

Annual Concurrency Forum

Meeting Goals

Fermilab, 4-6th February 2013

P. Mato / CERN

Outcome of the 2011 Workshop

- ❖ The first *Workshop on Concurrency in the many-Cores Era* was held at Fermilab in November 2011 to explore the possibility that interested High Energy Physics (HEP) institutions and projects collaborate on concurrent frameworks and applications R&D
- ❖ The outcome of the first workshop
 - ❖ Interest for common effort to make rapid progress on exploratory R&D activities during 2012
 - ❖ Aiming to share a common concurrency programming model
 - ❖ Identified a number of '**demonstrators**' to exercise different capabilities in a small scale with clear deliverables and metrics, and work in short cycles a few months long
 - ❖ Setup regular meetings **The Concurrency Forum** and **Web Site**
- ❖ Concurrency Forum meeting goals
 - ❖ Understanding common requirements and constraints
 - ❖ Share knowledge and learn from each other
 - ❖ As a community, identify what tool / library / model works and what does not for our applications

Demonstrators

- ❖ Stimulated big enthusiasm in the community!
- ❖ During 2012, we have had presentations and discussions on most of the identified *demonstrators*
 - ❖ People have been very active and responsive
- ❖ Main Topics
 - ❖ Parallelism in detector simulation
 - ❖ Heterogeneous computing (GPU, MIC)
 - ❖ Memory and parallelism
 - ❖ Parallelism in algorithms and frameworks

<https://concurrency.web.cern.ch>

Demonstrator	Summary	Contact Person(s)
'Whiteboard' service	Design and implementation of a 'service' able to schedule algorithms/modules based on their data dependencies	B. Hegner, P. Mato, D.Piparo
ng-go-gaudi	Investigate use of Go as a programming language to expose and harness concurrency in HEP	S. Binet
CET multicore framework	Use of MPI and OpenMP parallelism on a DAQ project based on Darkside-50	M. Paterno, J. Kowalkowski
libdispatch investigation	Further investigate use of libdispatch (GCD)	C. Jones
Evaluation of frameworks	Define standard set of modules timings and configurations, data samples, and set of module dependencies. Define torture tests that will help determine performance limits.	C. Jones
Multithreaded I/O	An attempt to integrate ROOT I/O subsystem and libdispatch by scheduling the different operations (disk I/O, compression, object serialization) concurrently	C. Jones, Ph. Canal
OpenCL for Physics Applications	Find a programming model to exploit current and future hardware for compute intensive tasks	T. Hauth, D. Piparo, V.Innocente
Virtualization	Can virtualization help in improving performance in many core environment	P. Buncic
TBB investigation for SuperB	Implement a small system with the Intel Threading Building Blocks (TBB), which is a library offering a rich approach to expressing parallelism in a C++ program	F. Giacomini
Data Locality in G4	Try different data flow organizations within G4 and measure the actual effect on the use and leverage of CPU caches	Ph.Canal
Adapting G4MT	Adapting G4MT prototype to an external concurrency model to be used by the experiment 's frameworks	J. Apostolakis
High performance simulation (vector prototype)	Develop a simulation aimed at achieving maximum performance on parallel architectures integrating fast and full simulation in the same framework. Study parallelism at different granularity levels.	F. Carminati
Object sharing	Demonstrate sharing of C++ objects among processes	P. Calafiura
Likelihood parallelization	Parallelization of likelihood functions for data analysis	A. Lazzaro
Track Seeding parallelization	Parallelization of the seed building in the CMS reco software	T.Hauth, V.Innocente, D.Piparo
Real-time Trigger using GPUs	Implement low latency and high throughput algorithms for NA62 online trigger	F. Pantaleo, V.Innocente

Annual Meeting Goals

Discuss progress made with each *demonstrator*, drawing **conclusions** for those that have completed their work programme and **identifying topics that still need to be investigated.**

Get an update from each experiment or major project on its **current views on parallelism issues**, such as choice of concurrency model and software technologies

This should be an opportunity for us all to discuss and see if we can converge on a common overall strategy

Possibility of **launching common development projects** with specific deliverables that can be of benefit to the whole community

Sessions Organized by Topic

- ❖ 1 - Experiments and Projects *Pere Mato*
- ❖ 2 - Multi-threaded Frameworks *Jim Kowalkowski*
- ❖ 3 - Libraries and Toolkits for a Parallel world *Benedikt Hegner*
- ❖ 4 - Technology Watch and Evaluations *Marc Paterno*
- ❖ 5 - Vectorization *John Apostolakis*
- ❖ 6 - Co-processors and Accelerators *Philippe Canal*
- ❖ 7 - Memory Optimization Techniques *Jakob Blomer*

Each topic will be introduced by the session convener, who will also lead the discussion foreseen at the end of each session

Summary of each session including the discussion will be presented at the Wednesday morning session by the session convener

Wednesday Morning

- ❖ Workshop Summary

- ❖ Main conclusions from each of the sessions
- ❖ Proposals for additional 'demonstrators' to continue the R&D phase on certain topics not yet well covered
- ❖ Proposals for collaborative development projects
- ❖ Feedback on Concurrency Forum (bi-weekly meetings, annual meetings)

- ❖ Informal discussions

- ❖ Time for the interested people to get together and discuss the details of the proposals and possible initial program of work