



Platform Competence Center Report

CERN openlab Major Review Feb 2014

Paweł Szostek, CERN openlab

On behalf of G.Bitzes, S.Jarp, P.Karpinski,
A.Nowak, A.Santogidis, P.Szostek, L. Valsan



CERNopenlab

- › **Manpower status**
- › **ICE-DIP**
- › **Research and technical activities**
- › **Teaching and workshops**
- › **Conferences, visits, interactions**
- › **Future plans**


› Arrivals

- Aram Santogidis – ICE-DIP Fellow (Oct)
- Przemyslaw Karpinski – ICE-DIP Fellow (Nov)

› Departures

- Przemyslaw Karpinski – studies at NUIM in Ireland and 18-month secondment at the Intel Exascale Labs in Paris

- › **5 Early Stage Researchers recruited and hired at CERN**
 - 1 in PH-ESE
 - 1 in PH-LBC
 - 1 in PH-ADT
 - 2 in IT-CF (openlab-PCC)
- › **2/5 completed University courses**
- › **4/5 went to ISOTDAQ 2014**
- › **Blog started**
- › **Secondments starting this season**
 - Santa Clara, Munich, Gdansk, Swindon, Paris
- › **Tight budget & visa issues**
- › **Looking forward to H2020**



Home

Blog

2014
06/01

Post #1: Intel Xeon Phi: OpenMP Stats

by Przemyslaw Karpinski at 03:42 PM in [ICE-DIP Intel Xeon PHI blog](#)

This is the first blog entry regarding Intel Xeon Phi and OpenMP programming. I hope you will find it interesting!

Intro

When I was reading OpenMP specification, I've noticed that there are some statistics which can be extracted programmatically. You can find proper descriptions for all of them in the "OpenMP Application Program Interface, Chapter 3: Runtime Library Routines".



Pawel Szostek - CERN openlab

Research and technical activities (1)

- › **Benchmarks - we fully revised and updated our benchmark base, in collaboration with PH**
- › **Reduced-size, efficient MLFit (w/ Vincenzo Innocente)**
- › **Multi-threaded Geant4 (production-like variant – w/ Andrea Dotti)**
 - Production version released and reported on at the **Concurrency Forum**
- › **Track Fitter (w/ Klaus-Dieter Oertel – Intel)**

Research and technical activities (2)

› **Compilers and tools**

- Intel tools now stable after repeated license problems

› **Performance tuning**

- Work on a custom profiler continues
- Further performance tuning packages made available (Linux kernel 3.11 + addons – unexplained 30% improvement on ATLAS code)
- Follow-ups concerning region based monitoring (after the GSoC project)

› **Intel Xeon Phi benchmarking – STREAM and micro-benchmarks**

Research and technical activities (3)

> Python code for ELF file analysis

- Pythonic **fast** address to source line mapping
- Pythonic symbol to offset mapping
- Static assembly code analysis
- Python bindings for Intel's XED + pyelftools integration + extra logic

```
import xed, elf, disass
e = elf.ELFFile("./test_elf")
hw_text = e.get_symbol_text('_ZN6C14helloWorldEv')
hw_insts = disass.disassemble_x64(hw_text)
bb = disass.get_basic_block(e, 0x81c
```

> Very rich experience with SWIG



Research and technical activities (4)

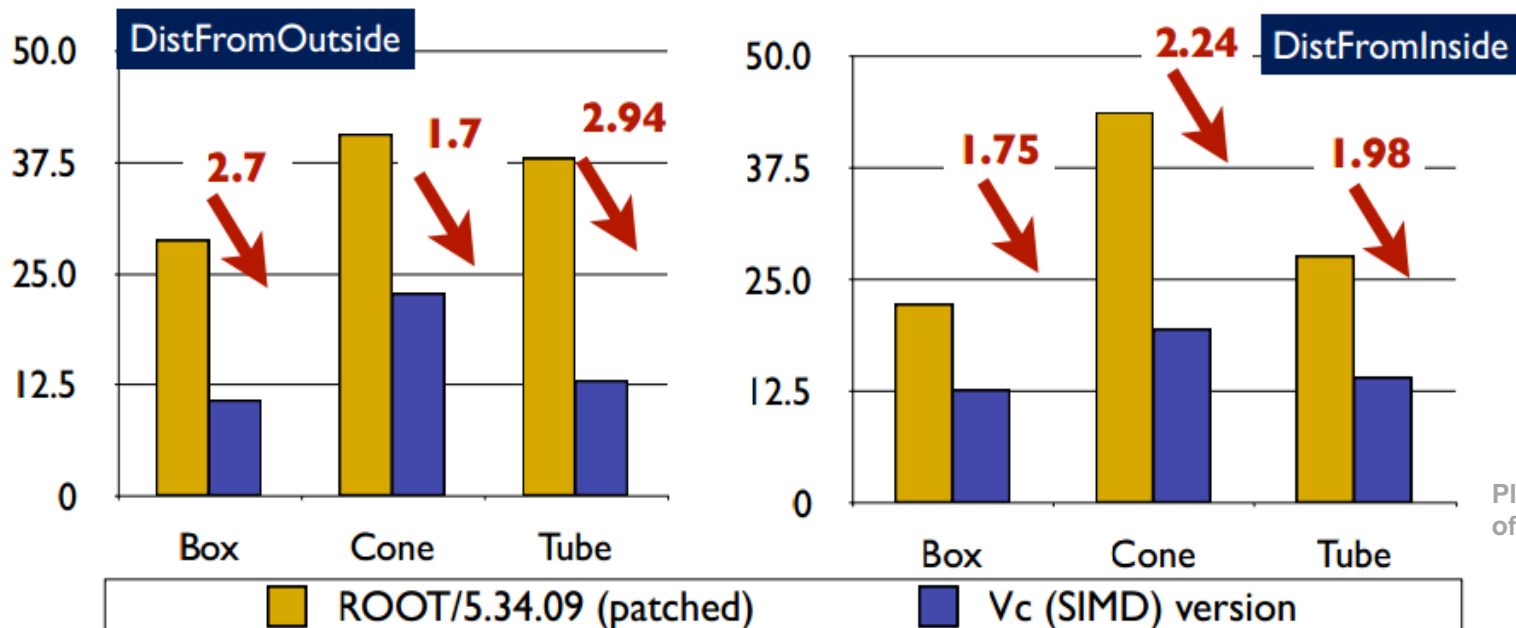
› **New Ivy Bridge-EX 4 Socket machine benchmarking**

- “supercomputer” in 4U chassis
- Requires outside-the-box approach
- Architectural details and results still under NDA

Research and technical activities (5)

> Geant-V collaboration

- Participation in the Geant-V project – aim is to build a high-performance simulation toolkit, exploiting the capabilities of modern hardware
- Multi-threaded and vectorized
- Initial geometry navigation on a toy detector achieves 4x speedup compared to ROOT version



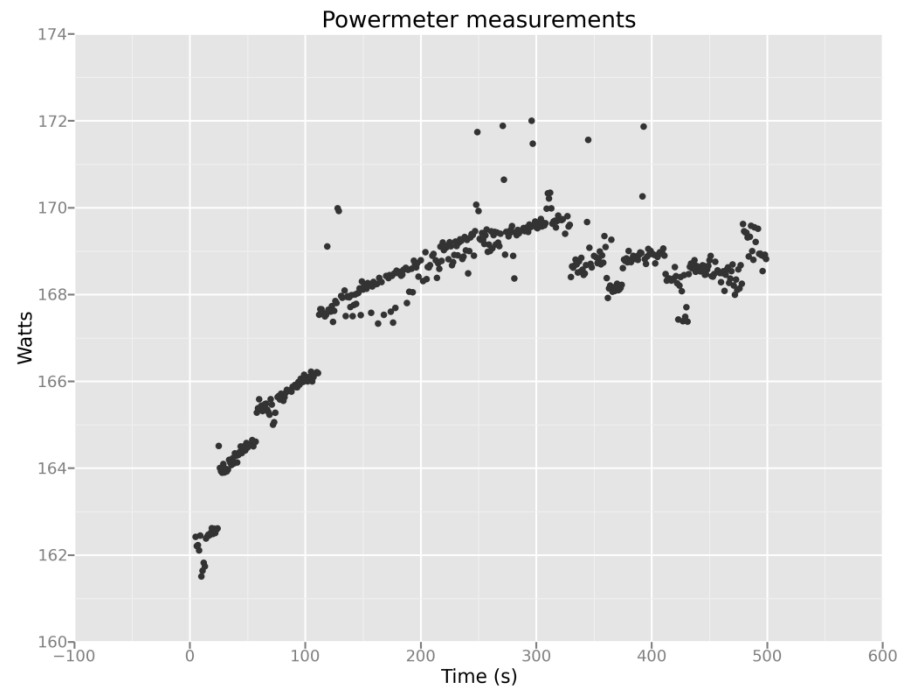
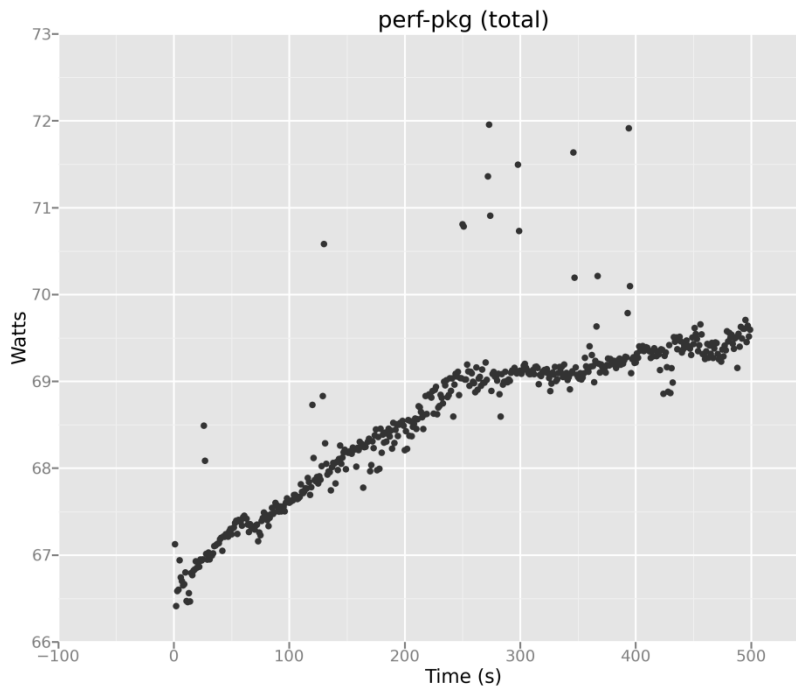
Plots by courtesy of Sandro Wenzel

Research and technical activities (6)

› Power consumption studies

- RAPL counters inside the CPU provide a way to measure its power consumption
- Comparing the total power consumption of a system and only the CPU's – what kind of relationship is there?
- How does the workload affect power consumption?
- Results are preliminary - work is ongoing

› Power consumption studies



Research and technical activities (8)

> Perf overhead studies

The overhead of profiling using PMU hardware counters [DRAFT]

February 2014

Authors:
Georgios Bitzes, Andrzej Nowak

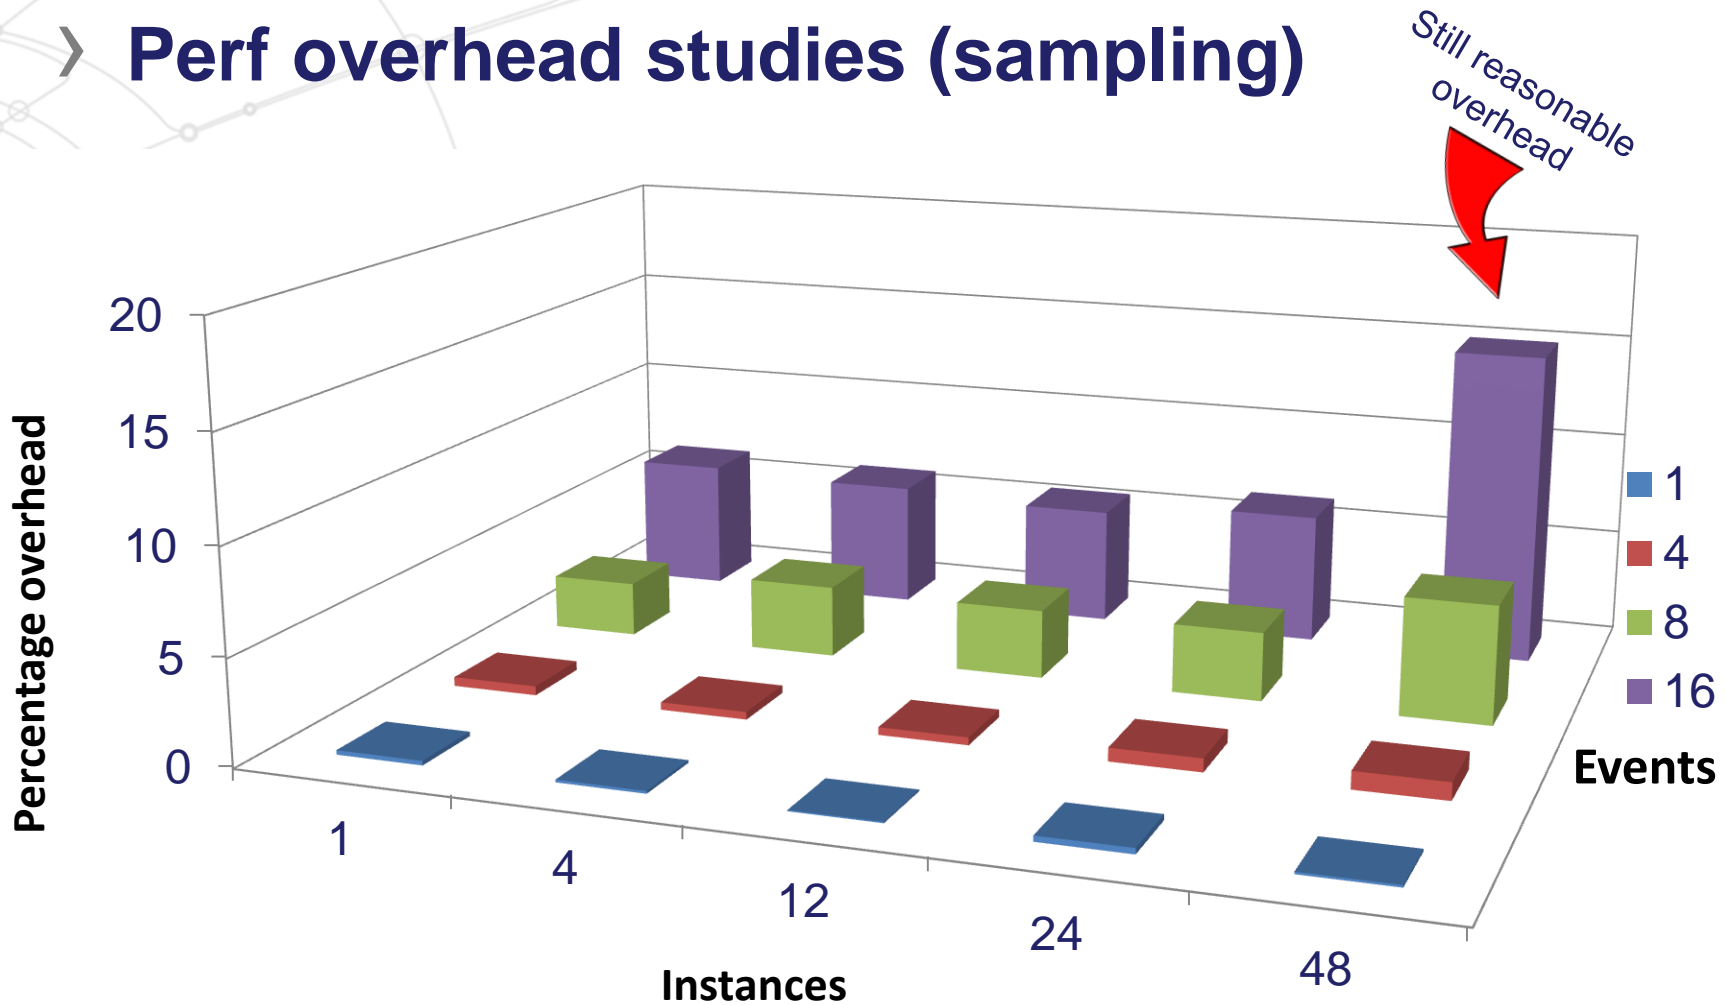
CERN openlab Report 2014



- Profiling methods: counting, sampling, PEBS, LBR
- Varying periods of sampling, number of events, number of cores
- Multiple benchmarks taken from SPEC, ROOT and Geant4
- **Very rarely the overhead goes over 25%, for counting mode principally negligible**
- **A report is ready**

Research and technical activities (9)


› Perf overhead studies (sampling)



Research and technical activities (10)

› On-going preparations for ARM64

- A system to appear in 1H2014
- Study of instruction set done
 - Vectors are similar in size to SSE
- Few details yet about micro-architecture (latency of math functions, etc.)
- Multiple compilers will be tested, in particular gcc and LLVM
- **Plan to compare to latest Atom (Avoton) servers**



Workshops and teaching (1)

- › **Parallelism, Compilers and Performance workshop – 6-8 November 2013**
 - ~35 participants
 - Speakers from CERN and Intel
- › **2nd CERN Advanced Performance Tuning Workshop – 20-22 November 2013**
 - ~50 participants
 - World-class speakers from ARM, Calxeda, CERN, Google, Intel, Versailles Exascale Computing lab
 - Spawned numerous follow-on activities throughout autumn and winter, as well as improvements in CERN code
 - <https://indico.cern.ch/conferenceDisplay.py?confId=280897>

Workshops and teaching (2)

- › Involvement in the iCSC 2014 (Feb 24-25)
- › 2nd Thematic CERN School of Computing coming up in June 2014
 - Limited, highly qualified audience
 - 1 week long, single topic focus

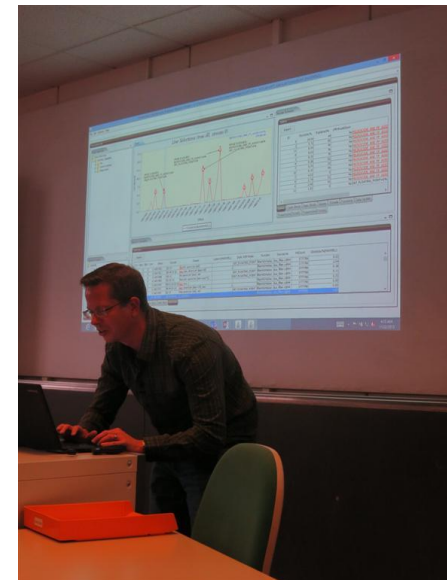


Conferences and talks

- › **ISC Big Data – 25-26 Sep 2014**
 - S. Jarp in Programme Chair
- › **CHEP'13, Amsterdam – 14-18 Oct 2013**
 - „Does the Intel Xeon Phi processor fit HEP workloads?” (A. Nowak)
 - „Beyond core count: a look at new mainstream computing platforms for HEP workloads” (P. Szostek)
- › **Atlas Software Week – 24 Oct 2013**
 - „An update on software for parallelism and heterogeneity” (A. Nowak)
- › **SC13, Denver – 17-22 Nov 2013**
 - S. Jarp attending, good sessions with key people from Intel
- › **ICE-DIP lectures (Jan 2014)**
 - „Introduction to trigger and data acquisition” (N. Neufeld / LHCb)
 - „DAQ implementation and practice” (P. Vande Vyvre / ALICE)
 - „Offline computing at CERN” (A. Nowak)

Visits and interaction

- › **Jeff Arnold, Intel**
- › **Stanislav Bratanov, Intel**
- › **Andres S. Charif-Rubial, Versailles Exascale Computing lab**
- › **Michael Chynoweth, Intel**
- › **Maria Dimakopoulou, Google**
- › **Stephane Eranian, Google**
- › **Al Grant, ARM**
- › **Hans Pabst, Intel**
- › **Robert Richter, Calxeda**
- › **Marie-Christine Sawley, Intel**
- › **Avinash Sodani, Intel**
- › **Karl Solsenbach, Intel**
- › **Mike Williams, ARM**
- › **Ahmad Yasin, Intel**
- › **Georg Zitzlsberger, Intel**
- › **Major Polish media visit**



Future work – next three to six months

- › **Looking forward to 2S „Haswell” systems and „Avoton” servers**
- › **Workshop and training period**
 - Two workshops in planning
 - Two schools in planning
- › **The ICE-DIP technical program moves to Intel**
- › **Releasing a profiling tool**
- › **Additional projects with Intel**
- › **Planning openlab V**
 - What are the future challenges? Will projects be predominantly mono-partner or multi-partner?
 - Contracts for all of us end around H2 2014

Q & A

Further questions?
pawel.szostek@cern.ch