

Platform Competence Centre activities overview

Sverre Jarp, Julien Leduc, Andrzej Nowak

January 28th 2010



PCC activity overview – outline

> News

- News from the quarter that passed

> Equipment status

- Status of the openlab cluster

> Processor and system investigations

- Scalability testing and beta systems evaluation

> Many-core activities

- Multi-threaded software development and scalability testing

> Compiler studies

> Performance monitoring

> Teaching

- Workshops, schools, etc

- > **New member of the team – Alfio Lazzaro (Fellow – starting Feb 16th)**
 - multi-threading studies
 - valuable practical experience with parallelization and real physics applications

- > **Continued collaboration with Intel**
 - openlab visited IDF 2009 in San Francisco and SC'09 in Portland – fruitful meetings concerning Xeon evolution, many-core etc.
 - Visits:
 - Richard Draycott's visit (September 17)
 - Visit from Intel Labs: Joe Curley, Pradeep Dubey, Mikhail Smelyanskiy
 - Rajeeb Hazra from Intel (exascale computing)
 - Planned visits:
 - 1 day visit of David Levinthal planned for Spring
 - PSR team to visit CERN on March 12th




- > **Current openlab-II built cluster is being actively used**
 - 64 Core2 blade systems being used for workshops, individual research of collaborating groups (experiments, frameworks) and other PCC activities
 - IA-64 systems still in use! Migrating all systems from Scientific Linux to Red Hat (SLC support ended)
 - Sizeable acquisition of Viglen Nehalem-EP systems (together with FIO – delivery Q1 2010)
- > **Until the systems arrive: 24 Nehalem-EP servers on loan**
 - Used for workshops and ongoing research in the domains of performance optimization, compilers and multi-threading
 - Equipped with performance monitoring tools, behaving better than the previous batch
- > **Several beta systems added, others “en route”**
- > **No production systems added since last meeting**

Processor and system investigations

- > **Evaluated numerous solutions from Intel**
 - Westmere-EP beta (6 cores per CPU, rather than 4)
- > **Evaluated extensively a Dunnington system (4 * 6 cores)**
- > **Nehalem-EX beta system (4 * 8 cores)**
 - Currently being investigated
 - Interesting results
- > **Other upcoming many-core solutions from Intel being investigated**
- > **Tukwila testing possibilities on the horizon**
 - Next generation Itanium architecture

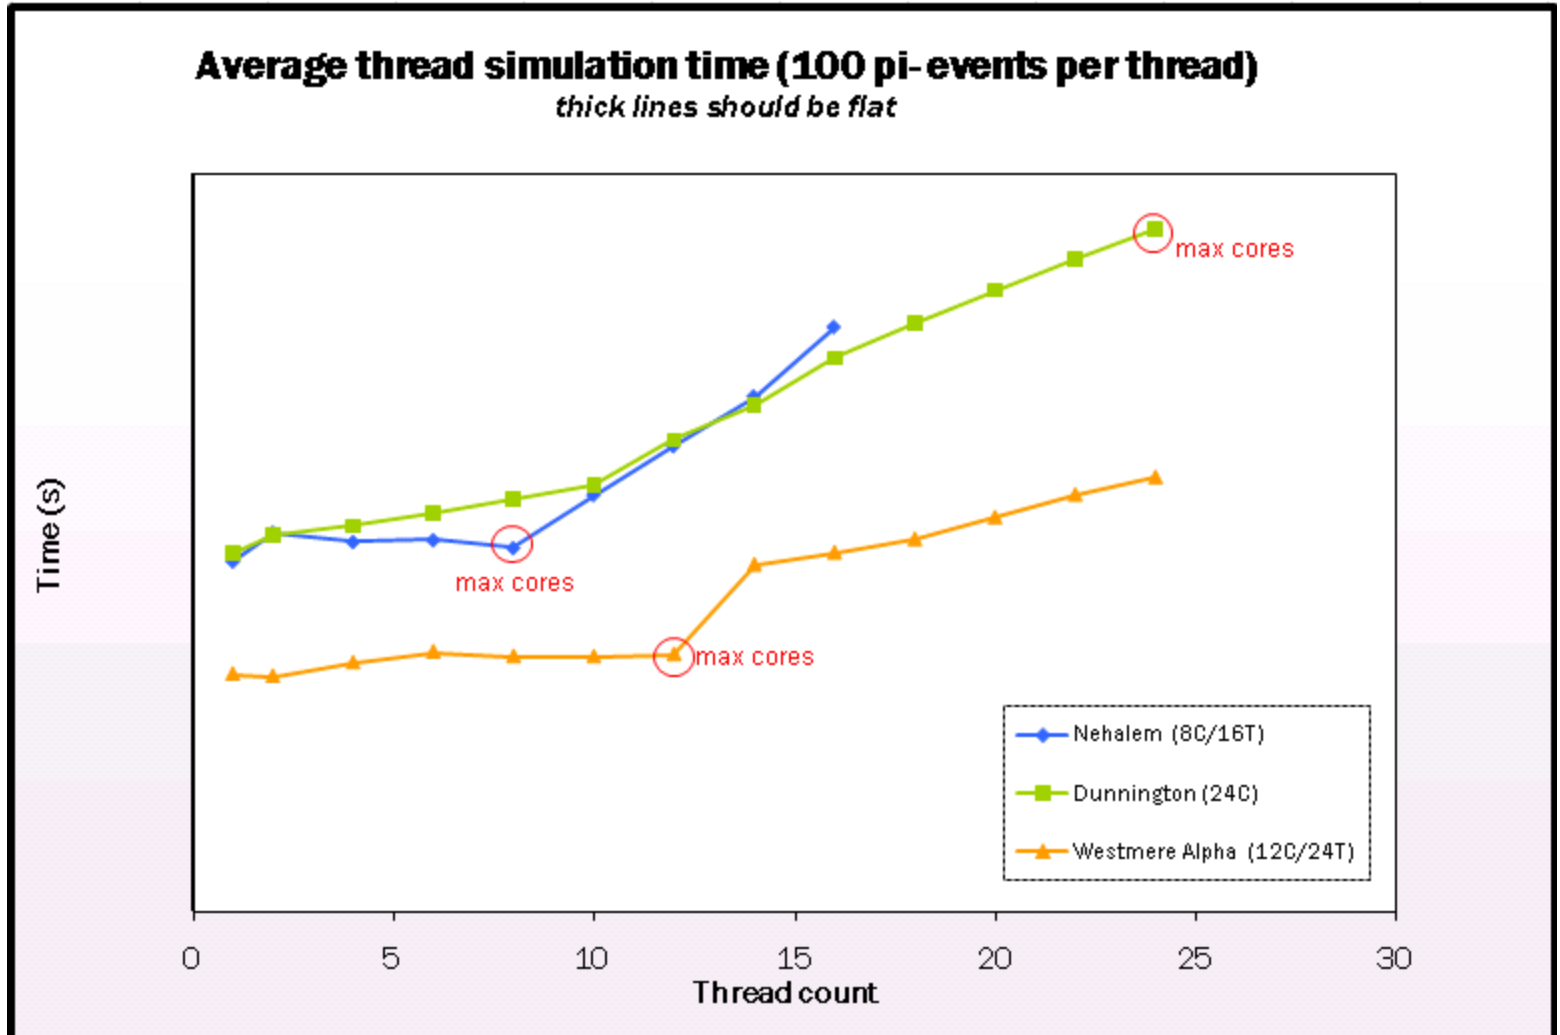
Nehalem-EX comparison

HEP SPEC 2006 64-bit NHM-EX



- > **Continued collaboration on a multi-threaded Geant4 prototype**
 - Scalability tests on various pieces of hardware (including Harpertown, Dunnington, Nehalem-EP, Nehalem-EX, Westmere-EP)
 - Preparing to expand tests to other hardware
- > **Multi-threaded version of “test40” being investigated**
- > **ALICE trackfitter to be revisited in the light of new technologies, such as Ct**
- > **ROOT minimization benchmark (from Alfio Lazzaro)**
- > **New activity: Regular benchmark papers to be published shortly after platform availability (see Nehalem-EP paper for reference)**

Multi-threaded Geant4 (not scaled)




```
void G4VUserPh
{
  if (newPart1
  if (newPart1
#ifdef G4VERBO
  if (verbos
    G4cout <
    G4cout <
    G4cout <
  }
#endif
  return;
}

// create new
if (newManage
// Add Proc
if (newPart
// Create
G4Particl
(G4P

if (gener
G4Proce
if (ion
newMa
} else
// no
newMa
}
```

> Evaluated 11.1 series compilers

- Main focus on Intel64

> Compilers (and other tools) are available from openlab CERN-wide – pre-installed and ready to use

- Maintained regularly – latest ICC and IFORT versions made available
- Ad-hoc support provided
- Interest is growing visibly

> Interesting 3-party discussions on floating-point consistency (Intel, openlab, accelerator people)

Performance optimization

- > **Continued performance studies of HEP software**
- > **Support provided for perfmon2, continued updates**
- > **More extensive investigations of Intel PTU**
 - Interesting functionality and capabilities
 - Some issues with common HEP software
 - Looking forward to a new version
- > **Performance monitoring strategy interviews conducted with all major players at CERN**
 - The 4 experiments: ATLAS, ALICE, CMS, LHCb
 - Two major frameworks: Geant4 and ROOT
 - Paper will be published soon
- > **Paper on useful performance monitoring techniques planned**
- > **Useful discussions w/ Intel and the perfmon2 author around the Nehalem Performance Monitoring Unit**

> Regular workshops still being held

- With the help of Jeff Arnold from Intel and two guest speakers (Lorenzo Moneta and Alfio Lazzaro)
- Interest maintained
- The program is constantly updated to match the newest technologies and trends; continued good attendance and interest
- Multi-threading and Parallelism
 - Nov 11/12 2009, May 4/5 2010
 - Multi and many-core technologies
 - Intel Threading Software tools discussed
 - Moved to Nehalem-EP systems



> **Computer Architecture and Performance Tuning**

- Oct 6/7 2009, Feb 9/10 2010
- Performance optimization
- Computer architecture
- Compilers
- Moving to Nehalem-EP systems

> **Performance optimization course also taught at the ESC'09 school**

- 12-17 October 2009, Bertinoro, Italy
- 25 attendees
- Full performance course taught, as well as some multi-threading elements

> Workshops for expert users organized in collaboration with Intel

- Reminder: “Inside the core” performance optimization workshop held on September 15/16
- Another workshop on many-core and multi-threading is going to take place on Feb 17/18
 - Best threading strategies
 - NUMA architectures
 - New tools from Intel
 - Hands on exercises
 - Led by experts from the Intel offices in Germany


> Inverted CERN School of Computing – mentoring program

- Mentorship for 2 lectures

Plans for the near future (1)

> Emphasis on many-core evaluations


- We're already in the many-core era!



Nehalem-EX (8 cores)



Larrabee (32+ cores)



22nm wafer

- Upcoming hardware to be tested
- Developing new benchmarks, assisting with parallelization activities

> Continued workshops and teaching

> Continued compiler and performance optimization activities

Plans for the near future (2)

> **Continued platform investigations**

- Regular benchmark papers to be published as platforms become available

> **Collaboration on Education with India (ICT) / Master Thesis Project Proposal**

- Develop a performance benchmarking framework, allowing benchmark run reproduction and comparison using data mining techniques

> **Equipment:**

- More frequent, annual upgrades of the openlab cluster planned, thanks to the collaboration with CF (“Computing Facilities” – ex FIO group)
- Continued slow phaseout of older IA-64 systems

Q & A



CERN
openlab