Numerical Computing Workshop - Introduction

Andrzej Nowak, CERN openlab

May 27th 2013

with material from Sverre Jarp and the speakers

CERN openlab / Intel Numerical Computing Workshop Spring 2013

About CERN openlab

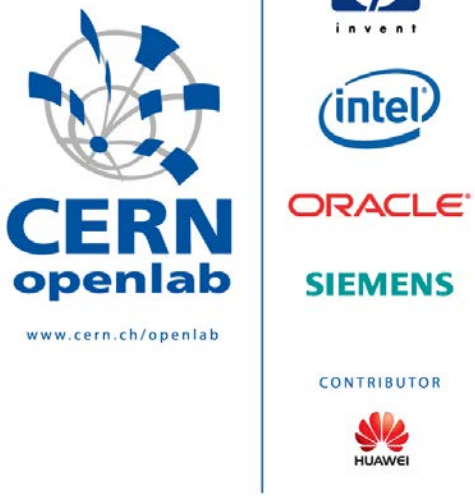

- **CERN openlab is a unique research partnership of CERN and the industry**
- **Objective: The advancement of cutting-edge computing solutions to be used by the worldwide LHC community**
  - Partners support manpower and equipment in dedicated competence centers
  - openlab delivers published research and evaluations based on partners' solutions – in a very challenging setting
- **Created robust hands-on training program in various computing topics, including international computing schools; Summer Student program**
- **http://cern.ch/openlab**

Why have this workshop?

- **In a constantly changing world:**
  - Hardware/software
  - Vectorization/Parallelization
  - New instructions (example: FMA)
  - New micro-architectures
    - Example: Throughput of VSQRTPD
  - Also: evolving standards
- **Highly qualified speakers on the following topics:**
  - IEEE754 floating-point format
  - Compilers
  - Math libraries
  - Developments in research

Previous instances of this class

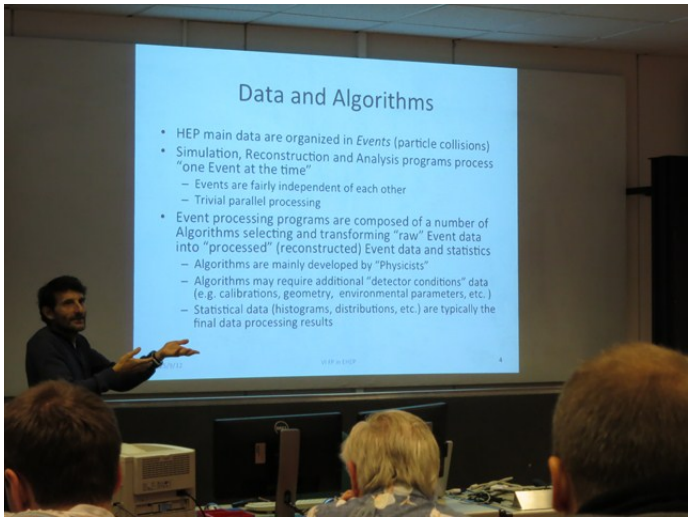

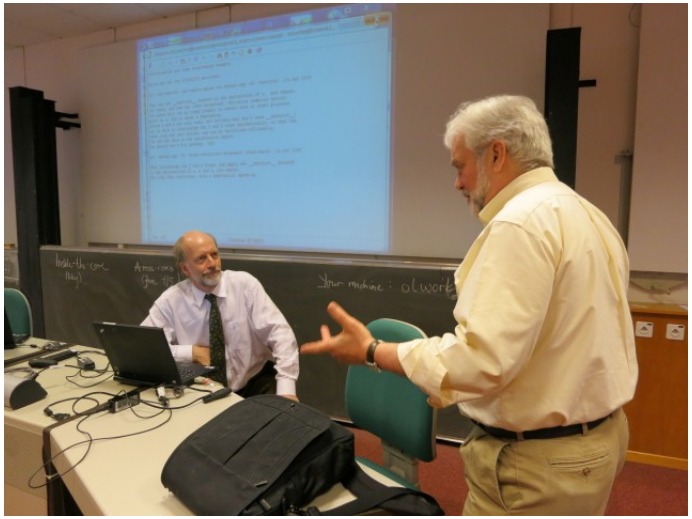

Workshop outline

> Practice

- Floating point in experimental data processing and in data analysis

> Understanding floating point

- Properties and algorithms
- Control in compilers and libraries

> Hands-on

- Lectures interleaved with exercises

> State of the art

- Target accuracy in HEP software
- Living with heterogeneity and non-CPU systems

Practice

- **How is floating point used and abused in HEP data processing and in data analysis?**
- **Challenges, opportunities, limitations**
- **Technical challenges:**
  - Libraries, compilers
  - Old algorithms for x87
  - Bugs
  - Vectorization

Understanding floating point

> IEEE754

- X86-64: Binary representation
  - Binary32, binary64
- Not all real numbers can be correctly represented
- The issues: Rounding and Reproducibility
  - Rounding modes and error control
- Cannot always trust what you know from your high-school math

> Jeff Arnold will talk extensively about this topic

- Which will also be supplemented by a set of exercises

Compilers

- **Compilers have a great influence on floating-point calculations**
  - Long list of compiler options which influence the calculations
    - Both “safe” and “unsafe”
  - Unfortunately, there is no standard for defining options

- **Martyn Corden will cover the Intel compiler and libraries and also compare to GNU**
  - Including a set of exercises

Math libraries

- **Equally vital for FLP calculations (at least at CERN), but the issues remain:**
  - Accuracy versus speed
  - Argument ranges
  - Rounding options
  - Precision: Average and maximum error

- **Bit-accurate libraries exist, but often a limited number of digits are sufficient.**

- **Danilo Piparo will cover HEP efforts**
- **Florent de Dinechin will cover this topic in great detail**

Heterogeneity and non-CPU techniques

- **Getting good floating point consistency on GPUs**
  - David Rohr
- **Numerical processing in FPGAs**
  - Florent de Dinechin

```
cpu_wrapper.cpp:
#define __DECL void
#include ``common.cpp``

void FitTracks() {
  for (int i = 0;i < nTr;i++) {
    FitTrack(n);
  }
}
```

```
gpu_wrapper.cpp:
#define __DECL __device void
#include ``common.cpp``

__kernel void FitTracksGPU() {
  FitTrack(threadIdx.x);
}

void FitTracks() {
  FitTracksGPU<<<nTr>>>();
}
```

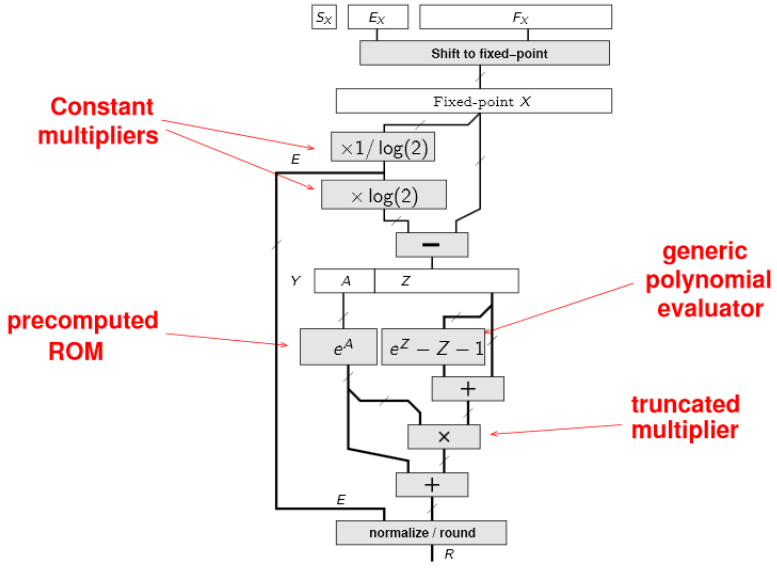

Conclusions

- This workshop will help in full understanding and control of floating point
- Issues
- Algorithms
- Compilers
- Libraries
- Heterogeneity
- “Make sure you are in the driver’s seat”

CERN
openlab