

22nd DESY Linux User Meeting

Xeons: benchmarks and user experiences at H1

Bengt Wessling

5.2.2003

Overview

- What's special about Xeons?
- Which CPUs have Hyperthreading?
- How does Hyperthreading look like?
- Comparing with HT on and off
- Conclusions: desktop, workgroup and farm machines

What's special about Xeons?

Note: no SMP-capable non-Xeons for P4-series anymore
→ buy Xeons to get SMP

Introduction of Hyperthreading ("Simultaneous Multithreading"):
1 physical processor → 2 logical processors
so: dual-Xeon looks like 4 CPUs

idea: better usage of existing resources on cpu

- CPU die only 5% larger
- increase in heat dissipation low
- only small additional power consumption
- good marketing: "Adobe Photoshop runs 30% faster"

Which CPUs have Hyperthreading?

all P4s: Xeons and P4.

```
cat /proc/cpuinfo
flags: fpu vme de pse tsc msr pae mce cx8 apic sep mtrr
      pge mca cmov pat pse36 clflush dts acpi mmx fxsr
      sse sse2 ss ht tm
              ^^
```

Activation of HT:

- all Xeons and newer P4s only.
- enable in bios
- requires SMP-kernel
- disable on boot with `noht`

Hyperthreading: how does it look like?

2 logical processors:

- have own registers, interrupt handling
- share all other units: integer, floating point, SSE...
- share caches
- "compete" for physical computing resources
- ideal for multithreaded applications: same program code, same memory

Nearly no HEP code is multithreaded.

Rumors: some code may experience slowdown.

So: better measure it for yourself.

Comparing with HT on and off

Machines: Dual-Xeon 2.4GHz, 1GB RAM

Tasks:

- 1.) generating Monte Carlo ($\approx 100\%$ CPU consumption)
- 2.) data production (95% – 100% CPU)
- 3.) kernel compilation

Idling machines: no differences

Loaded machines (one physical CPU idle): no differences

Overloaded machines (three major jobs on two CPUs, tested only with HT):
nice behavior: +15% CPU for test job with 22% more CPU usage due to extra job

→ for single jobs: no slowdown/acceleration: as hoped/expected

→ not answered: throughput, latency

Conclusions: desktop, workgroup and farm machines

Desktops:

- run with non-SMP kernels (right?)
- HT on modern P4s: need P4-SMP kernels
- interactive work profits from HT

WGS:

- at H1: not meant to run larger jobs anyway
- interactive work profits from HT

Farm nodes:

- HT breaks SMP affinity (HT-aware scheduler patch exists)
- most jobs CPU bound: no resources left for other jobs
- no recommendation to use HT without kernel patch