

# Platform Competence Centre

CERN openlab II Quarterly Review  
9 October 2007

Sverre Jarp  
CERN openlab CTO

Main contributors:  
A.Hirstius, R.Jurga, A.Nowak,





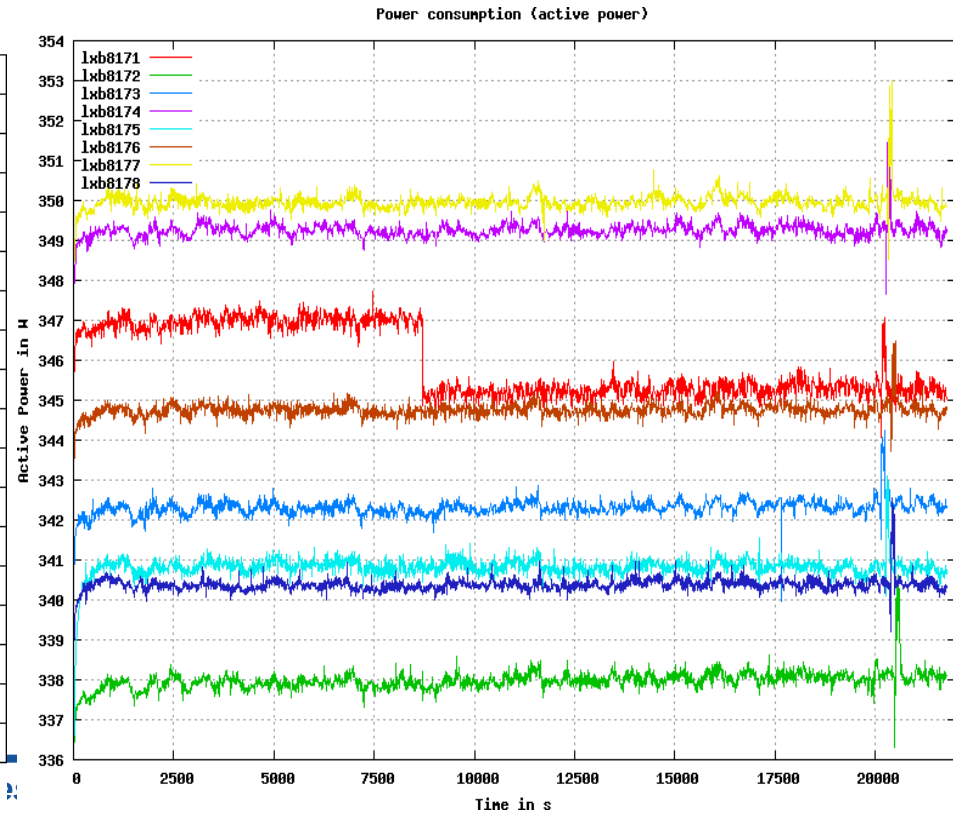
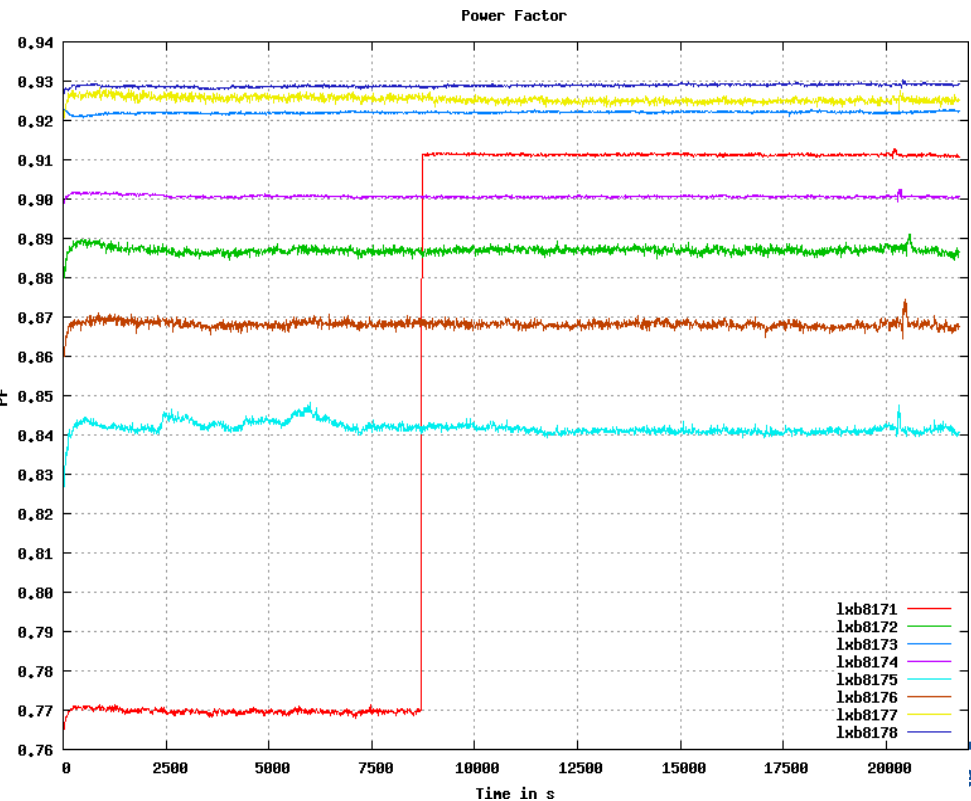
- Power measurements
- “Silent” data corruption
- New hardware
- Itanium usage
- Beta testing
- Compiler project
- Multithreading
- perfmon2/pfmon, gpfmon
- Conclusions



CERN  
openlab

# Power Measurements

- Extensive power measurements by a summer student - Piotr Jurga
  - Using “standard” batch nodes (from different suppliers)
    - 8 “identical” machines per supplier were tested
    - Differences between machines are much larger than previously assumed!
    - “strange” behavior of some machines (e.g.: sudden jump in the power factor)




- The silent data corruptions were “discovered” about a year ago on production storage and batch systems
  - Data is being corrupted “on-the-fly” with no warning or error message
  - P.Kelemen/FIO wrote a probe to collect data on those corruptions
  - Active involvement of A.Hirstius
    - Probe also run on Itanium cluster
      - Corruptions seen on disk server with RAID array
  - Any application that does not use checksums or other methods of checking dataset consistency will not detect the silent corruptions!!
- Since some of the corruptions may be caused by problems with the memory it's possible that network devices are affected as well!
  - The problem: How to find out if this is the case??

- Upgrade of 100 CPUs
  - Intel Itanium2 “Madison” to “Montecito” Dual-Core
    - 1.6 GHz
  - Included upgrade of 50 HP mainboards!
    - Very work intensive – the upgrade took several days
    - One technician from HP, part time
    - Most of the work done by openlab fellows
    - Some main boards with problems on arrival, replaced last week
- Two laptops and one desktop from Intel
  - To be used in vPro testing
- Five Xeon servers (Dual-core, 2.66GHz) arrived from HP last week
  - Already operational

- The Itanium cluster is quite extensively used for multiple activities
  - Two groups run parallel jobs using MPI based on Voltaire's Infiniband switches
  - 20 machines: Computational Fluid Dynamics
    - TS Dept.
  - 30-40 machines: Accelerator studies
    - AB dept.
  
- In addition:
  - Some systems in use by Procurve team
  - Some systems in use by IT's security team
    - Correlating data from the CERN firewall
  - Several systems used by ourselves
    - Benchmarking, compiler testing, etc.
  
- Finally, we plan to retire old 1.3 GHz single-core systems

- Access to “Harpertown” in August
- Penryn family member
  - 45 nm technology
  - Dual Core and Quad Core
  - Large L2 cache (6MB/2 cores)
  - Fast Super Shuffle Engine
  - Fast Radix-16 Divider
  - Enhanced VT-x
  - Clock speed > 3 GHz
  - Higher bus speeds
  - Increased Performance and Energy Efficiency
- CERN/HEP will benefit from several of these enhancements
- Also
  - NDA discussions about Intel’s future plans
    - Involving partners from IT (especially FIO) and the LHC experiments



- Effort continues as before around three compilers:
  - Intel 10.0, open64 4.0, gcc 4.2/4.3
  - IA-64/Itanium and Intel-64/Xeon
  - New snippet from summer student (Gabriel Esteves)
- Still very active in reporting bugs into Intel's problem database
  - See quote from Pat Gelsinger
  - ROOT and Geant4 used for checking new releases
    - Correctness and performance
- Interesting meeting in Portland prior to CHEP
  - Evolution of existing compilers
  - Compilers for new (derivative) architectures
- Intel talk on Thursday:
  - Ct
    - New language based on Nested Data Parallelism (NDP)





## Multi-threading and Parallelism WORKSHOP

4th-5th of October 2007, CERN

A second instance of the Multi-threading and Parallelism Workshop will be held on the 4th and 5th of October 2007 at CERN. Experts from Intel will lead the two day event and help you improve your knowledge by explaining the key intricacies of parallel programming and presenting the most efficient solutions to popular multi-threading problems.

### Event highlights:

- Day 1, Fundamental aspects of multi-threaded and parallel computing:
  - The move to multi-core and its impact on software
  - Important parallelism and multi-threading concepts
  - Threaded programming methodology and scalability issues
  - OpenMP and POSIX Threads discussion
  - CERN-specific parallelism related topics
- Day 2, Hands-on lab
- Q&A with Intel experts - all topics, from beginner to advanced

The workshop is co-organized by CERN openlab and Intel for users affiliated with CERN. Registrations are based on a first-come first-served basis by sending e-mail to [openlab\\_workshops@cern.ch](mailto:openlab_workshops@cern.ch). 50 seats are available for day 1 and 40 seats for day 2 - please indicate which are of interest to you in your registration message. Go to <http://cern.ch/openlab> for details. All registrations will be confirmed by e-mail.




<http://cern.ch/openlab>

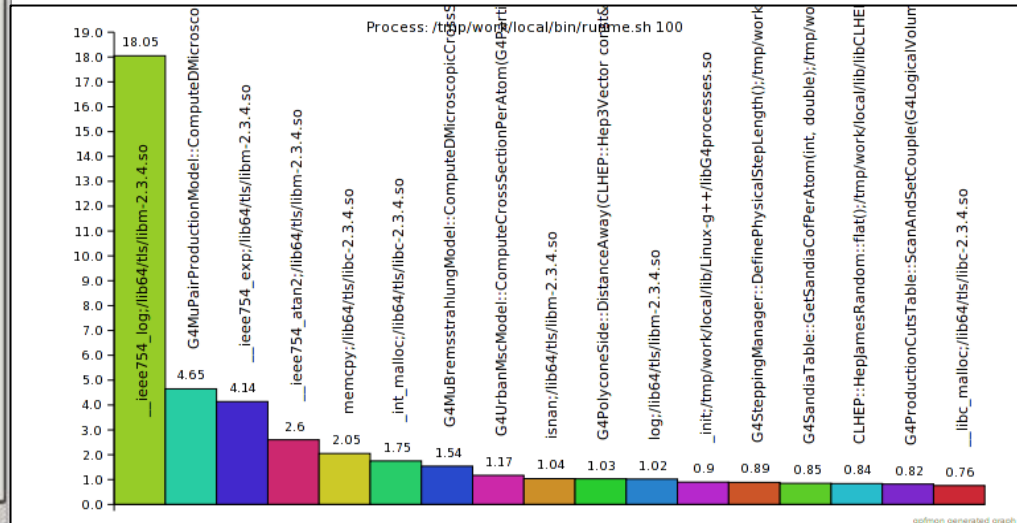
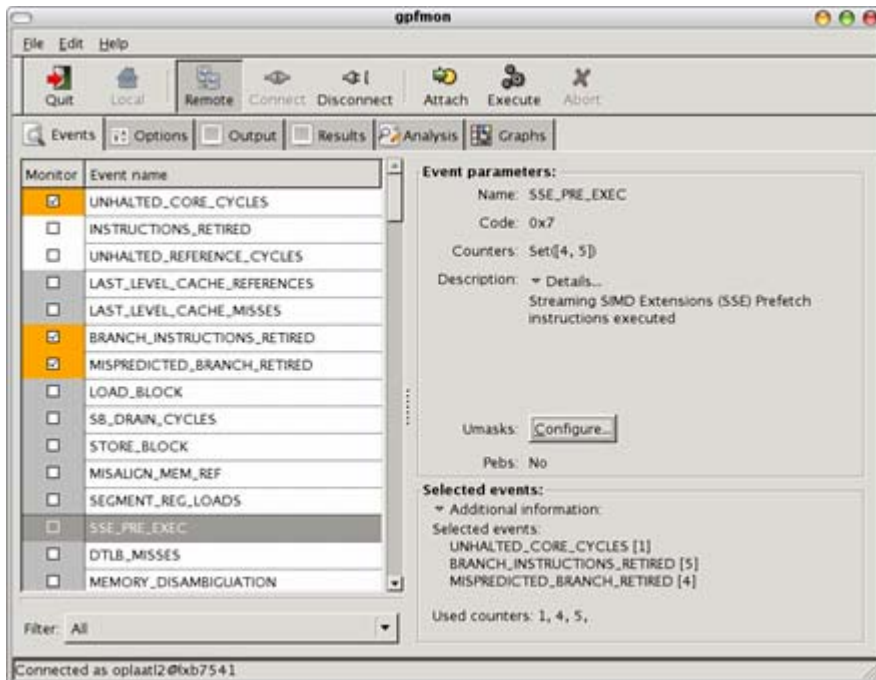


- Two MT workshops held
  - 31 May/1 June (40/20 participants)
    - Article in September CNL
  - 4 & 5 October (50/40 participants)
  - Thanks to H.Cornelius and H-J.Plum who taught on both occasions
    - Emphasis on openMP, pThreads, TBB
      - Intel ThreadChecker, ThreadProfiler, VTUNE as tools
    - Three CERN presentations as well
      - S.Jarp, E.McIntosh, F.Rademakers
    - A lot of background preparation done by A.Nowak and A.Hirstius
- Already a waiting list for 2008
- Transfer teaching responsibility to openlab
  - A.Nowak, S.Jarp
  - In close collaboration with Intel

- perfmon2
  - Ubiquitous interface to hardware performance monitoring units (PMUs) in Linux kernel
    - Most likely, now in kernel version 2.6.24
    - Author (S.Eranian/HP Labs) initially targetted Itanium only
  
- User interface: pfmon
  - CERN additions developed for correct symbol resolution
    - Loading/unloading of shared libraries
    - Transfer of control
  - Also bug fixes
  - Active involvement of J.Cañizales Díaz (student)
  - Several production machines updated with the new software
  
- Presentation/paper at CHEP07
  - “perfmon2: A leap forward in performance monitoring” (Jarp, Jurga, Nowak)



- GUI on top of pfmon:
  - More intuitive approach
  - Still in heavy development by A.Nowak



- Very active period
  - Broad portfolio of activities
  - Lots of new results
  - Great relationship with our partners
- We expect all activities to continue “full steam” in the coming quarter
  - and well beyond!