Current status of performance measurements

CERN openlab II quarterly review

20 September 2006

Andreas Hirstius

What is done where

Benchmark:

- SPEC2000 version 1.3
  - SPECInt with icc and gcc
  - SPECfp with icc
    - some benchmarks are written in FORTRAN90 – unsupported by gcc
  - SPECbase measurements
    - 1, n/2, n, 1.5*n parallel jobs, manually started
    - Benchmarks run independently – out-of-sync
  - SPECrate measurements
    - 1, n/2, n, 1.5*n parallel jobs, controlled by script
    - Benchmarks are running in-sync (n=# of cores)
- SPEC2006 version 1.0 has just arrived
  - runtime is ~6 - 8 times longer

Machines:

- 8-way Itanium2 1.6GHz (HP rx7620) ➞ 8 cores
- 4-way Montecito 1.4GHz (Tiger4 based) ➞ 8 cores
- Dual Dempsey 3.2GHz (64-bit) ➞ 4 cores
- Dual Woodcrest 2.66GHz (64-bit) ➞ 4 cores
- … and others

First remarks

- SPEC2000 ran only on Itanium out-of-the-box with the config files provided by SPEC
- Significant effort necessary to it get to work on EM64T and even on IA32 machines !?!?
- Some benchmarks do not run properly in 64-bit mode on EM64T!

- gcc was tested with “-O2” because this is the max. optimisation the physics programs use
  - 252.eon from SPECint has problems with “-O2” (EM64T/IA32)
  - “-O0” was also tested because some code might still be compiled with it

- icc was tested with two optimisation levels
  - “-fast” == “-xP -O3 -ipo -no-prec-div -static”
  - “-fast” and profile guided optimisation

And now many numbers ----->

SPEC results – SPECint Itanium

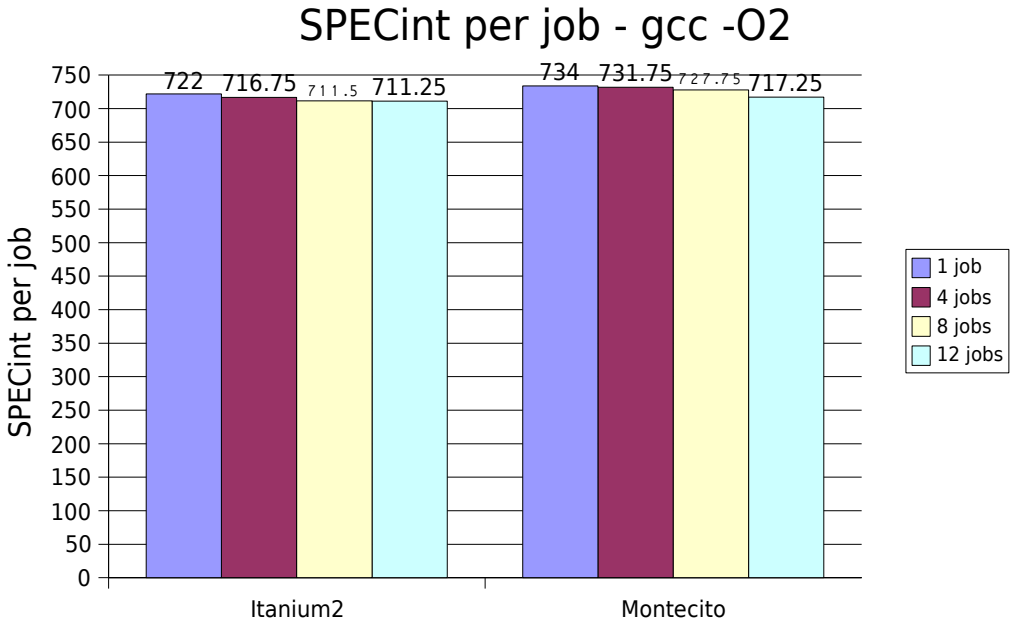

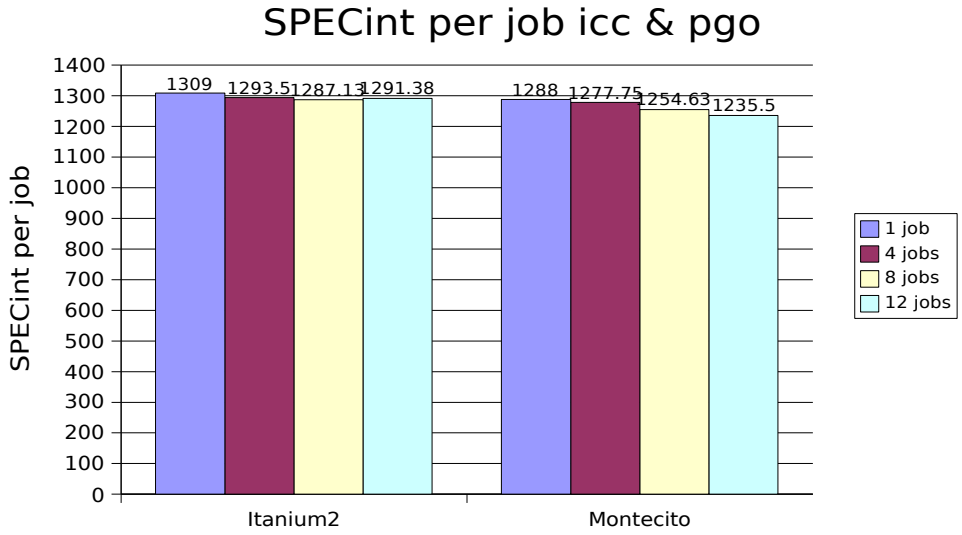

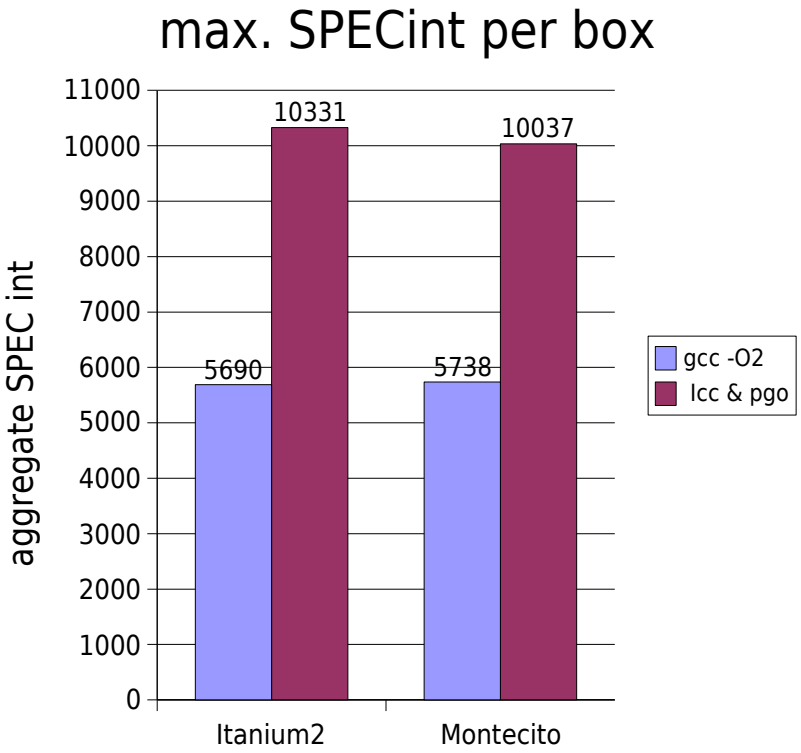

SPEC results – SPECint Xeon – gcc

Woodcrest delivers with gcc and CERN settings (*-O2 -fPIC -pthread*):

- ~1700 SPECint for a single job (2.666GHz)
- ~6000 SPECint for a dual-socket system (4 cores, 6 jobs, 2.666GHz)
- Per GHz:
  - 108% faster than Dempsey, 155% faster than Nocona (1 job)
  - 83% faster than Dempsey, 346% faster than Nocona (max. SPECint)
- 20% faster in 64bit mode compared to 32bit mode!

SPEC2000 - gcc CERN settings - 1 job

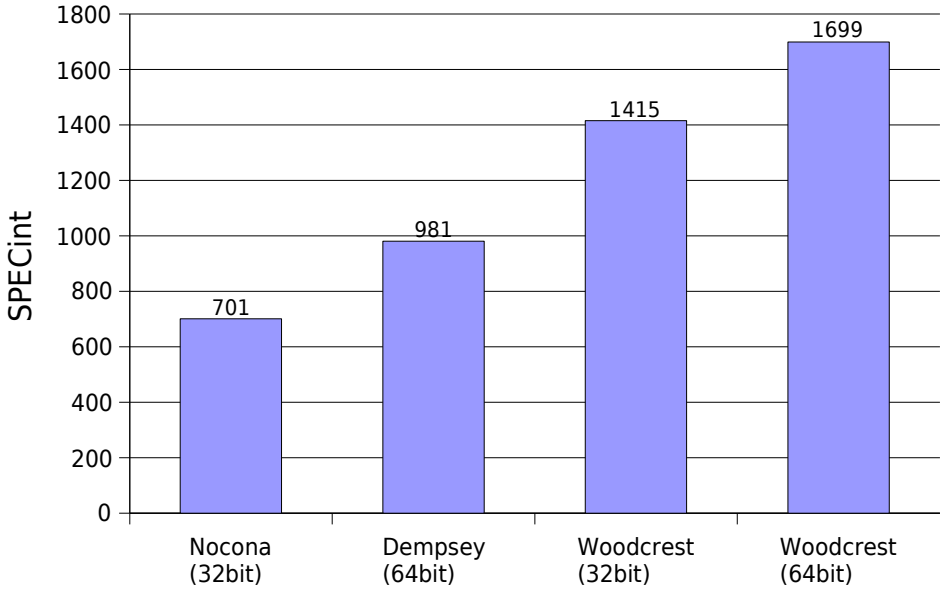

SPEC2000 - gcc CERN settings - 1 job

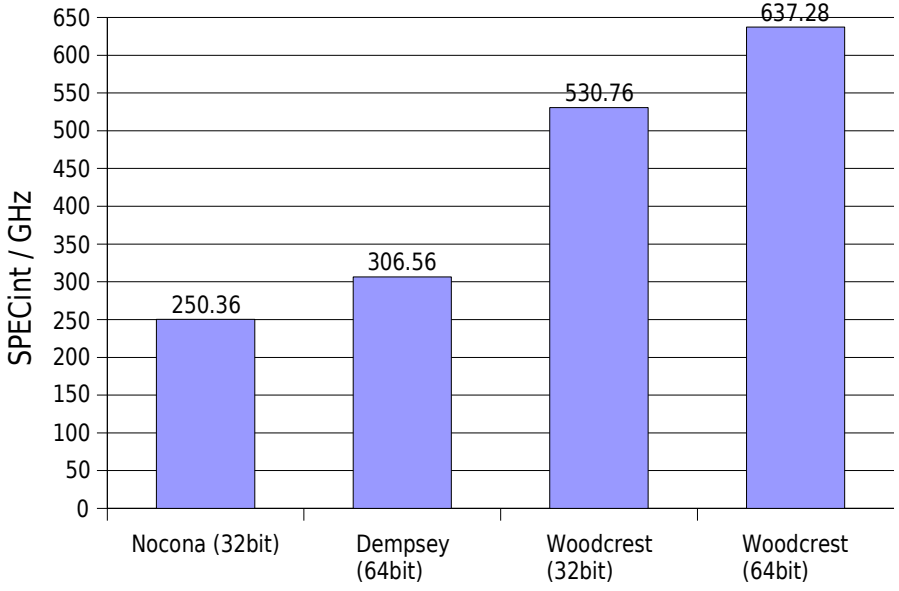

SPEC results – SPECint Xeon – icc

Woodcrest delivers with icc and max. optimization:

- ~2600 SPECint for a single job (2.666GHz)
- ~9300 SPECint for a dual-socket system (4 cores, 6 jobs, 2.666GHz)
- Per GHz
  - 113% faster than Dempsey, 138% faster than Nocona (1 job)
  - 82% faster than Dempsey, 319% faster than Nocona (max. SPECint)

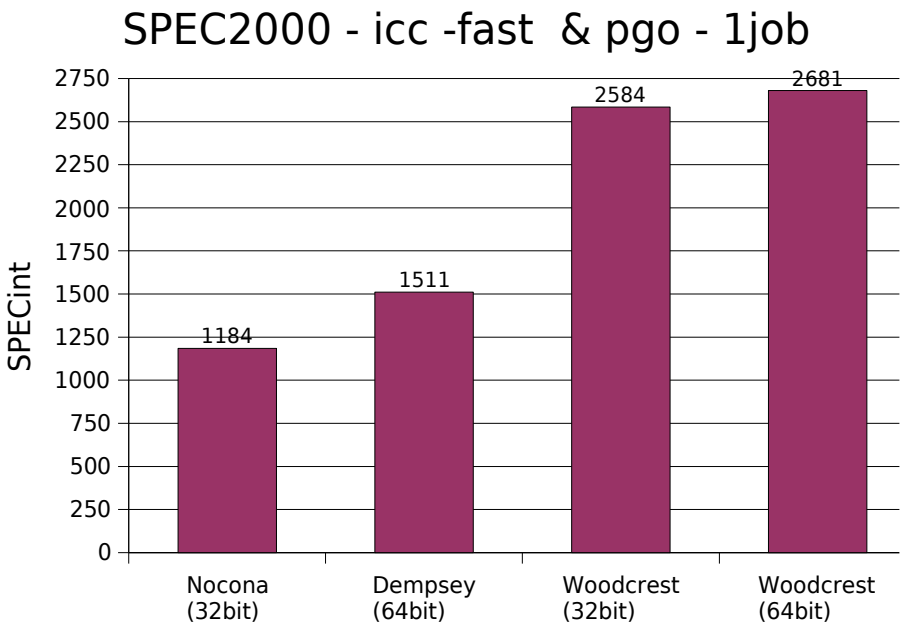

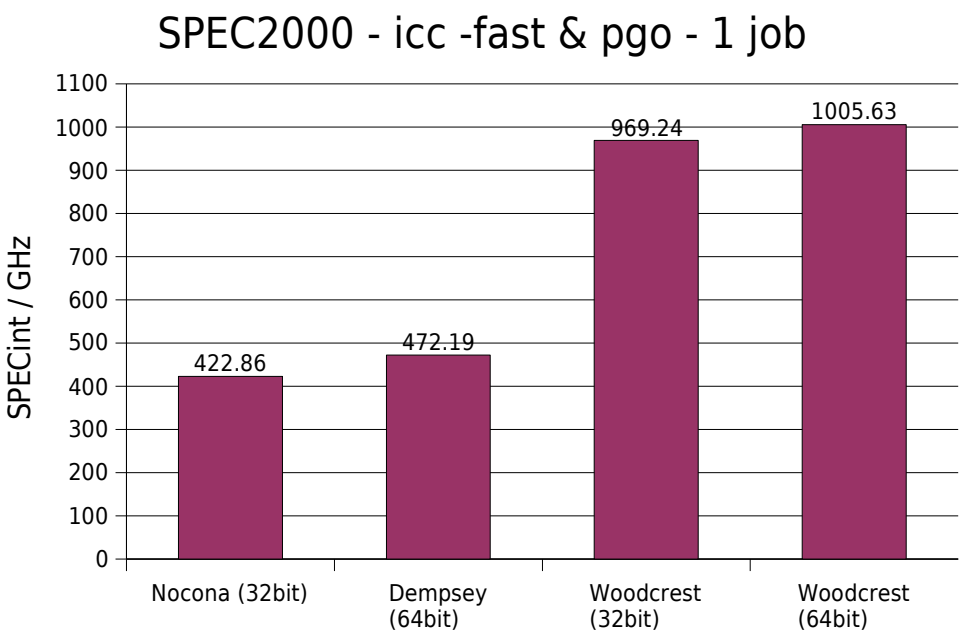

SPEC results – SPECint Xeon – icc vs. gcc

- gcc CERN → icc & pgo

➔ >60% gain

➔ Significant performance improvement without much effort
(at least in theory)

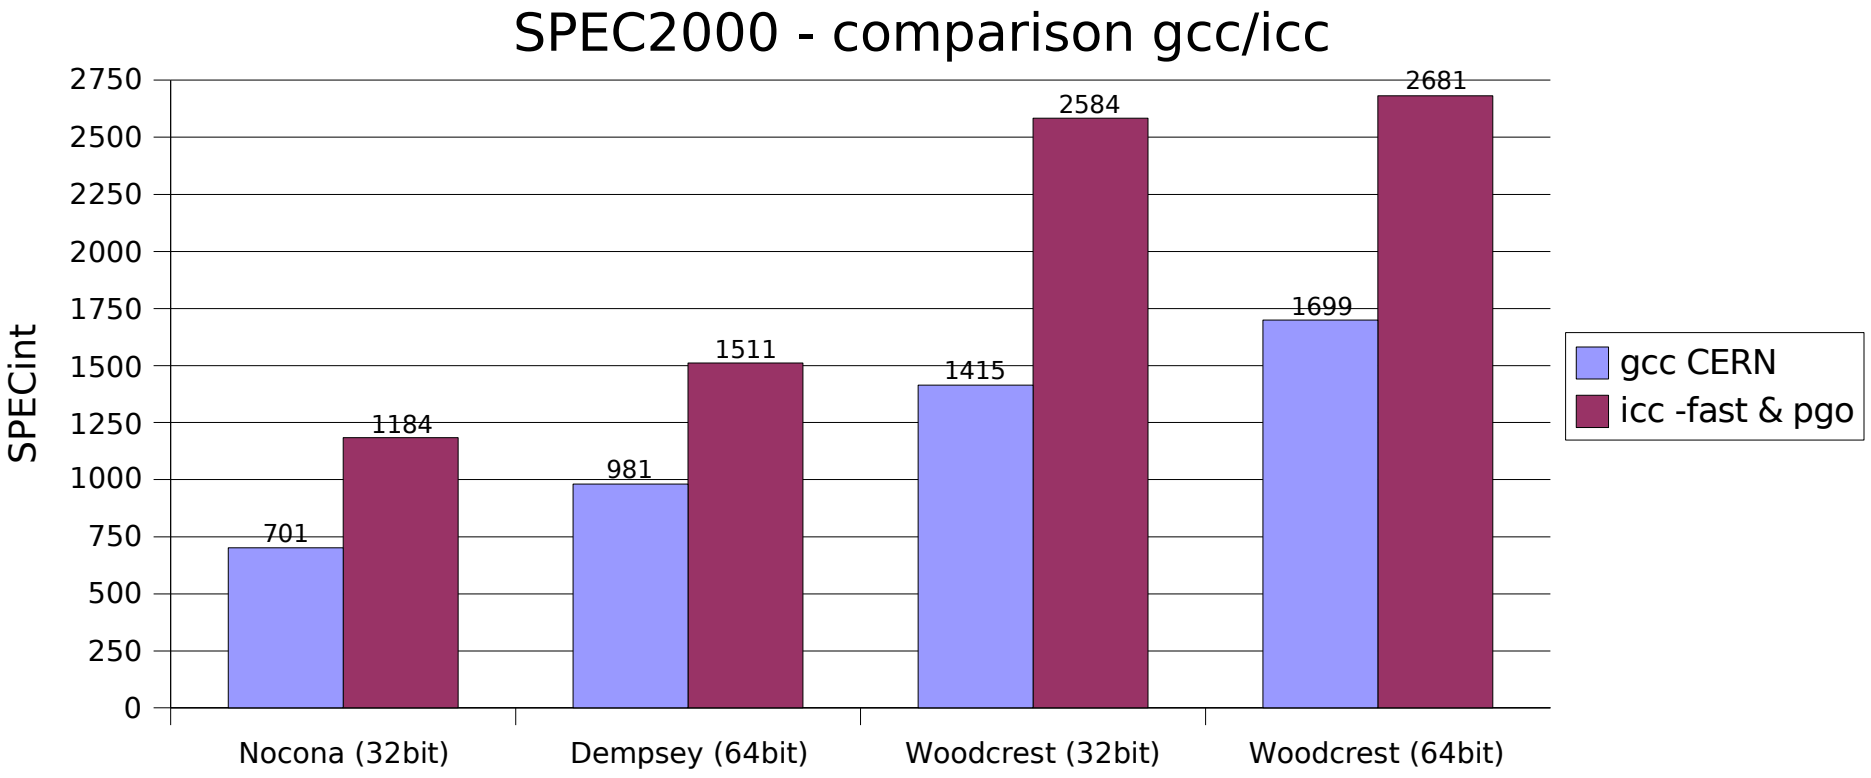

SPEC general conclusions

SPECbase vs. SPECrate

- Multiple SPECbase jobs show better scaling than SPECrate
- With oversubscribed CPUs SPECbase usually still sees an increase in aggregate performance
- SPECrate is significantly worse with 1.5*n jobs than with n jobs

Most likely explanation:

- when running in-sync (SPECrate) all jobs hit the same "problematic" code at the same time leading to long idle time of the CPUs
- when running out-of-sync (multiple SPECbase) the "problematic" code of one job is covered up by "normal" code from another job which finally results in higher aggregate performance

➡ For our purposes multiple SPECint are more appropriate

A word on power consumption

Measurement of the power consumption of

- The **Dempsey** system
  - Cores based on ancient Netburst architecture
  - 4 GB FB-DIMMs (DDR2 based) (~10W per DIMM)
  - 3 disks (~11W idle, ~16W active)

- The **Woodcrest** system

➔ Cores based on Core 2 architecture
  - 4 GB - FB-DIMMs (DDR2 based) (~10W per DIMM)
  - 3 disks (~11W idle, ~16W active)

	Woodcrest	**Dempsey**
Max. utilisation	~290W	~410W
SPECint/Watt icc & pgo	31.9	14.9
SPECint/Watt gcc -O2	20.5	9.8

Power consumption and Hyperthreading

A summerstudent was looking at power consumption and performance depending on Hyperthreading (HT)

- 2 jobs on a Dual 3.6GHz Irwindale (4GB RAM)
  - with HT on performance was very low ➞ the scheduler scheduled both jobs on the same physical CPU most of the time
  - with a “trick” the scheduler moved the jobs to different physical CPUs

➔ with HT the performance drops by ~10%, but power consumption drops by ~20%

➔ gain of ~10% in performance/Watt ⇨ but exactly why ?!?!?!?

EON sub-bmk; 2 jobs; SPECint