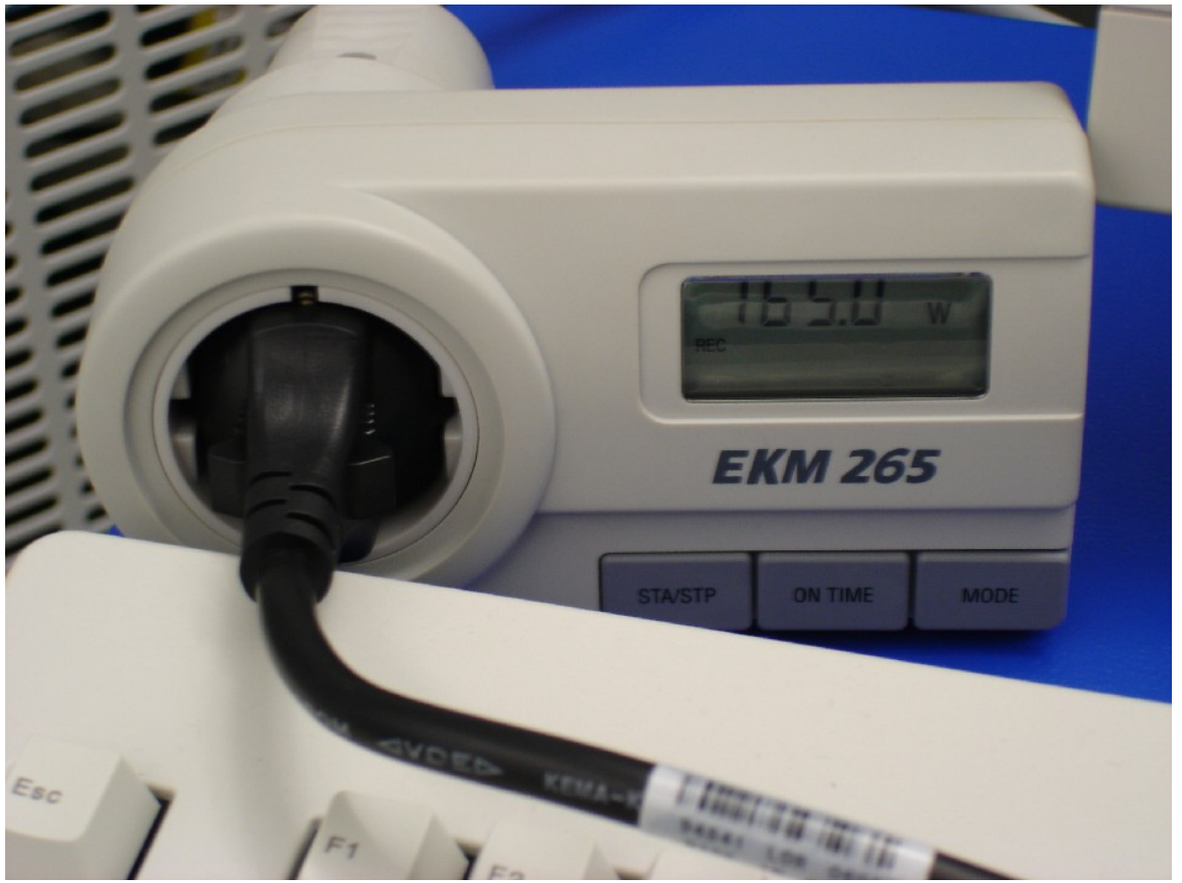

Power Consumption and Heat Dissipation:

➔ What about the accuracy of a cheap wattmeter?

Power Consumption and Heat Dissipation:

- ➔ What about the accuracy of a cheap wattmeter?
- ➔ GridKa acts like an electrical instantaneous water heater:

Power Consumption and Heat Dissipation:

- ➔ What about the accuracy of a cheap wattmeter?
- ➔ GridKa acts like an electrical instantaneous water heater:

Important note:
Okay, there is nevertheless a difference!
You can't run a benchmark on a water heater.
;-)

Power Consumption and Heat Dissipation:

- ➔ What about the accuracy of a cheap wattmeter?
- ➔ GridKa acts like an electrical instantaneous water heater:
- ➔ The energy quantity, which is led out from the area, can be computed from the velocity of flow and the water input and return temperature.

  The result corresponds to the measured power input!

How Does One Measure Correctly?

➔ Electric power consumption (W) shown in the next figure:
- → heating of the machine room,
- → capacity of cooling system / water chillers,
- → electricity bill.

➔ Total power input (measured in VA) can be larger:
- → bigger electric installations required,
- → heating of cables and transformers (outside of the machine room),
- → penalties for electricity.

How Does One Measure Correctly?

➔ Machines were loaded with one instance of cpuburn-1.4 per CPU core.

No additional file accessing software was running (worker nodes).

Electric Power Consumption (Whole System, W per CPU Core)

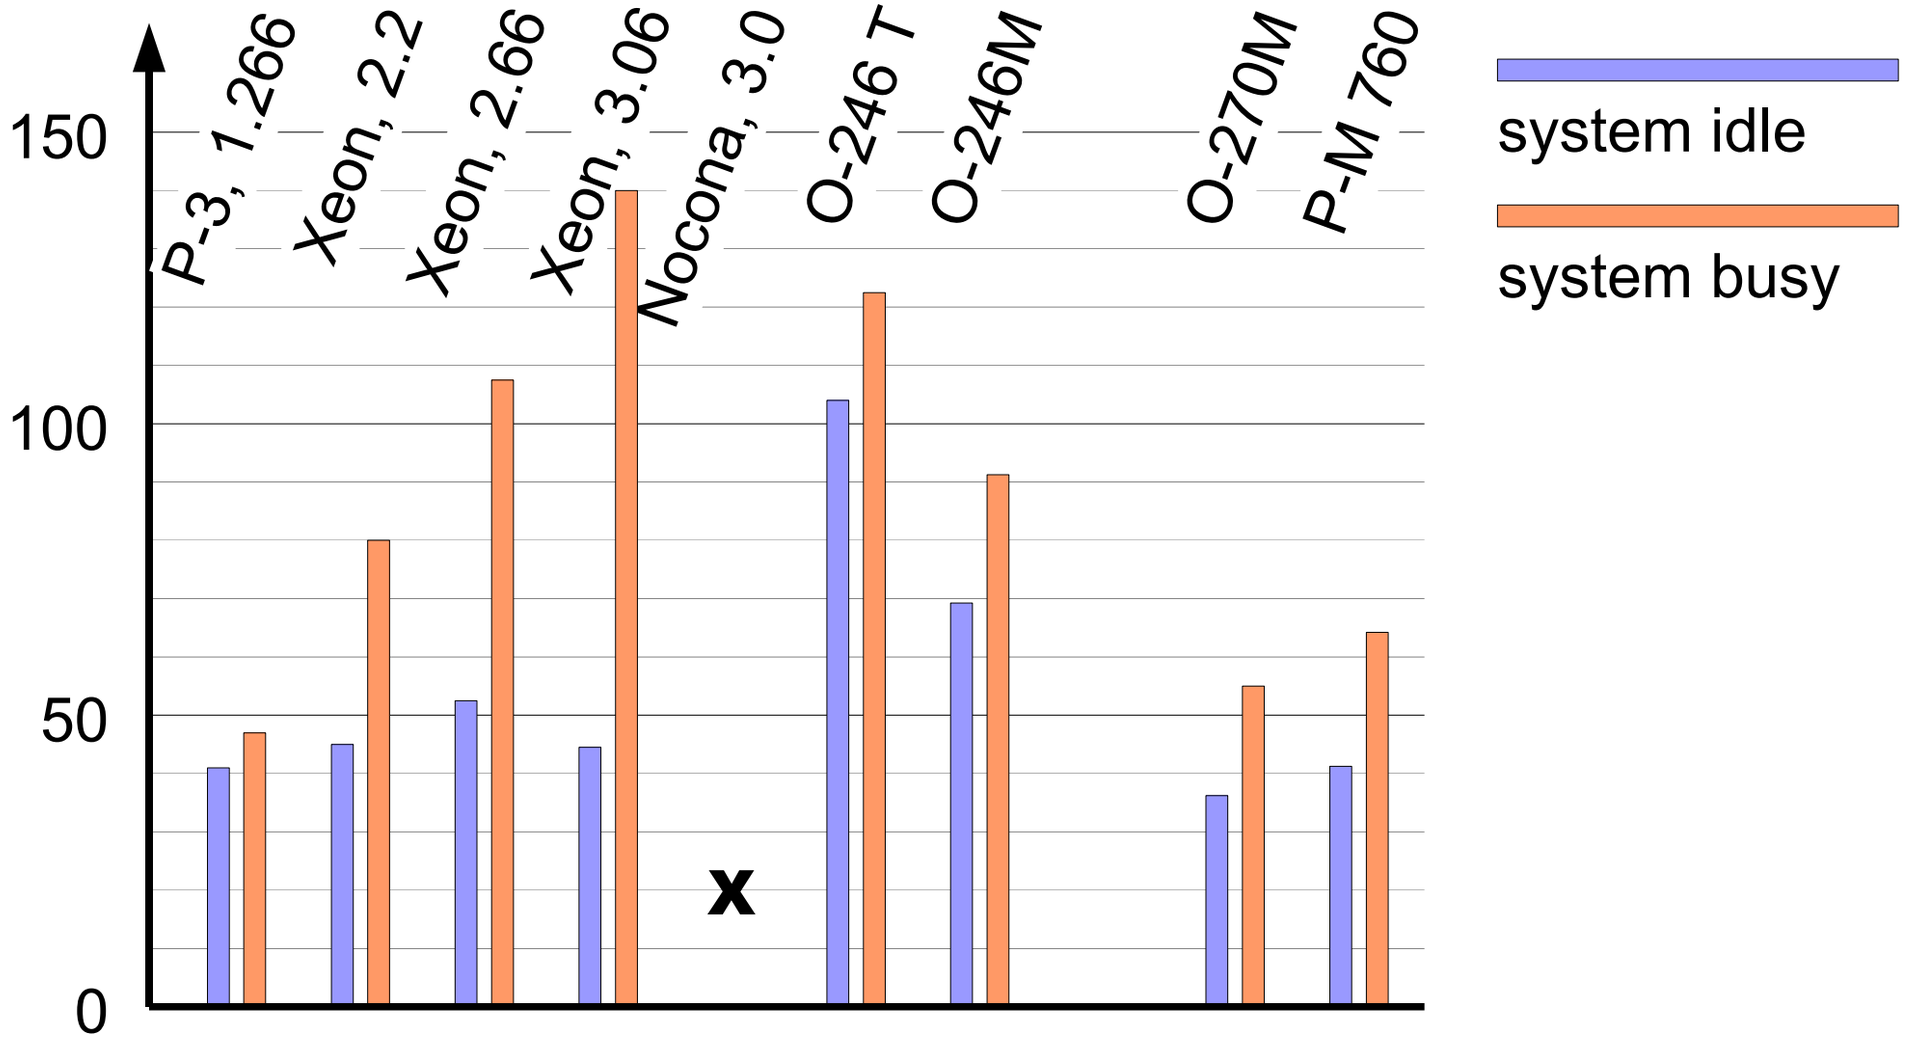

Thermic Efficiency (SPECint_base2000 per W)

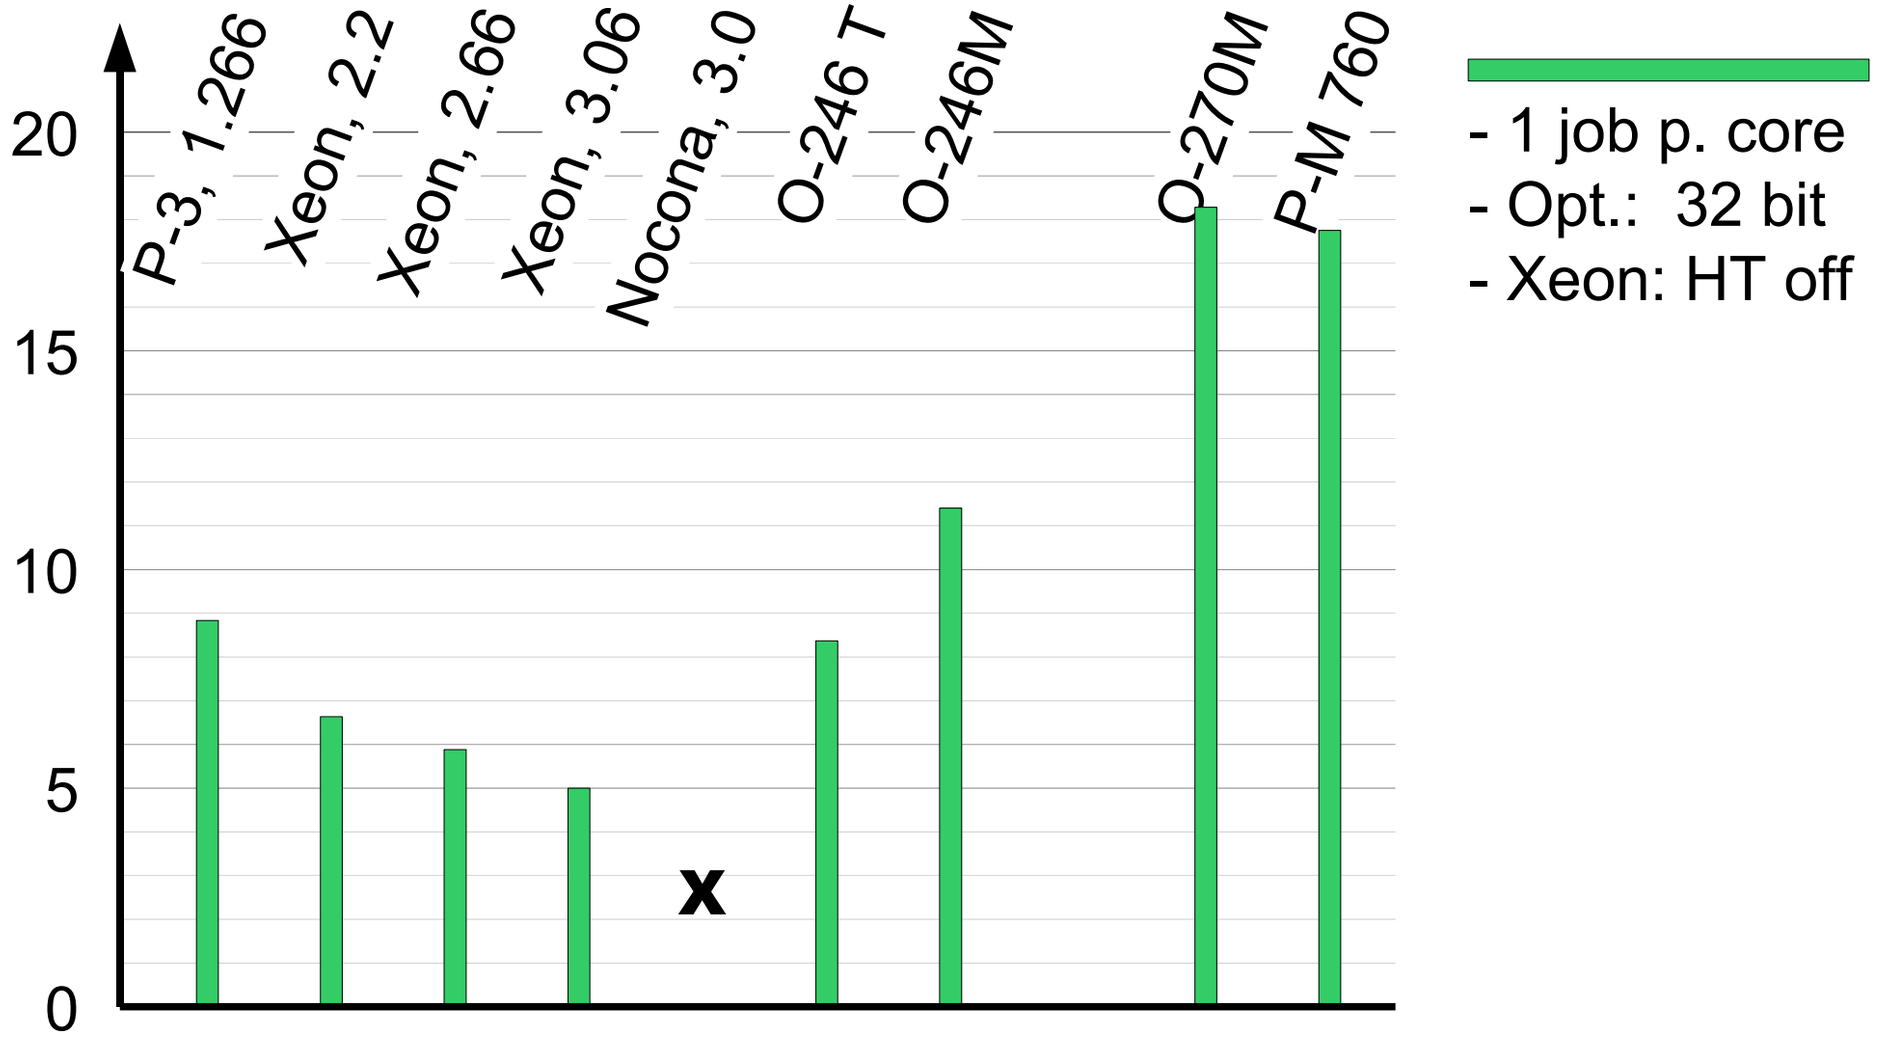

Conclusions:

Conclusions:

- ➔ Run benchmarks in order to compare different hardware architectures.
- ➔ Values published at present, for exemple at spec.org, are not directly comparable.
- ➔ Official HEP benchmark web page?

Conclusions:

- ➔ The arithmetic performance increased in the last two years less strongly than expected (Moore's law) – reasons are: HT, 64 bit, cache size, ...
- ➔ At present the CPU speed raises again strongly (more efficient architectures, dual core, ...).

Conclusions:

➔ Have an eye on the electric power consumption. The compute speed vs. electric power input ratio of today's boxes varies by a factor of more than 3.

Comments, Questions?