

# CERN openlab II

Platform  
Competence Centre  
January 2008

Main contributors:  
Andreas Hirstius  
Sverre Jarp  
Andrzej Nowak



- Second CERN/Intel multi-threading and parallelism workshop
  - 4<sup>th</sup>-5<sup>th</sup> of October 2007 (1 day lectures, 1 day labs)
  - 5 lecturers (2 Intel + 3 CERN), 45 participants, 20 people oversubscribed
  - Survey: 100% respondents' expectations met
  - Next workshop: Late Spring 2008
- Licenses for the Intel Threading Tools (and other SW products) available
  - Linux tools available publicly on AFS – no installation required
- Advances in Geant4 parallelization experiment (by Northeastern University, MA, USA)





# MT Workshop pictures

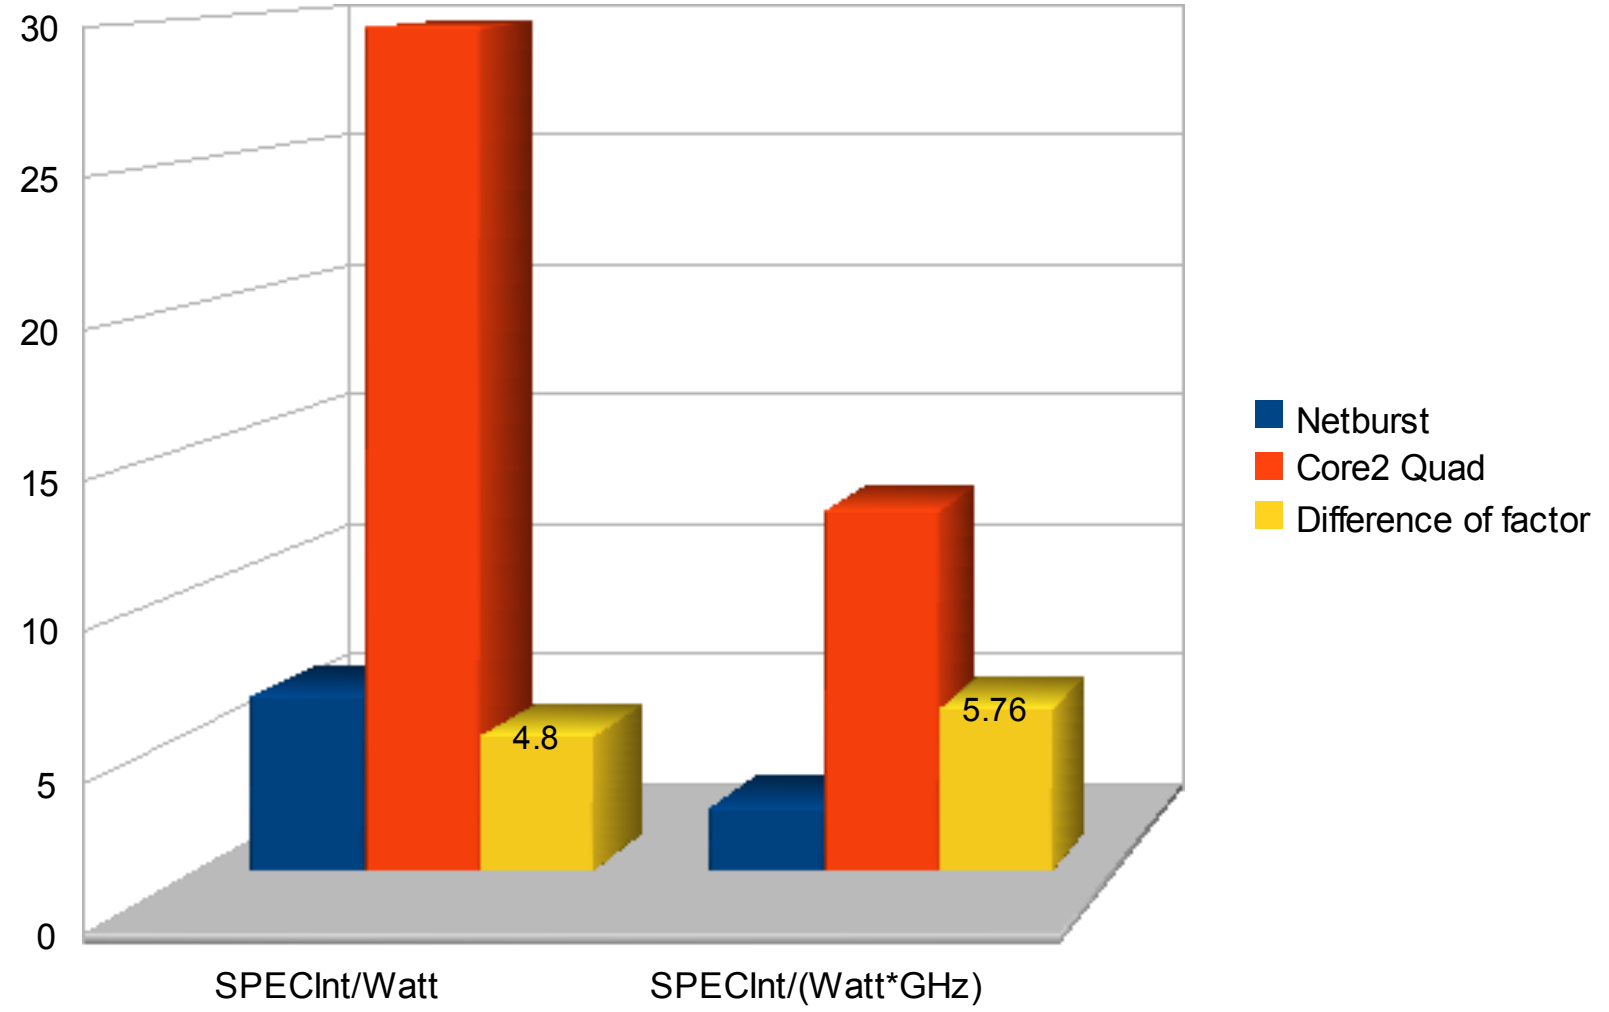


- Starting a collaboration with the parallelism R&D project in CERN's Physics Department
- Started a benchmarking collaboration with IT/FIO
- Two new benchmarks: LHCb and ALICE (HLT)
  - Håvard Bjerke (openlab) added multithreading (using TBB) to the ALICE benchmark available for both IBM CELL and Intel x86
- Performance/Architecture courses being prepared:
  - March 2008 (dry run; by invitation only)
  - Summer 2008 (CERN School of Computing)
  - Fall 2008 (CERN)

- Pfmmon – the perfmon2 Linux client
  - CERN's commitment: symbol correlation and updates for complex CERN software – ROOT, etc
  - ~12-15 man weeks coding (5 weeks real time)
    - Andrzej Nowak (CERN) + Stéphane Eranian (HP Labs)
  - Main part of CERN's contribution finalized, changes committed to CVS, YE2007
  - Perfmon2 is still on the way to become the standard Linux performance monitoring interface, minor advances made
  - Additional contributions on CERN's behalf expected
    - Testing, new features

- “Ad-Hoc” Working Group on improving power efficiency in the CERN computing centre
  - Improve efficiency as much as reasonably possible!
- Initiated a very detailed study of Core2 based systems
  - Different processors; different memory
- On-going discussions on new CC design
- Paper on power efficiency of computer centres in collaboration with Intel.
  - Will be shown to the Italian ministry of the environment

## Comparison Netburst vs. Core2



- Participation in LHC wide Benchmark Project
  - Goal: develop one or more benchmarks that represent(s) HEP workloads the best way possible
  - Starting point is SPECint2000 with CERN settings
  - All experiments are involved
  - Other HEP sites are involved (e.g. DESY, Karlsruhe)
  - Perfmon2 will be used to take a really close look
  
- Perfmon2
  - Work on infrastructure around perfmon2
    - Get it more “production” ready as we wait for final inclusion in mainline kernel and certain patches to pfmon



- Another try to get (high) into the TOP500 list
- Using as many of the ~1300 quad-core systems as possible
- Working closely with Sergey Shalnov (Intel)
  - Using his “hybrid” version of High Performance Linpack
    - OpenMP on the node
    - MPI between nodes
    - Significant reduction in communication
    - Should be significantly better than pure MPI in a large system with relatively slow interconnect – like ours
- Initial test done over Christmas

- Still very active on Intel-64 and IA-64 platforms
  - Compilers: Intel icc 10.1 & GNU gcc 4.2.2
- Benchmarks from ROOT and Geant4
  - 9 stress tests, 16 snippets, and 3 G4 examples
- 2008:
  - Target is further improvements in execution time
    - Zero regressions
  - Special emphasis on additional compiler options when using O2
    - Since CERN programmers never use O3
  - Expand to more complex benchmarks
    - Multithreading/TBB + SSE
  - Compiler expert from Intel in September/October

- Close collaboration (especially with Intel) on new activities
  - We are keenly interested in the move from multi-core to **many-core**
    - Benchmarks submitted to Intel
    - Scalability tests run on simulators
    - Encouraging results with Geant4-derived benchmarks (CMS)
  - Also review of ISA extensions
  - On the software side, we now participate in the review of the Ct language specifications
    - Initial document received (yesterday!)

- Complete the TOP500 run
- Intensive work on power efficiency projects
  - In particular the systematic power analysis of Core2 based systems
- Intensive work on benchmarking
  - Technical Student starts March 1<sup>st</sup>
- 10Gb Ethernet
  - The new Intel 10Gb NICs have been ordered
  - Promise to deliver excellent performance
- Preparations for the move of the opencluster
- Performance/Architecture class
- Accelerate “new activities”

# Q&A