

Table of Contents

Multi-Core R&D Project.....	1
General.....	1
Latest Additions.....	1
"Recent" Summary Reports.....	1
"Recent" Event of interest.....	1
Track 1:.....	1
Objective.....	2
Deliverables.....	2
Current/Recent Activities.....	2
Previous works.....	2
Track 2:.....	2
Objective.....	2
Deliverables.....	3
Current/Recent Activities.....	3
Track 3:.....	3
Objective.....	3
Deliverables.....	3
Current/Recent Activities.....	3
Track 4:.....	3
Deliverables.....	4
Current/Recent Activities.....	4

Multi-Core R&D Project

Optimize LHC physics software applications to run on multi-core cpus

This project is part of the CERN SFT [group](#) in the PH Division and conducted in collaboration with the OpenLab [and](#) the LHC experiments

The project ended in 2011. A new project Concurrent Programming Models and Frameworks [has](#) been initiated that will take-over the activities carried during this R&D.

General

- Meetings [group](#)
- mailing-list archive [group](#)
- Action Items
- Blogs
- In case of further questions

Latest Additions

- Reducing memory footprint using jemalloc
- Reconciling OOD with DOD
- Auto-vectorize trigonometric and trascendental functions

"Recent" Summary Reports

- HEP software on Multicore: present experience and perspectives: given at the SuperB Computing R&D Workshop 2011 [group](#)
- collection of presentations and reports on the activities in year 2010 [group](#)
 - ◆ Please refer to these reports for a detailed description of the activities listed in the sections below
- Report (pdf) [group](#) at the 2nd Workshop on adapting applications and computing services to multi-core and virtualization (2010)
- Report from the Workshop on "CHALLENGES FOR UNDERSTANDING THE QUANTUM UNIVERSE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE" [group](#)
- Report (pdf) [group](#) at first Workshop on adapting applications and computing services to multi-core and virtualization (2009)
- Report (pdf) [group](#) at the Kickoff Workshop on multi-core and virtualization (2008)

"Recent" Event of interest

- Workshop on Concurrency in the many-Cores Era, November 2011 [group](#)
- ATLAS workshop on Future computing in particle physics: June 2011 [group](#)
- An introduction to OpenCL: tutorial by Tim Mattson (Intel Labs) [group](#)
- Workshop on HPC on Hybrid platforms at INRIA [group](#)
- SuperB Computing R&D Workshop 2011 [group](#)

Track 1:

Objective

- Investigate current and future multi-core architectures.
- Evaluate tools to measure performance.
- Develop a measurement and analysis methodology.

Deliverables

- Assessment of industry trend in multi-core architectures.
 - ◆ see reports above
 - ◆ *google* for Nehalem, SandyBridge, Bulldozer,, NuMa, multicore, manycore, gpgpu
- Recommendations on tools, metrics and methodology to assess the performance of LHC physics application software on such architectures
 - ◆ CMSSW Performance Monitoring on processors based on the Intel Core and Nehalem Microarchitectures by *Daniele Kruse*
 - ◆ see OpenLab reports and lectures at INFN school referenced below
- Instrumentation software
 - ◆ Perfmon2 by *OpenLab*
 - ◆ igprof [↗](#) by *Giulio Eulisse & Lassi Tuura*
 - ◆ Gaudi Profiler [↗](#)
 - ◆ CMSSW Profiler

Current/Recent Activities

- Investigate core i7 architecture and benchmark it with particular attention to hyperthreading (report by VI)
- See also OpenLab documents [↗](#)
- Establish a measurement methodology based on perfmon2
 - ◆ * CMSSW Performance Monitoring on processors based on the Intel Core and Nehalem Microarchitectures by *Daniele Kruse*
- Instrument LHC software (CMSSW, Gaudi, Geant4)
- Performance Analysis on Intel® Corei7 [↗](#) by David A Levinthal (Intel)
- ptuview [↗](#) by Matti Kortelainen

Previous works

- Analysing CMS software performance using IgProf, OProfile and callgrind [↗](#) by Lassi Tuura, Giulio Eulisse and VI
- The investigation of CMS software through an automated performance measurement and analysis system. by Danilo Piparo, Robin Moser and VI

Track 2:

Objective

- Measure and analyze performance of current LHC physics application software on multi-core architectures
- Identify bottlenecks
- Prototype solutions at the level of *system* and *core* libraries

Deliverables

- Reports on performance of current LHC physics application software
 - ◆ CMSSW Performance Monitoring on processors based on the Intel Core and Nehalem Microarchitectures by *Daniele Kruse*
 - ◆ OpenLab documents [↗](#)
- Recommendations on best practices to avoid bottlenecks and best exploit multi-core architectures
 - ◆ courses given at the INFN school on "Architectures, tools and methodologies for developing efficient large scientific computing applications" in 2009 [↗](#), 2010 [↗](#)
- Eventual materialization in software library components to implement them
 - ◆ ksm.tar.bz2: ksm module for RHEL 5.2 (version march 2009)
 - ◆ modified version of the cephes library to support auto-vectorization in GCC

Current/Recent Activities

- Evaluation of tcmalloc and *huge-pages* report by VI
- Evaluation of KSM (report by VI, and consolidated evaluation on march 2009)
- Workshop with Intel
- Evaluation of jemalloc as a mean to reduce the memory footprint on multicore

Track 3:

Objective

- Investigate solutions to parallelize current LHC physics software at *application framework* level
- Identify reusable design patterns and implementation technologies to achieve parallelization
- produce prototypes

Deliverables

- Recommendations on reusable design patterns and implementation technologies to use to achieve parallelization:
 - ◆ Study on shared memory for the LHC using Boost Interprocess library [↗](#)
 - ◆ G4report-11-04.txt: Progress report on multithreaded Geant4
 - ◆ see also conference reports referenced in the "activity summary" above
- Eventual materialization in software library components to implement them
 - ◆ Gaudi parallel [↗](#)

Current/Recent Activities

- parallelization of Gaudi using python
- PROOF-lite
- parallelization of ATLAS event processing framework
- parallelization of CMS event processing framework

Track 4:

- Investigate solutions to parallelize algorithms used in current LHC physics application software
- Identify reusable design patterns and implementation technologies to achieve effective high granularity parallelization
- produce prototypes

Deliverables

- Recommendations on reusable design patterns and implementation technologies to use to achieve effective high granularity parallelization
 - ◆ Reconciling OOD with DOD
- Eventual materialization in software library components to implement them
 - ◆ Parallel Minuit2 (*google* for it)

Current/Recent Activities

- Vectorization in reconstruction Algorithms (by VI and CMS team)
- port of RootFit to GPU (by OpenLab)
- Evaluation of maximum likelihood fits on GPU devices using CUDA by Felice Pantaleo. Given at "Facing the Multicore Challenge II" workshop in Karlsruhe, September 2011 [↗](#)

-- VincenzoInnocente - 21 Oct 2007

-- MarcMagransDeAbril - 04 Jul 2008

This topic: LCG > MultiCoreRD

Topic revision: r35 - 2015-03-15 - VincenzoInnocente



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)