Moving out of SI2K

How INFN is moving out of SI2K as a benchmark for Worker Nodes performance evaluation

Michele Michelotto at pd.infn.it

SI2K frozen

- SI2K is the benchmark used up to now to measure the computing power of all the HEP experiments
  - Computing power requested by experiment
  - Computing power offered by a Tier-0/Tier1/Tier2
- SI2K is actually the SPEC CPU Int 2000 benchmark
  - Came after Spec89, Spec Int 92 and Spec Int 95
  - Declared obsolete by SPEC in 2006
  - The new benchmark from SPEC is CPU Int 2006

Problems with SI2K

- Impossible to find SPEC Int 2000 results for the new processors (e.g. the not so new Clovertwon 4core)
- Impossible to find SPEC Int 2006 for old processor (before 2006)
- It doesn’t match the HEP program behaviour like it did in the past

Nominal SI vs real SI

- SI2K results for the last generation processor afftected by inflation
- So CERN (and FZK) started to use a new currency: SI2K measured with “gcc”, the gnu C compiler and using two flavour of optimization
  - High tuning: gcc –O3 –funroll-loops–march=$ARCH
  - Low tuning: gcc –O2 –fPIC –pthread
- CERN Proposal: Use as site rating the “Real SI” increased by 50%
  - Actually this make sense only for a short period of time and for the last generation of processor

Nominal SI vs real SI

- CERN Proposal: Use as site rating the “Real SI” obatained by nominal SI increased by 50%
  - Actually this make sense only for a short period of time and for the last generation of processor
- Run n copies in parallel
  - Where n is the number of cores in the worker node

Too many SI2K

- Take as an example a worker node with two Intel Woodcrest dual core 5160 at 3.06 GHz
- SI2K nominal: 2929 – 3089 (min – max)
- SI2K on 4 cores: 11716 - 12536
- SI2K gcc-low: 5523
- SI2K gcc-high: 7034
- SI2K gcc-low + 50%: 8284

Even more

- Actually all the gcc results in the previous slide are on i386 (32bit)
- if you would like to know how your code is running on 64 bit machine, you can measure Specint INT 2000 with gcc on x86_64.
- So the worker node with two Intel Woodcrest dual core 5160 at 3.06 GHz
- SI2K nominal: 2929 – 3089 (min – max)
- SI2K on 4 cores: 11716 - 12536
- SI2K gcc-low: 6021
- SI2K gcc-high: 6409
- SI2K gcc-low + 50%: 9031

A scale factor

- All these numbers would be only annoying in a world with a unique architecture, in which only clock improves in time
- You would be able to find a fixed ratio between all those number.
- But the ratio depends on CPU producer (intel vs AMD) and processor generation (old xeon vs new “core” Xeon

The nominal SI2K

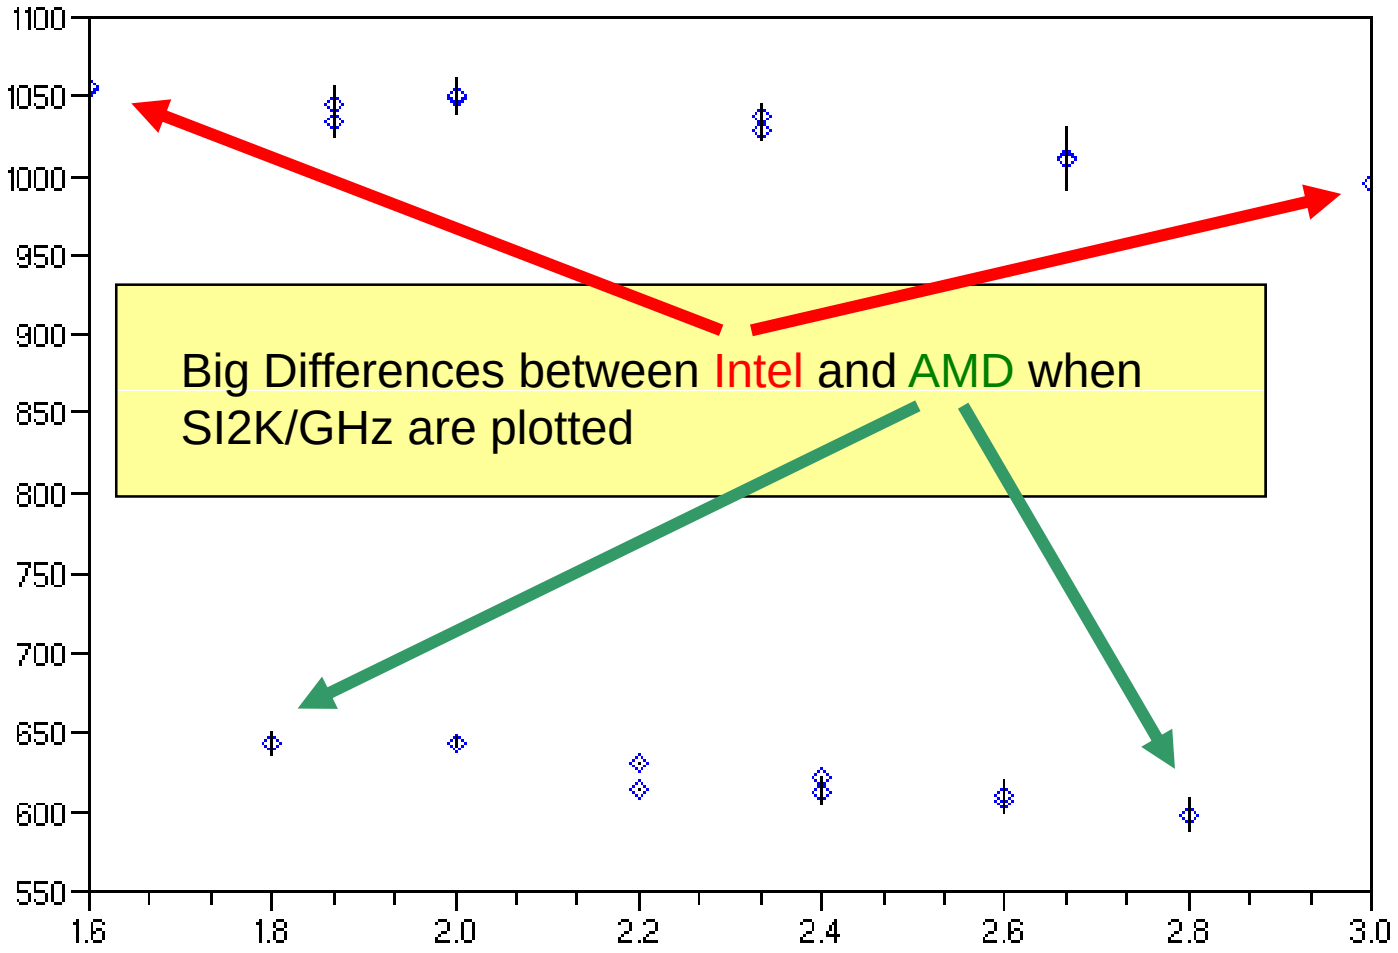

The nominal SI2Krate

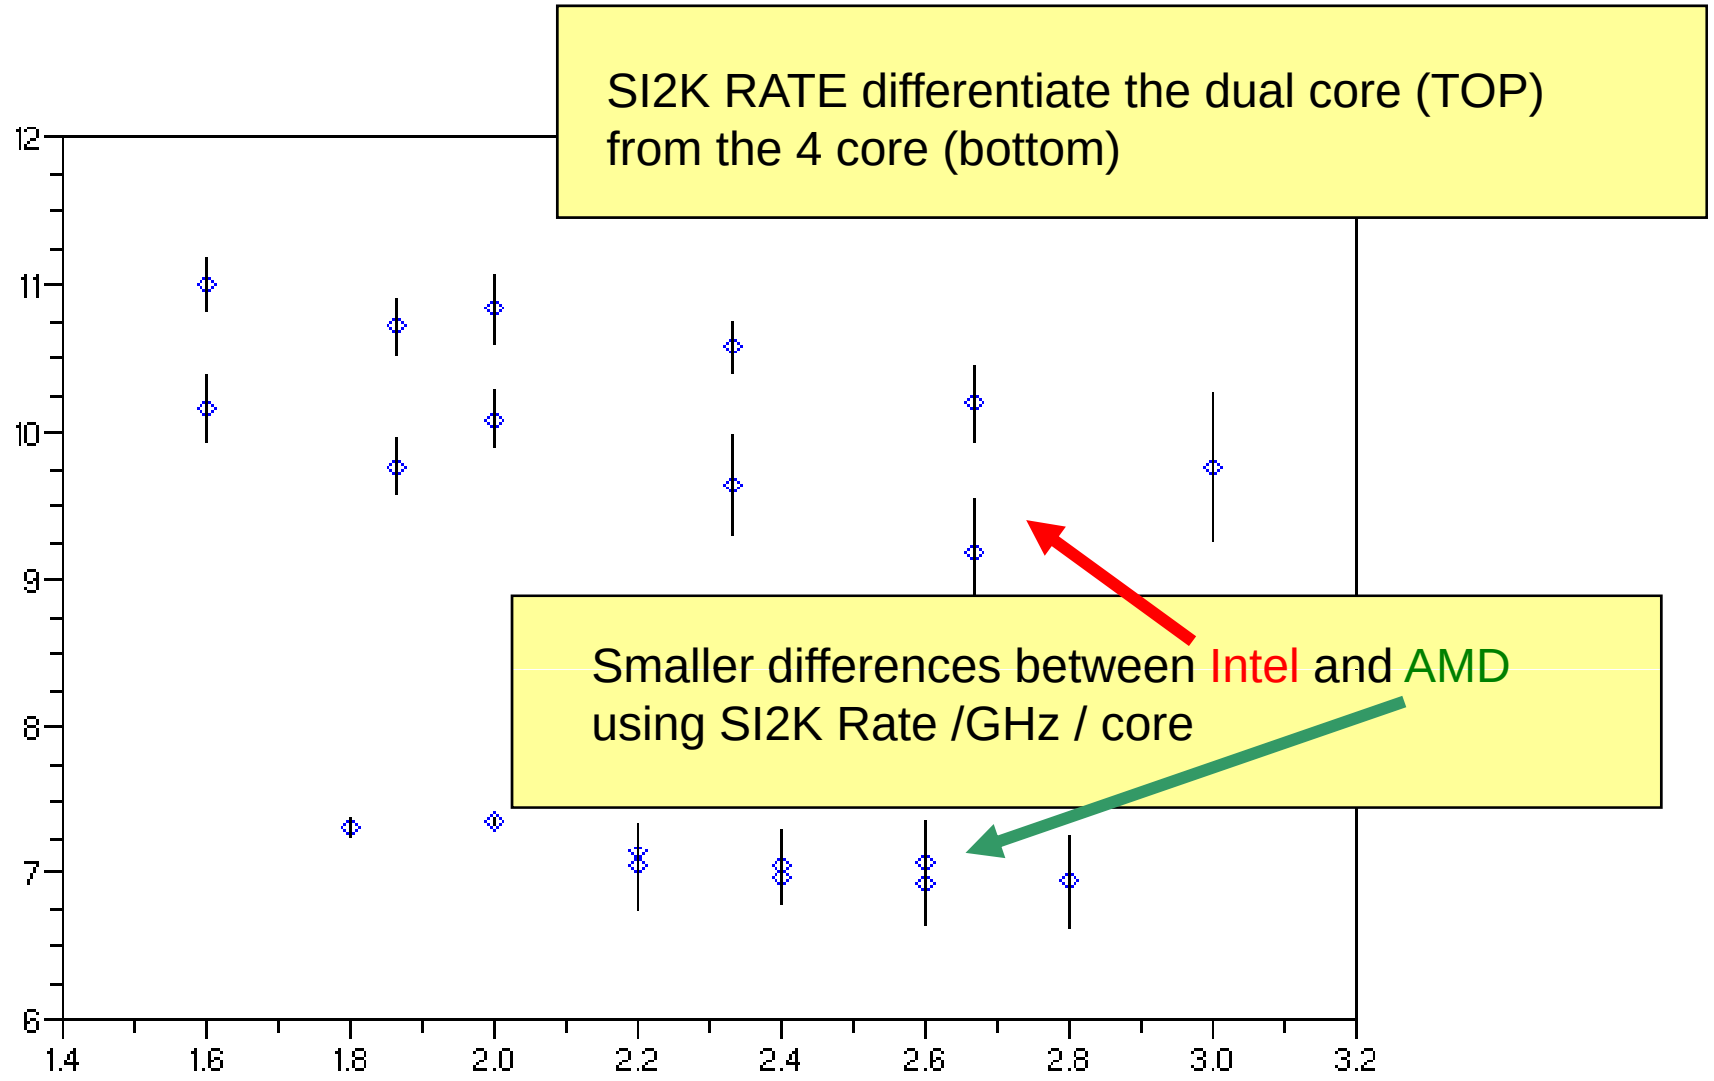

The CERN SI

- A TEMPORAY Solution:
- The intel-amd gap now is smaller
- Intel xeon scores -50%
- Amd opteron scores only -29%
- Big differences between 32 and 64 bit
- SI2K/core/Ghz taken measured by M.Alef and pubblished on the HEPIX site at Caspur

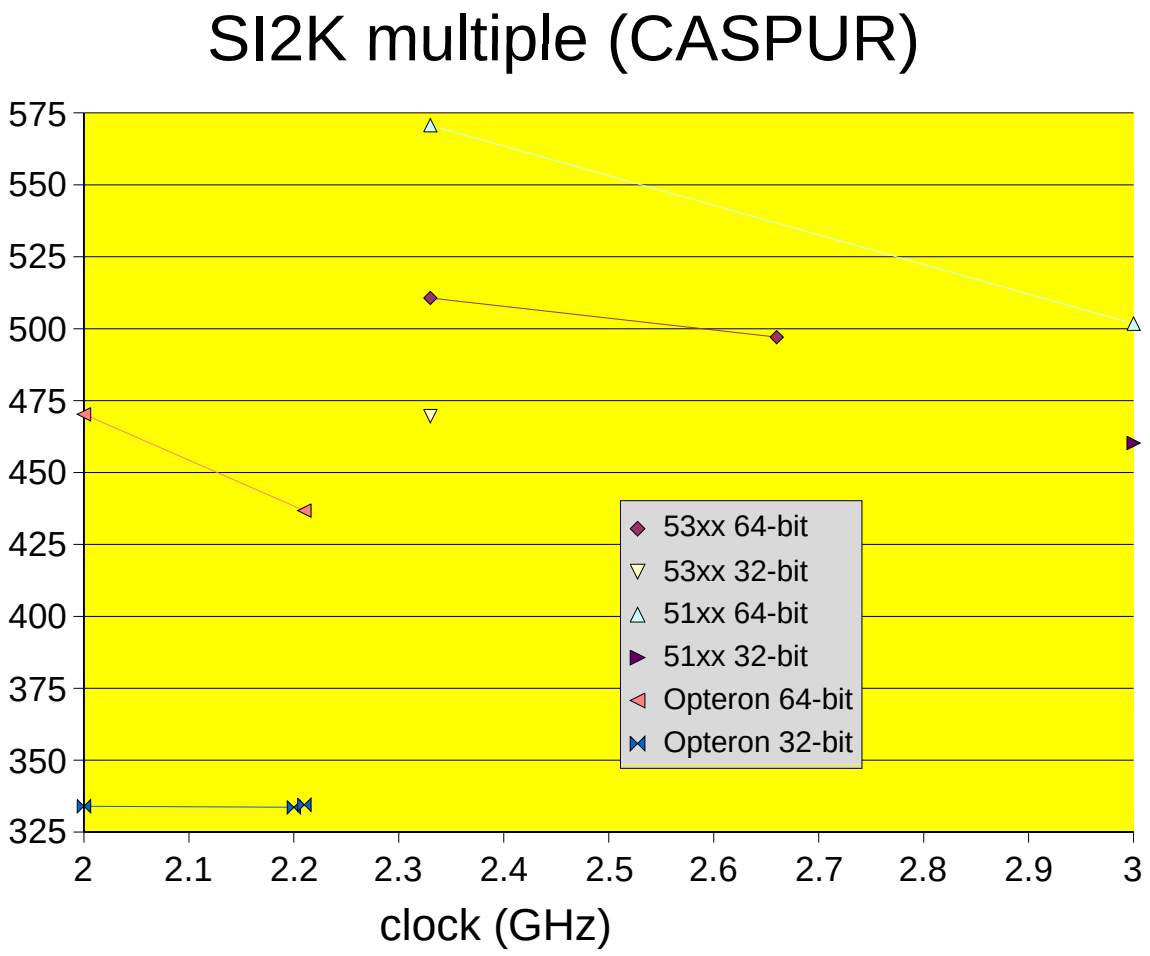

SI2K6rate

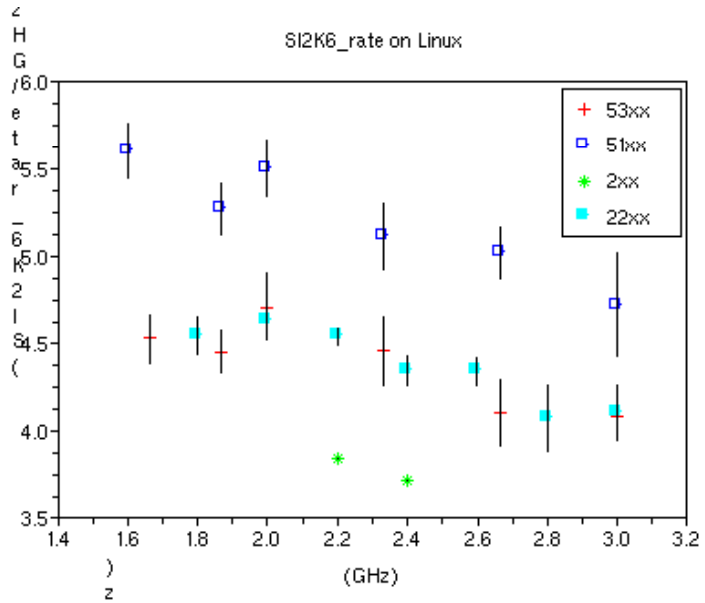
SI2K6_rate on Linux
53xx
51xx
2xx
22xx
(GHz)

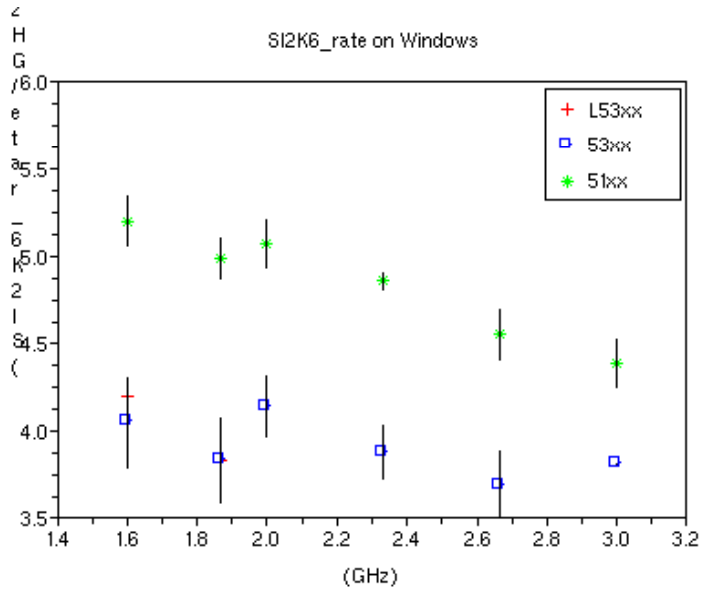
SI2K6_rate on Windows
L53xx
53xx
51xx
(GHz)

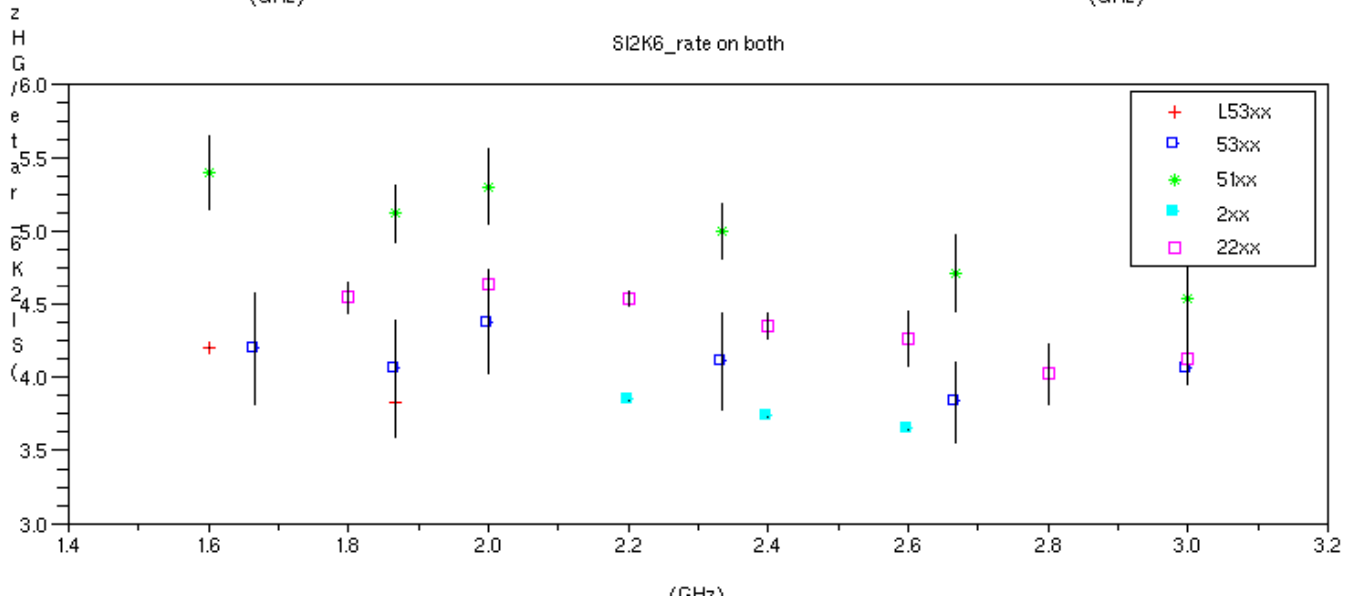
SI2K6_rate on both
L53xx
53xx
51xx
2xx
22xx

Babar TierA Results

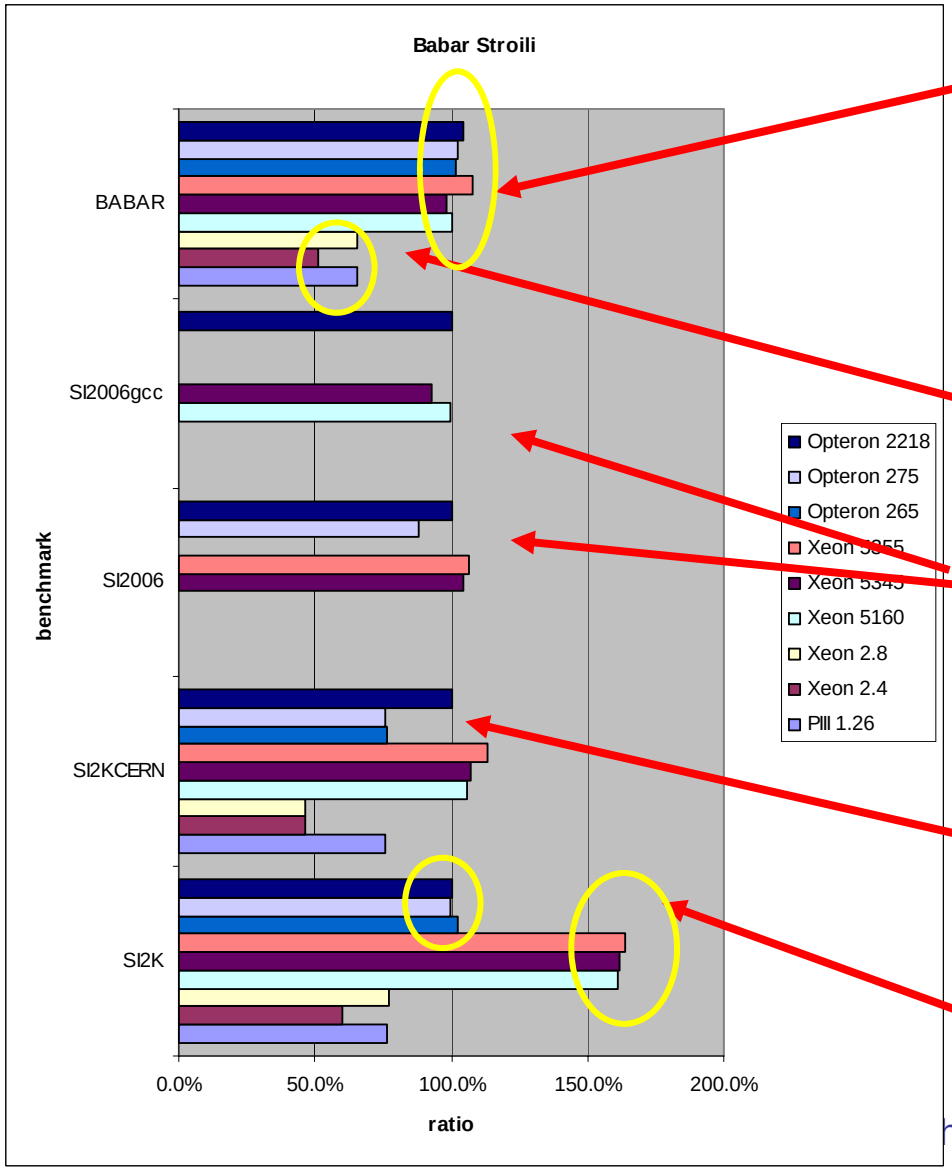

- If you normalize by core and clock all new processors have the same performance
- Doubling the older generation cpu
- SI2006 matches this pattern (pubblished and gcc ratio constant)
- SI2000-cern better than SI2K nominal
- SI2000 clearly does'nt work

CMS sw SIM and Pythia

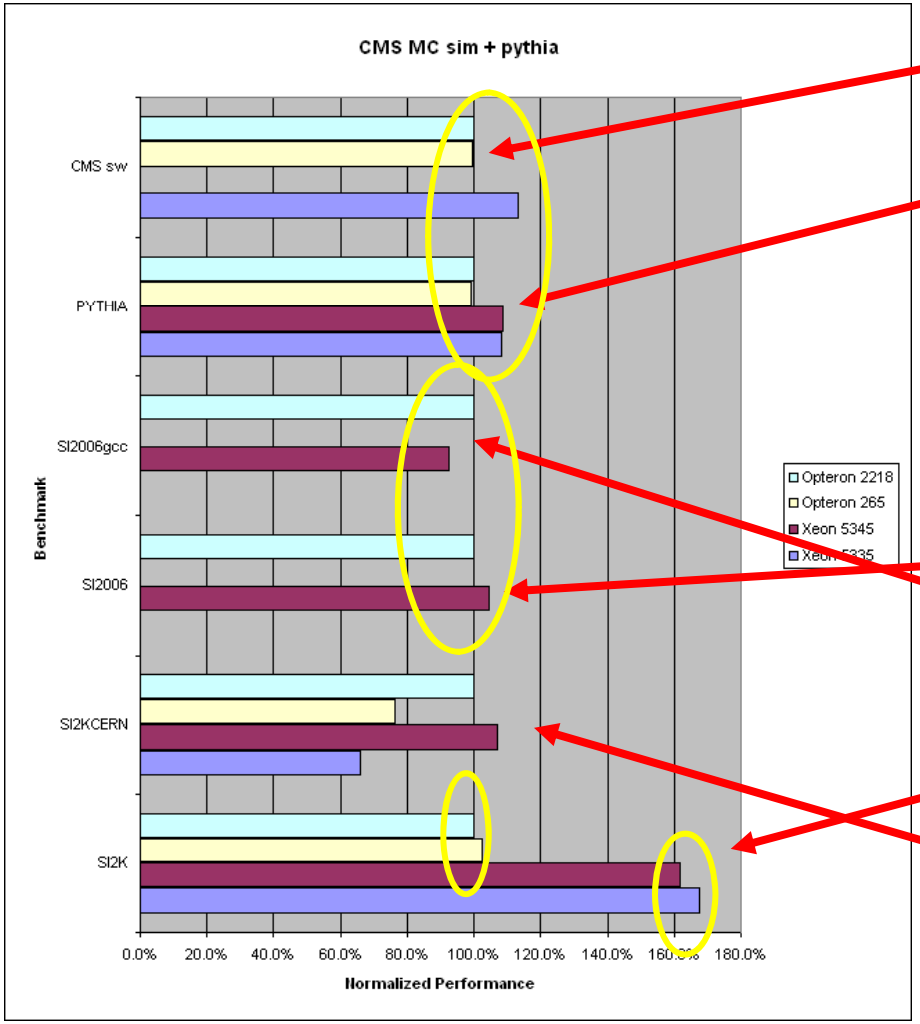

- CMS Montecarlo simulation (32bit) and Pythia (64bit) show the same performance once normalized
- Both Specint 2006 pubblished and Specint 2006 with gcc show the same behaviour
- SI2K clearly wrong. SI2K cern better but not as good as SI2006

What should be do?

- We should stop using SI2000 pubblished
- **SI2k CERN** a temporary solution but is still makes use of a 7 years old benchmark that is not mantained by SPEC (v1.3)

Why SI2006

Best candidate is Spec INT 2006 rate
measured with gcc

- Spec INT 2006 because it's new, you can find pubblished results, will have bug fixes, and will be portable to new platform compiler
- RATE because you can take in account of scalability issues with the whole machine architecture
- Measured with GCC: to keep the environment as close as possible as the experiments.

What is missing

- All the official number are still expressed in terms of SI2K (nominal)
- So funding agencies are still bounded to use SI2K for all formal agreements
- A general agreements on the next benchmark is still MISSING

Final remarks

- The working group is still collecting results (A. Crescente, L. Dell'Agnello, D. Salomoni, F. Rosso, A. Sartirana, R.Stroili, D.Salomoni, M.Morandin, A.De Salvo, T.Zani)
- NB: Results from SI2006 change with
  - Version of gcc (gcc3 vs gcc4)
  - 32bit vs 64bit
  - Tuning gcc –O2 –fPIC –pthread vs gcc –O3
- So it's difficult to have an homogeneous set of results
- Will add Atlas and Babar results when available