Stephan Wiesand
February 5, 2003

Some benchmarks from physics on recent ia32 systems

including first hyperthreading experience

Numbers compiled by P. Wegner, DV Zeuthen

Benchmark list Amanda experiment

Architecture (single/dual)	Results	
	Amanda Simulation	Amanda Reconstruction
Host: oceanide1, new desktop PC, 533 MHz FSB Pentium4, 2.40GHz, 512 KB Cache, 256 MB Memory	1625.460u, 10.530s	138.680u, 0.140s
Host: cube3, farm node, 400 MHz FSB (Dell) XEON-P4, 2.4 GHz, 512 KB cache, 2 GB Memory	1626.200u, 1.480s	131.250u 0.320s
Host: minerva, theory PC, 400 MHz FSB XEON-P4, 2.00GHz, 512 KB Cache, 2 GB Memory	1968.150u, 4.620s	167.110u, 0.120s
Host: euterpe, network services Pentium III, 1266MHz, 512 KB Cache, 1 GB Memory	3158.580u, 9.370s	154.830u, 0.740s
Host: pub4, public login host Pentium III (Coppermine), 800 MHz, 256 KB Cache, 1 GB Memory	4778.030u, 31.170s	253.820u, 4.940s
Host: ice53, farm node Pentium III (Coppermine), 800 MHz, 256 KB cache 512 MB	4889.670u, 17.010s	250.290u, 0.310s

Benchmark list hyper threading vs non hyper threading

Architecture	Results			
	Amanda Simulation	Amanda Reconstruction	Theorie form3	
Host: cube3, farm node, 400 MHz FSB XEON-P4, 2.4 GHz, 512 KB cache, 2 GB Memeory Non hyper-threading kernel	1626.200u 1.480s	131.250u 0.320s	Time = 61.33 sec Generated terms = 35999900	
Host: cube8, farm node, 400 MHz FSB XEON-P4, 2.4 GHz, 512 KB cache, 2 GB Memeory Hyper-threading kernel	1617.420u 1.530s	154.210u 0.430s	Time = 62.52 sec Generated terms = 35999900	

Dirac Operator Benchmark

32/64-bit Dirac Kernel, LQCD, using pre-fetch, SSE2 (Martin Lüscher, CERN):

P4, 1.4 GHz, 256 MB Rambus, 256 KB cache

Time per lattice point:

0.926 micro sec (1503 Mflops [32 bit arithmetic])

1.709 micro sec (814 Mflops [64 bit arithmetic])

P4, 2.4 GHz, 256 MB Rambus, 512 KB cache

Time per lattice point:

0.631 micro sec (2207 Mflops [32 bit arithmetic])

1.200 micro sec (1159 Mflops [64 bit arithmetic])

Dirac Operator Benchmark

32/64-bit Dirac Kernel, LQCD, using pre-fetch, SSE2

Tests from Leibnitz Rechenzentrum Munich

Pentium IV 3.06GHz. with ECC Rambus, FSB533

single precision: per lattice point: 0.509 micro sec (2736 Mflops)

double precision: per lattice point: 0.997 micro sec (1396 Mflops)

Pentium IV 2.53GHz. with Rambus 1066 memory, FSB533

single precision: per lattice point: 0.570 micro sec (2443 Mflops)

double precision: per lattice point: 1.078 micro sec (1290 Mflops)

Xeon at 2.4GHz. with PC2100 DDR SDRAM memory, FSB400(?), 4-way machine (!)

single precision: per lattice point: 0.753 micro sec (1849 Mflops)

double precision: per lattice point: 1.495 micro sec (930 Mflops)

Numerical Algebra Benchmarks

Add_assign_field (daxpy), square_norm (Martin Lüscher, CERN):

Architecture	Results: Time per lattice point, MFLOPs			
	add_assign_field		square_norm	
	32-bit	64-bit	32-bit	64-bit
P4, 1.4 GHz, 256 MB Rambus, 256 KB cache	0.133 micro sec 362 Mflops	0.264 micro sec 182 Mflops	0.046 micro sec 1042 Mflops	0.097 micro sec 497 Mflops
P4, 2.4 GHz, 256 MB Rambus, 512 KB cache	0.126 micro sec 382 Mflops	0.250 micro sec 191 Mflops	0.043 micro sec 1115 Mflops	0.086 micro sec 557 Mflops

Conclusions

- recent P4 systems are in average much faster than P3
  - but for some real world tasks P3 Tualatin does surprisingly well
  - when on a budget, the 1.4 Ghz version may be an option
- P4 with fast memory can still gain considerably from higher frequencies
  - even above 2Ghz

- hypedthreading is not a booster, but neither does much harm
  - even gcc compiled code seems to use pipelines well without it
  - moderate slowdown for some applications
  - does not hamper system stability (25 CPU months w/o problems)