

# CERN openlab III


## Major Review Platform CC



Sverre Jarpe  
Alfio Lazzaro  
Julien Leduc  
Andrzej Nowak

- 3 workshops already held this year:
  - Computer Architecture and Performance Tuning: 17/18 February and 22/23 September
    - Performance optimization
    - Computer architecture
    - Compilers
  - Multi-threading and Parallelism: 4/5 May
    - Multi and many-core technologies
    - Intel Threading Software tools demonstrated
  - New lecture added about NUMA specific issues in both flavors
  - Jeff Arnold from Intel's compiler group continues as a regular teacher
  - David Levinthal from Intel gave a short version of his performance monitoring lecture
  - High demand: 35 attendees in September
- Additional workshop planned for this year:
  - Multi-threading and Parallelism 10/11 November

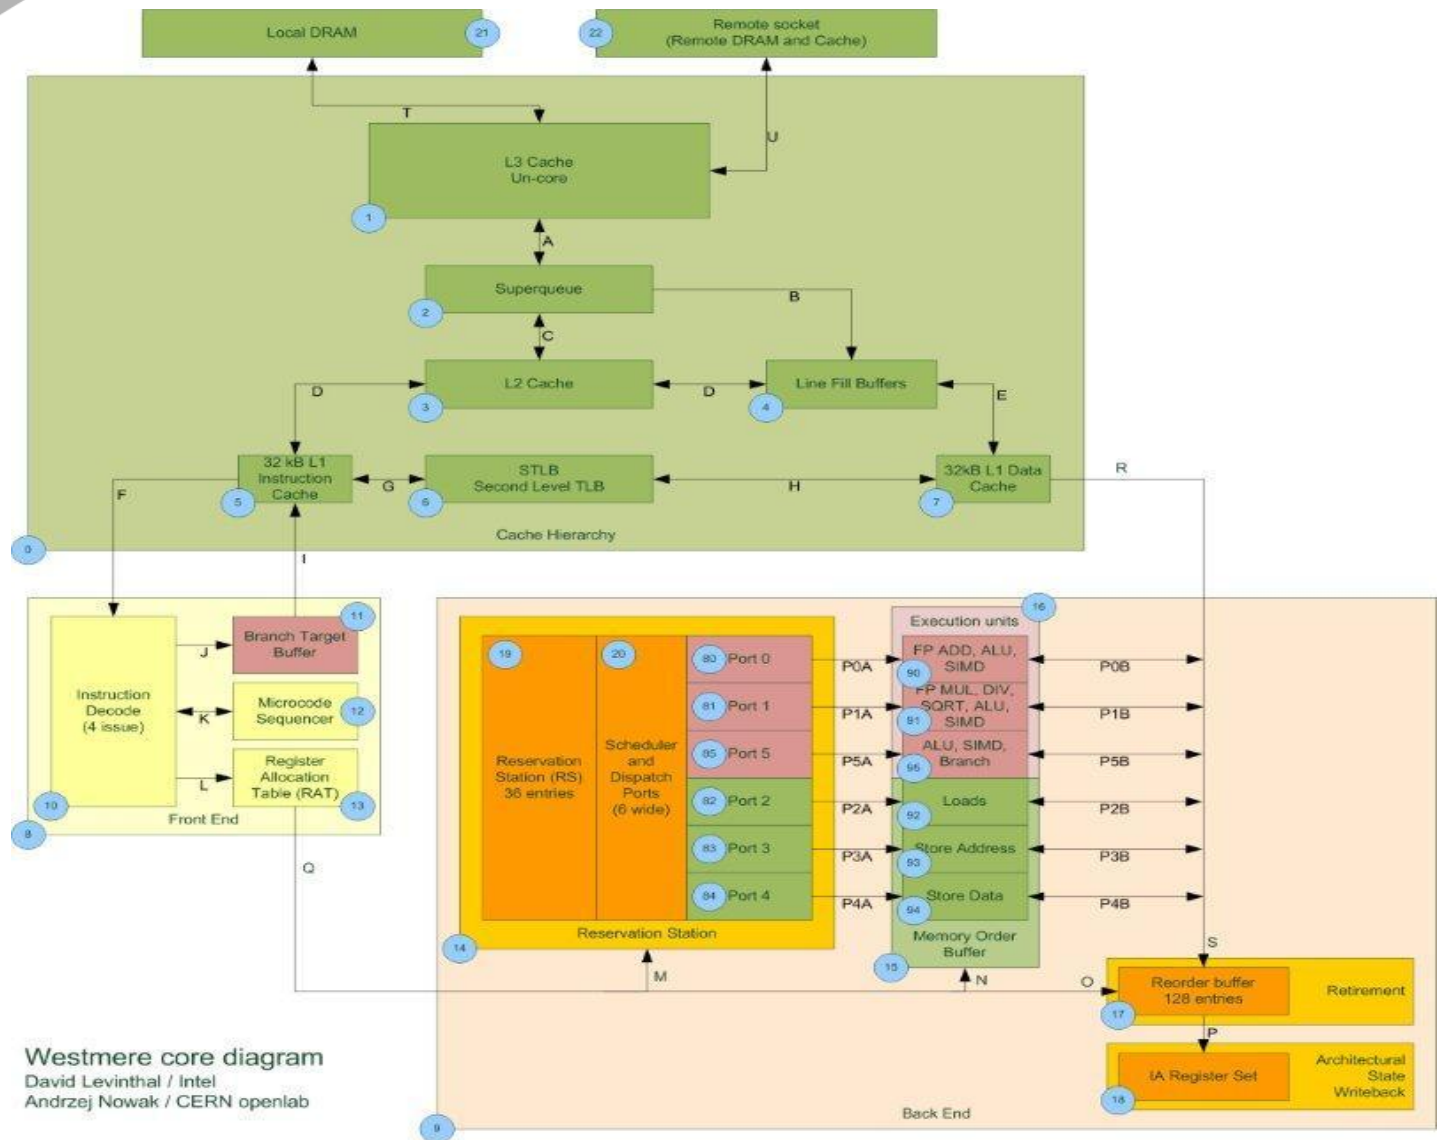
- Teaching at the CERN School of Computing 2010 (August, Uxbridge, UK)
- Revised performance tuning program to be offered at the INFN computing school (November, Bertinoro, IT)
- In addition:
  - openlab/Intel Special workshop “Manycore Optimization” (17-18 February) Levent Akyil, Herbert Cornelius, Mario Deilmann
  - IT seminar and exercise session driven by Kathleen Knobe on “Concurrent Collections” (7 May)
  - 2 workshops on advanced performance monitoring given by David Levinthal (6 May + end of July)
    - 20 participants from CERN
    - Centered on Intel’s Performance Tuning Utility (PTU) and advanced software optimization strategies



- Continued perfmon2 deployments
  - In close contact with the author
  - Continued contribution to development
- Paper on performance monitoring strategies at CERN published
- Close collaboration with David Levinthal
  - PTU 4 now compatible with perfmon2
  - Extensive PTU 4 experiments (including summer student work)
  - Westmere and Nehalem event maps and analysis facilities created
  - Next generation tool being developed jointly
- Live background monitoring solution still considered
- Continued teaching efforts with tangible effects



# Performance tuning activities (2)



- 60 New Viglen L5520 based systems installed beginning of July
  - Standard CERN systems (CDB, IPMI management...)
  - 3 \* 1TB disks in RAID0
  - Perfmon2
  - Used for all the workshops (up to 40 system used already)
- New Westmere systems expected in November
  - Around 60 dedicated new systems
- Many old Itaniums to be withdrawn

# Evaluation of the Intel Westmere-EP server processor (1)



## Evaluation of the Intel Westmere-EP server processor

Sverre Jarp, Alfio Lazzaro, Julien Leduc, Andrzej Nowak  
CERN openlab, April 2010 – updated version 1.1



### Executive Summary

In this paper we report on a set of benchmark results recently obtained by CERN openlab when comparing the 6-core “Westmere-EP” processor with Intel’s previous generation of the same microarchitecture, the “Nehalem-EP”. The former is produced in a new 32nm process, the latter in 45nm. Both platforms are dual-socket servers. Multiple benchmarks were used to get a good understanding of the performance of the new processor. We used both industry-standard benchmarks, such as SPEC2006, and specific High Energy Physics benchmarks, representing both simulation of physics detectors and data analysis of physics events.

Before summarizing the results we must stress the fact that benchmarking of modern processors is a very complex affair. One has to control (at least) the following features: processor frequency, overclocking via Turbo mode, the number of physical cores in use, the use of logical cores via Simultaneous Multi-Threading (SMT), the cache sizes available, the memory configuration installed, as well as the power configuration if throughput per watt is to be measured. We have tried to do a good job of comparing like with like.


In summary, we see good scaling with the core count. We observed a very appreciable throughput increase of up to 61% when using the in-house benchmarks, compared to the previous processor generation. Our variant of the SPEC benchmark rate, “HEPSPEC06”, gives 32% more throughput. HEPSPC per watt is measured to improve by up to 23% which is less than the improvement when going from Harpertown to Nehalem (36%). Benefits of SMT were seen to be of similar significance as in the previous processor generation.

## Early production level Westmere CPU (X5670) was extensively tested


- Shrink of the 5500 series CPUs “tick”
- 6 cores (12 threads)
- Turbo
- SKU limited to X only
  - With Nehalem-EP we had the full range

## Dual socket system hosting 12 cores (24 threads) and 24 GB of memory

## The final version of the paper was ready the day when the Westmere was launched



HEPSPSPEC06 performance comparison  
SMT-off Turbo-on



Efficiency comparison across  
3 generations of CPUs





## Evaluation of the Intel Nehalem-EX server processor

Sverre Jarp, Alfio Lazzaro, Julien Leduc, Andrzej Nowak  
CERN openlab, May 2010 – version 1.1



### Executive Summary

In this paper we report on a set of benchmark results recently obtained by the CERN openlab by comparing the 4-socket, 32-core Intel Xeon X7560 server with the previous generation 4-socket server, based on the Xeon X7460 processor. The Xeon X7560 processor represents a major change in many respects, especially the memory sub-system, so it was important to make multiple comparisons. In most benchmarks the two 4-socket servers were compared. It should be underlined that both servers represent the "top of the line" in terms of frequency. However, in some cases, it was important to compare systems that integrated the latest processor features, such as QPI links, Symmetric multithreading and over-clocking via Turbo mode, and in such situations the X7560 server was compared to a dual socket L5520 based system with an identical frequency of 2.26 GHz.

Before summarizing the results we must stress the fact that benchmarking of modern processors is a very complex affair. One has to control (at least) the following features: processor frequency, over-clocking via Turbo mode, the number of physical cores in use, the use of logical cores via Symmetric MultiThreading (SMT), the cache sizes available, the configured memory topology, as well as the power configuration if throughput per watt is to be measured. We have tried to do a good job of comparing like with like.

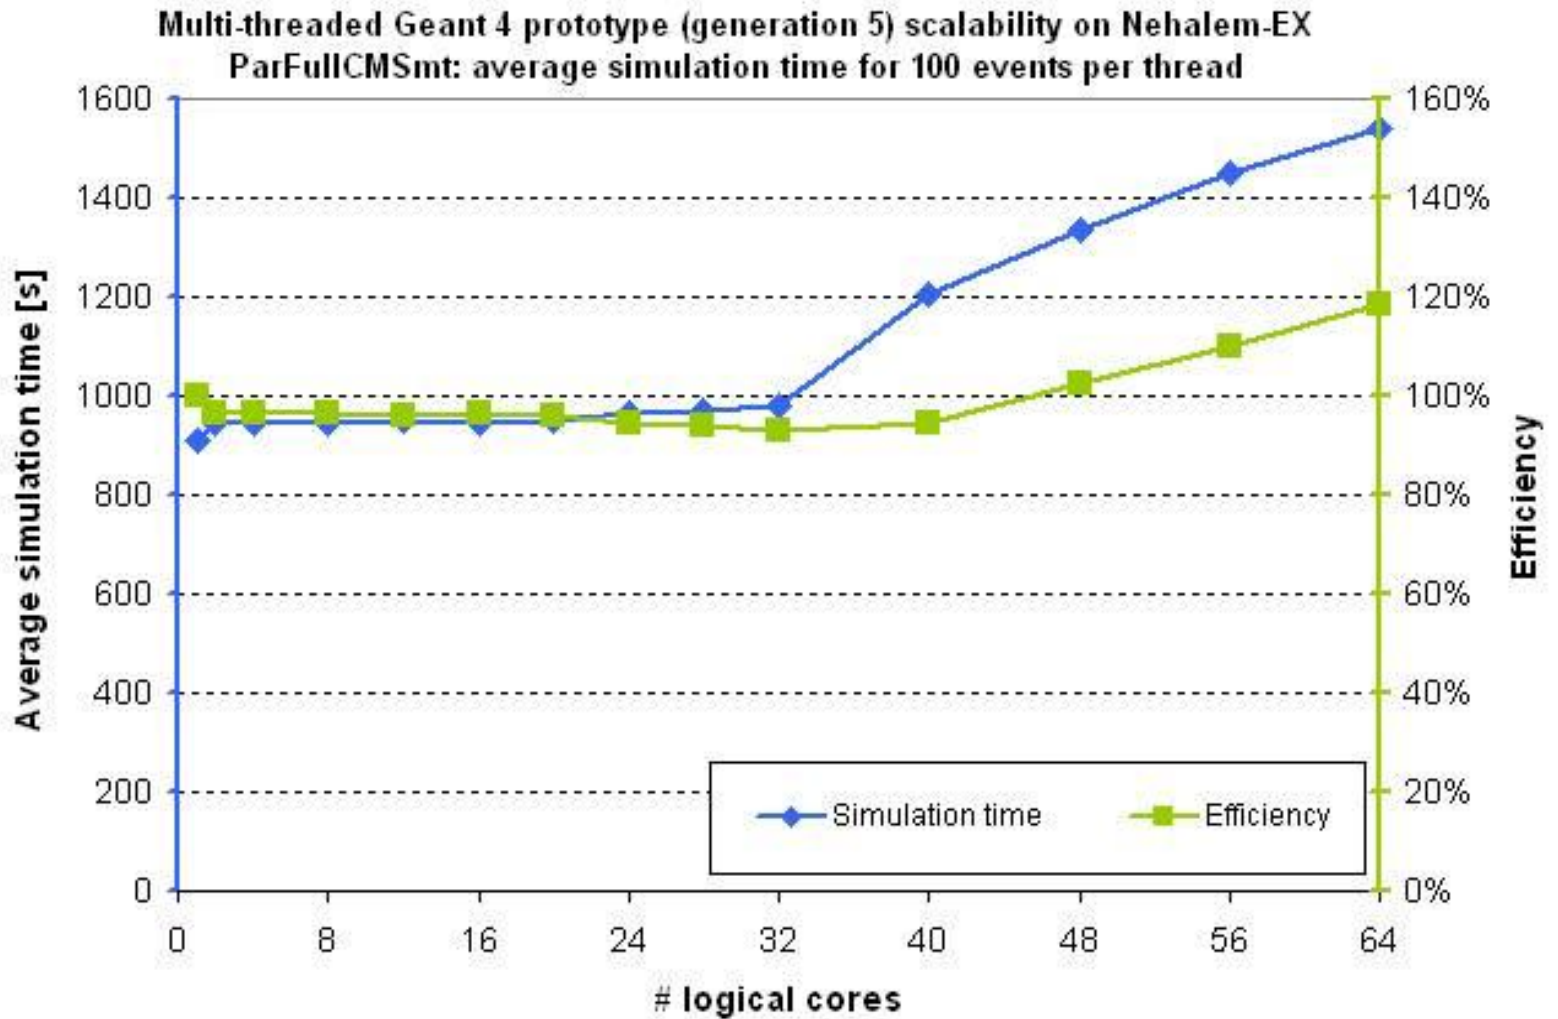
In summary, we saw a broad range of results. Our variant of the SPEC benchmark rate, "HEPSPEC", gave a stunning 3x overall improvement on the new server, thanks to good scaling with the 32 cores and a 26% additional gain when enabling SMT. In-house data analysis and simulation benchmarks showed throughput increases in the range of 11 to 60%. Oracle database tests will follow. Finally it should be mentioned

## • Early production level Nehalem CPU (X7560)

- Expandable version of the 5500 series CPUs
- 8 cores (16 threads)
- Turbo

## • 4 socket system counting 32 cores (64 threads) and 128GB of RAM

## • The final version of the paper was ready the day when the Nehalem-EX was launched




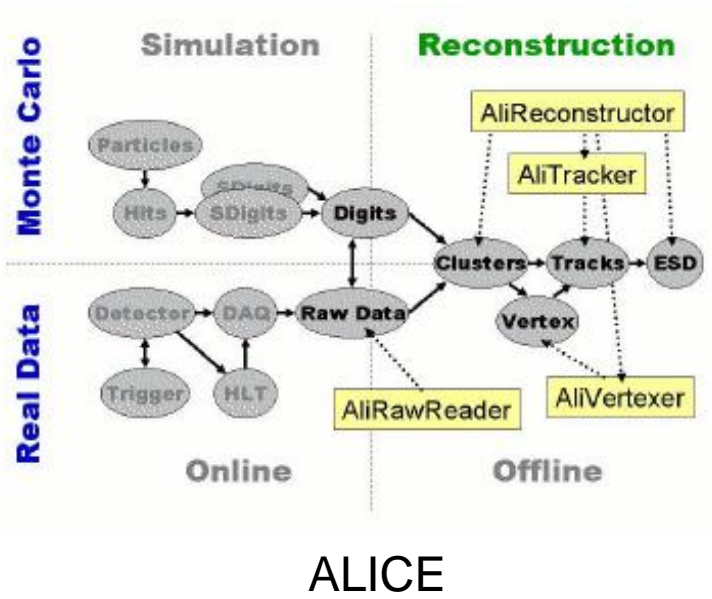
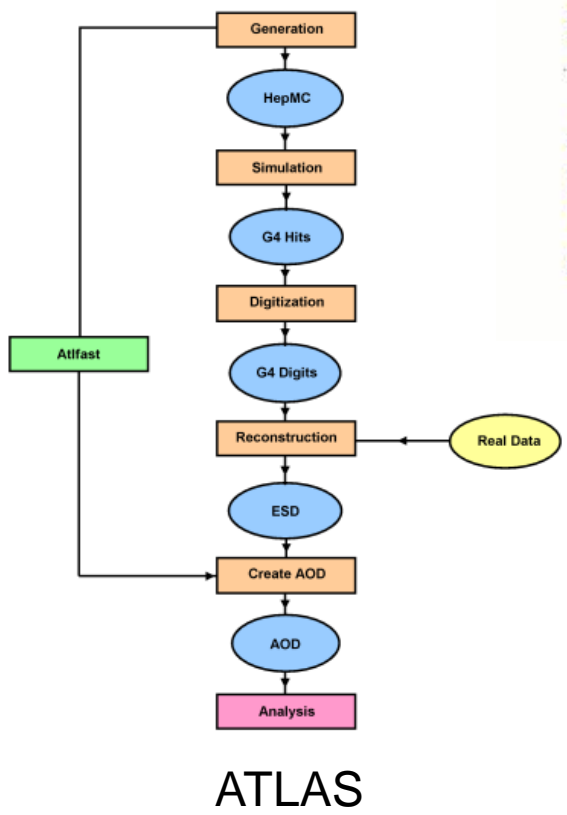
- Westmere, Nehalem-EX papers required each to run benchmarks during several CPU days
- Often repeated on early production systems
  - Hardware and firmware upgrades, bugs

## Oplabench framework would lighten our load for those tasks

- Ambitious project developed in collaboration with India
  - Inaugural project with few resources
- Imon Banerjee Master project
  - 1.5 months at CERN (mid July-mid August)
    - Attended CSC
  - Works on this project until the end of her Master thesis (end April 2011)

### Goals of the project:

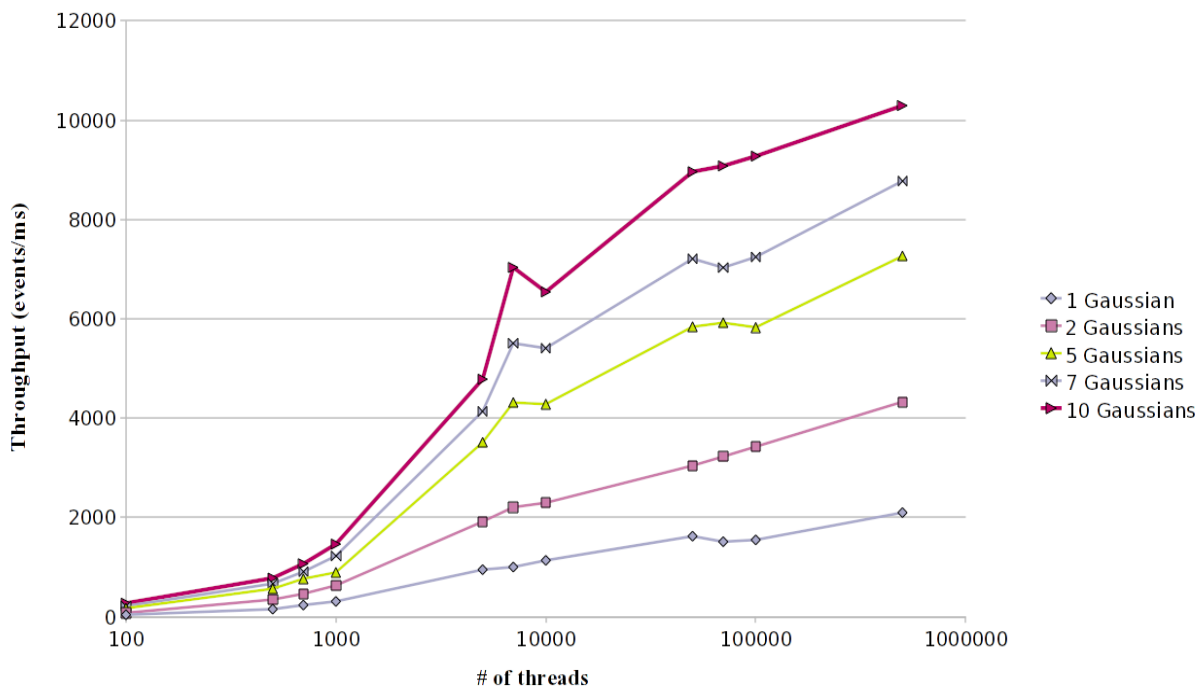
- Evaluate the opportunity to develop interesting projects in this context
- Ensure reproducibility of benchmarks, recording the parameters in a database
  - Hardware configuration: CPU, memory, harddrive
  - Firmware: mainly BIOS version and configuration
  - OS: OS release, software installed from OS repositories
  - Software environment required for the benchmark
  - Benchmark code
- Could store the results to automatically generate graphs, csv files...
- Could work on par with some standard local tools
  - For SLC5/RH based linux distributions installation: AIMS for OS installation at CERN, Kickstart for other sites



- Extraction, maintenance and performance monitoring of the latest versions of the 4 major software experiment frameworks at CERN
- Now ready for regular use in openlab
  - Trials in November

# Accelerators: CUDA inside ROOT/RooFit

Felice Pantaleo Summer Student





GPUs show impressive performance on some HEP-like codes, when lots of parallelism is available

- Data communication penalized on the PCIe bus
- Exact comparison with CPU impossible with CUDA, OpenCL implementation would allow to run the same code on GPU and CPU

Paper submitted to IPDPS 2011: "Evaluation of likelihood functions for data analysis on Graphics Processing Units"

- Very interesting ISA and architecture
- Main work on three pieces of software:
  - ALICE Trackfitter (online)
    - Vectorized, threaded (pthread, OpenMP)
  - Multithreaded Geant4 prototype (offline simulation)
    - Fully threaded (pthread), hard to vectorize
  - ROOT analysis being ported to FreeBSD





**CERN**  
openlab




- CERN openlab is evaluating top of the range Atom processors
  - Dual core
  - Hyperthreading
  - 64bit
- Recent test: Atom Pineview D510 @ 1.66GHz
- But no significant changes measured from the initial technical paper on N330
- Need to improve power measurement procedures
  - two identical systems dedicated to Atom benchmarking

Would allow for an accurate comparison of the 2 latest releases of Atom CPUs



- **icc**
  - 12.0 beta evaluation
    - Main focus intel64
    - Includes better vector support (AVX)
    - Several bugs and regressions reported
- **Performance comparisons with gcc 4.5.1**
  - This version includes -flto (similar to -ipo)
- **Good FP seminar at CERN by Jeff Arnold (May)**



- Visitors from Intel

- HPC: Tim Mattson
- Performance: David Levinthal
- Compilers and techniques: Jeff Arnold, Kathleen Knobe
- PSR: Tom Garrisson, Balint Fleischer, Willem Wery
- Tech: Herbert Cornelius, Mario Deilmann, Klaus-Dieter Örtel, Andrey Semin, Christopher Dahnken, Levent Akyil
- Mgmt/support: Stephan Gillich, Claudio Bellini, Jean-Faust Mukumbi
- Power and thermal: Mike Patterson

- Visits/trips by openlab

- ACAT (Jaipur)
- ISC (Hamburg)
- IDF (San Francisco)
- ERIC (Braunschweig)

- Expecting Sandy Bridge EP
- Completing icc beta
  - Expecting 12.0 production release (November)
- KNF upgrade
- Intel/openlab expert workshop: ArBB, Amplifier XE (November)
- 4 papers to be presented at CHEP