Platform Competence Centre activities overview

Havard Bjerke, Sverre Jarp, Andrzej Nowak

January 29th 2009

openlab III – outlook and staffing

- **openlab III started on January 1st 2009**
- **Exciting new technologies ahead of us:**
  - New Xeon generations
  - Larrabee
  - Atom
  - New compilers and other Intel software tools
  - many more
- **Foreseen composition of the Platform Competence Center:**
  - Sverre Jarp (supervisor, openlab CTO)
  - 1 staff member – Andrzej Nowak (Intel)
  - 1 fellow – Julien Leduc (Intel; March arrival)
  - technical and summer students – Axel Busch (March arrival), others

Equipment

openlab II related developments

- All openlab machines moved to a cooler place in the computing center (some failures along the way)
- New 512-core Core 2 cluster installed – 64 blade systems from HP
- Some older Itanium 2 machines being retired

Remaining openlab II contributions

- 4 Dunnington systems
- 32 Nehalem-EP systems (Spring 2009)
- 4-8 new Itanium (Tukwila) systems

Possible openlab III contributions

- Nehalem-EX systems

Initial emphasis in openlab III (1)

- **Intel Nehalem – plans for 2009**
  - Expansion to Nehalem-based Xeons
  - Thermal and performance measurements
  - Study in detail the impact of multi-threading
- **Intel graphics processor**
  - already some work done in 2008
    - SDK experimentation,
    - Training
    - Initial analysis of the planned hardware and its potential
    - Testimonial published
  - plans for 2009:
    - Hardware
    - OS support

Initial emphasis in openlab III (2)

- **Intel Atom**
  - Done in 2008
    - Power consumption analysis performed on a consumer-grade system
    - Performed some extrapolations relating to power and cost
    - Paper published on openlab website
  - plans for 2009
    - More tests with newer hardware/chipset
    - General usability discussions with Intel
- **Close monitoring of the whole low-power CPU ecosystem**
- **New Intel software technologies (Ct and others)**

Teaching (1)

- **Two workshops held in Autumn**
  - Multi-threading and Parallelism workshop (Nov 11 – Nov 12)
  - Computer Architecture and Performance Tuning workshop (Oct 14)
- **Recent changes:**
  - The multi-threading workshop has been re-designed to align better with the students' expectations
  - Cemented collaboration with Technical Training
- **Cyclic repetition of both classes**
  - At least 4 workshops planned for 2009

Teaching (2) – Future developments

- **Expansion of the architecture/tuning workshop to 2 full days with more focus on architecture & compilers**
- **Intel involvement in both workshops (Jeff Arnold)**
- **Participation in the CERN School of Computing**
- **Participation in other workshops and courses is being discussed**
  - i.e. ESC'09 – INFN, Italy,
  - CHEP'09 – talks on perfmon, Atom and many-core

Performance monitoring

> Perfmon

- v3 published by Stephane Eranian (Google) – minimal patchset for the Linux kernel
- Rival patch posted to LKML by Ingo Molnar
- Situation unclear, long battle ahead
- Pfmon reworked “under the hood”
- Due to firmware problems, perfmon doesn’t run on any of our new HP systems

> Related report published by G. Balazs

- With help from Intel we’re trying to understand some interesting results related to bus traffic events

Solid State Drive testing

- **32 GB NAND SLC drive received from Intel – enterprise class**
- **Compared with a 7200RPM SATA drive in streaming and with an STEC MACH 8 in random ops**
- **Streaming:**
  - Reads, writes – ~4x faster than SATA
  - Random reads – ~3.8x faster than SATA
  - Random writes – ~5x faster than SATA
- **Random ops (filesystem like – 4kB blocks; tested by FIO)**
  - Excellent random access properties
  - Excels at reads
  - Roughly equivalent to STEC MACH 8
- **Other tests performed by Eric Grancher**

Multi-threading

- **Geant4 work progressing well**
  - Xin Dong reports 22 MB per thread!
- **Candidate report submitted for review**
- **Jeff Arnold's work at openlab**
  - 2.5 mo sabbatical spent at openlab
  - measured memory consumption of the LHCb framework
  - A report is on the way

Thermal measurements

- **Comprehensive report submitted by Gyorgy Balazs**
  - Components measured: CPUs, RAM, Drives

Intel Compiler project

- **Intel icc 11.0 beta completed**
  - Several new (mainly performance-related) incidents reported back to Intel
  - Now awaiting 11.1 beta
- **Presentation given by Intel specialist on in-order optimization on ATOM**
  - Rather different from Itanium; but also promising
- **New activity:**
  - Compare ratios when SPEC2006 C++ programs are compiled and run on both Intel and AMD CPUs using either the Intel or the GNU compiler

Alice High-Level Trigger tracking

> Latest developments

- Merged with AliRoot framework
- Interactive 3D visualization
- Multithreaded tracking with TBB
- Algorithm optimization
- SIMDization

HLT – Visualization

- Interactive visualization crucial to understanding the effect of algorithm changes – developed together with Intel Bruehl

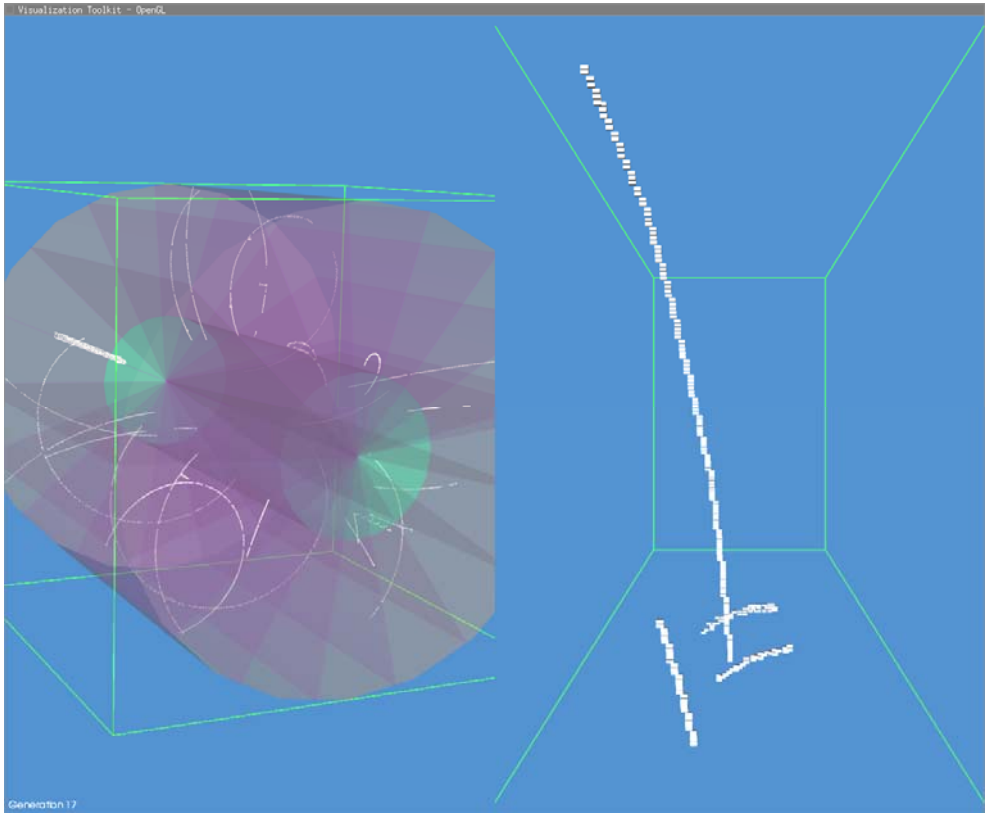

HLT – Algorithm optimizations

- **Shortcut the Cellular Automaton**
  - Only execute one iteration
  - Find tracks during the Cellular Automaton
  - Order of magnitude speedup
- **The event data is digitized in order to allow integer SIMD calculation**
- **Lowest precision possible for data structures**
  - More efficient use of cache, memory
  - Higher throughput for SIMD instructions
- **Use relative positions for lower precision**
  - Absolute neighbor position has > 8 bit precision
- **Allows calculating the neighbors of 16 cells in parallel using SSE**

HLT - Results

- Shows opportunities for parallelism in tracking
- Can exploit both multi-core and SIMD concurrency
- Not linear scaling, but benefits from an increase in hardware concurrency
- Plan to use the software in the HLT cluster
- A report from Havard is forthcoming
- Participation in an Alice workshop in February

Q & A

CERN openlab