Status of Computing Performance Monitoring

V. Daniel Elvira for the G4 Computing Performance Task team

Fermi National Accelerator Laboratory

The G4 Computing Performance Task

The G4CP task is part of the G4 Testing & QA group:

- Monitor time and memory performance through the development cycle
- Identify problems and opportunities for code improvement and optimization
- Communicate results and findings to the appropriate G4 group leaders and the Steering Board

Web page link through the Testing & QA's pages:

https://oink.fnal.gov/perfanalysis/g4p/admin/task.html

(Description of the benchmarking/profiling protocol and a link to a results page)

Team: Ph. Canal, D. Elvira, K. Genser, S. Jun (FNAL)

Profiling, Time/Memory Monitoring

Geant4 Benchmarking and Profiling

The Geant4 (G4) Computing Performance Task Force (G4CP Task Force) is part of the Geant4 testing and QA group. The charge of the G4CP Task Force is to monitor G4 software through its development cycle for expected and unexpected changes in computing performance, identify problems and opportunities for code improvement and optimization, and communicate the results and findings to the appropriate G4 working group leaders and the Steering Board.

We are suggesting the following Geant4 benchmarking/profiling protocol for profiling/timing of the official, reference and candidate releases:

Contents

1. Fast: FAST is a set of tools for collecting, managing, and analyzing data about code performance. Instructions for use of the FAST toolkit is available at FAST project page
2. IgProf: Ignominous Profiler is a simple tool for measuring and analysing application memory and performance characteristics. For more information, see IgProf home page
3. HPC Performance Tools: HPCToolkits OpenISpeedShop TAU

1. SimplifiedCalo (A. Dotti),
2. cmsExp (Standalone G4 with CMS GDML geometry and realistic B field by S. Jun)

Profiling, Time/Memory Monitoring

Profiled all G4.9.6 versions and every reference release towards G4.10

*) the first version starting the multithread (MT) migration
**) the reference tag 9.6.r06 corresponds to the published 10.0-beta
***) switching the cxx flag from O2 -DNDEBUG to -O2 -g -fno-omit-frame-pointer

2) Profiling Results

Geant4 Version	Application	Performance		Summary	
9.6.r08	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r07O2g***	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r07	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r06**	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r05	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r04	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r03	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r02*	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.r01	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.p02	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6.p01	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6 (sl6)	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM
9.6 (sl5)	SimplifiedCalo	Simple Profiler	Memory Profiler	CPU	MEM

Profiling:

- Call path, function call, library call
- Graphical display of call path
- Info in file for stat analysis

Benchmarking:

- CPU time/event
- Total mem/run

Profiling, Time/Memory Monitoring

CPU Time in seconds/event

Geant4.9.6.r08 SimplifiedCalo 01

Sample	Physics List	B-Field	Energy	CPU/event
Higgs->ZZ	FTFP_BERT	4.0T	14 TeV	380.9300

Sample	Physics List	B-Field	Beam Energy			
			1 GeV	5 GeV	10 GeV	50 GeV
Electron	FTFP_BERT	0.0T	0.0185	0.0932	0.1861	0.9178
		4.0T	0.0238	0.1198	0.2394	1.1832
Pion-	FTFP_BERT	0.0T	0.0335	0.1468	0.2808	1.2598
		4.0T	0.0373	0.1663	0.3203	1.4548
	QGSP_BERT	4.0T	0.0374	0.1631	0.3179	1.4052
	QGSP_BIC	4.0T	0.0325	0.1532	0.2955	1.3511
	LHEP	4.0T	N/A	N/A	N/A	N/A
Proton	FTFP_BERT	4.0T	0.0287	0.1705	0.3299	1.5219
Anti-Anti_Proton	FTFP_BERT	4.0T	0.0873	0.2272	0.3779	1.5555

Processor: AMD Opteron(tm) Processor 6128 (CPU: 2000 MHz, Cache: 512 KB)
Total Memory: 66007532 kB

Time & Memory monitoring on varied selection of input samples: particle type, energy, physics lists

Profiling, Time/Memory Monitoring

Total Memory in Counts/1,000,000 (Mbytes)

Geant4.9.6.r08 SimplifiedCalo 01

Sample	Physics List	B-Field	Energy	After First Event	After Last Event
Higgs->ZZ	FTFP_BERT	4.0T	14 TeV	355.072	18280.1

Sample	Physics List	B-Field	Beam Energy/Process							
			1 GeV		5 GeV		10 GeV		50 GeV	
			First Event	Last Event	First Event	Last Event	First Event	Last Event	First Event	Last Event
Electron	FTFP_BERT	0.0T	57.2087	59.2854	57.3669	62.9564	57.756	67.2786	57.8991	104.875
		4.0T	57.2701	59.3826	57.4701	63.424	57.8767	67.8267	57.8572	105.003
Pion-	FTFP_BERT	0.0T	57.8302	137.114	58.3068	375.199	58.7227	597.012	62.1217	2049.04
		4.0T	57.7951	136.59	58.7208	377.699	58.7349	605.811	61.8229	2018.9
	QGSP_BERT	4.0T	58.6727	137.468	59.2609	354.357	59.6125	599.601	85.6904	1866.82
	QGSP_BIC	4.0T	130.197	871.129	156.66	3077.34	135.03	5259.16	188.437	18611
	LHEP	4.0T	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Proton	FTFP_BERT	4.0T	57.6862	120.622	58.5796	412.722	58.802	678.727	62.7171	2478.43
Anti-Anti_Proton	FTFP_BERT	4.0T	57.9079	211.732	59.0813	480.648	58.4718	725.861	61.4891	2393.44

Processor: AMD Opteron(tm) Processor 6128 (CPU: 2000 MHz, Cache: 512 KB)
Total Memory: 66007532 kB

Time & Memory monitoring on varied selection of input samples: particle type, energy, physics lists

New observables added to protocol

Based on input from members of the Collaboration:

- Number of tracks and steps to determine if changes come from physics or geometry (requested by A. Dotti)
- physics lists for major releases (requested by A. Ribon)
- Initialization time (requested by M. Asai)
- More memory monitoring stats (vsize, RSS, shared)

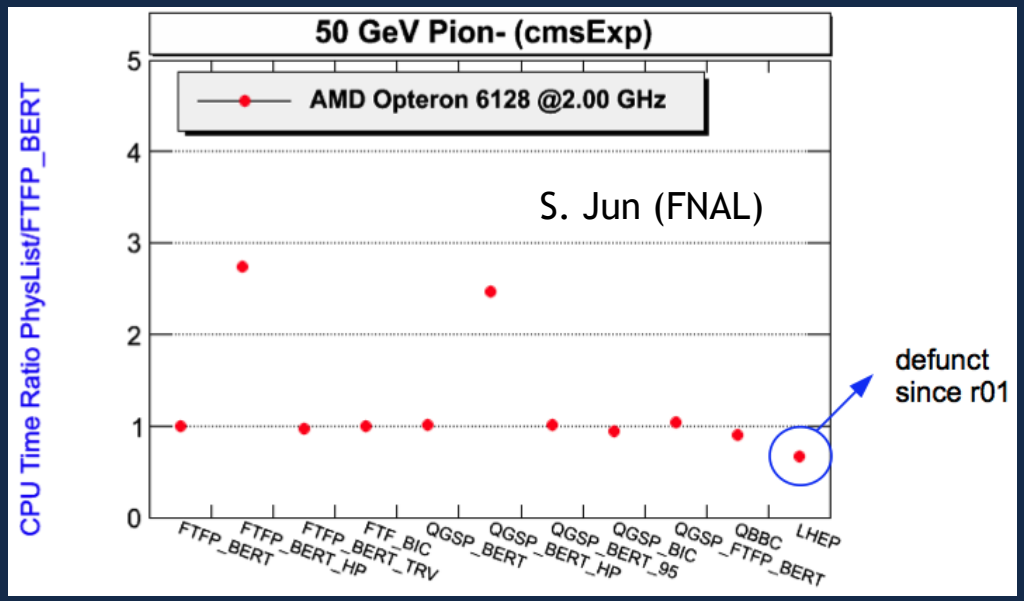

Profiling, Time/Memory Monitoring

New G4CP protocol to test MT capability through the G4.10 release process was implemented as of April 2013:

- Sequential Mode Monitoring (cmake flag off – MT part not compiled)
- Multithread Mode Monitoring (MT flag on – single thread test) to evaluate the MT penalty
- Scalability test. Throughput (#evts/sec versus # of threads) to verify linearity)

6) Geant4 MT Performance

Geant4 Version	Application	Performance		
9.6.r08	cmsExpMT	Summary	OpenISpeedShop	HPCToolkit
9.6.r07	cmsExpMT	Summary	OpenISpeedShop	HPCToolkit
9.6.r06	cmsExpMT	Summary	OpenISpeedShop	HPCToolkit
9.5.p01mt	cmsExpMT	Summary	OpenISpeedShop	HPCToolkit

Profiling & Monitoring Results

G4.9.6.rX (single thread mode)/G4.9.6 ratios

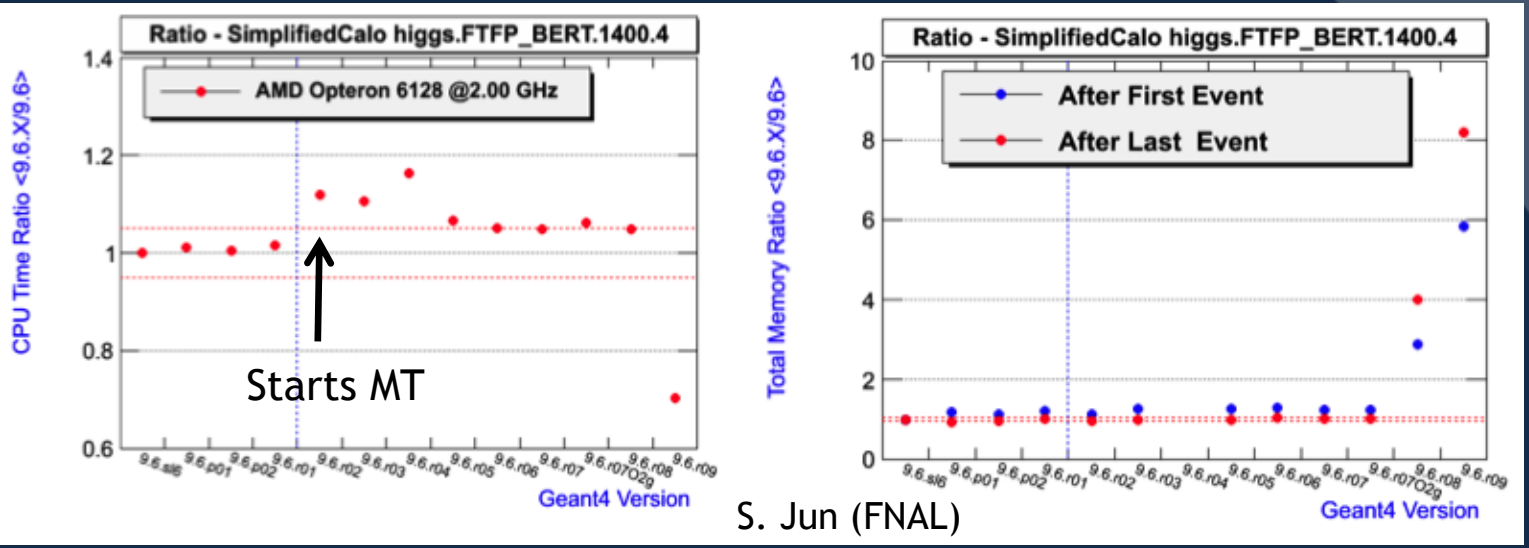

S. Jun (FNAL)

Performance penalty of 5% in G4.9.6.r>2 (single thread mode) mostly due to overhead in hadronic physics code

S. Jun is researching and testing open source performance tools for multithreaded applications: HPCToolKit, Open SpeedShop, TAU

Profiling & Monitoring Results

G4MT scaling tests: throughput versus # of cores

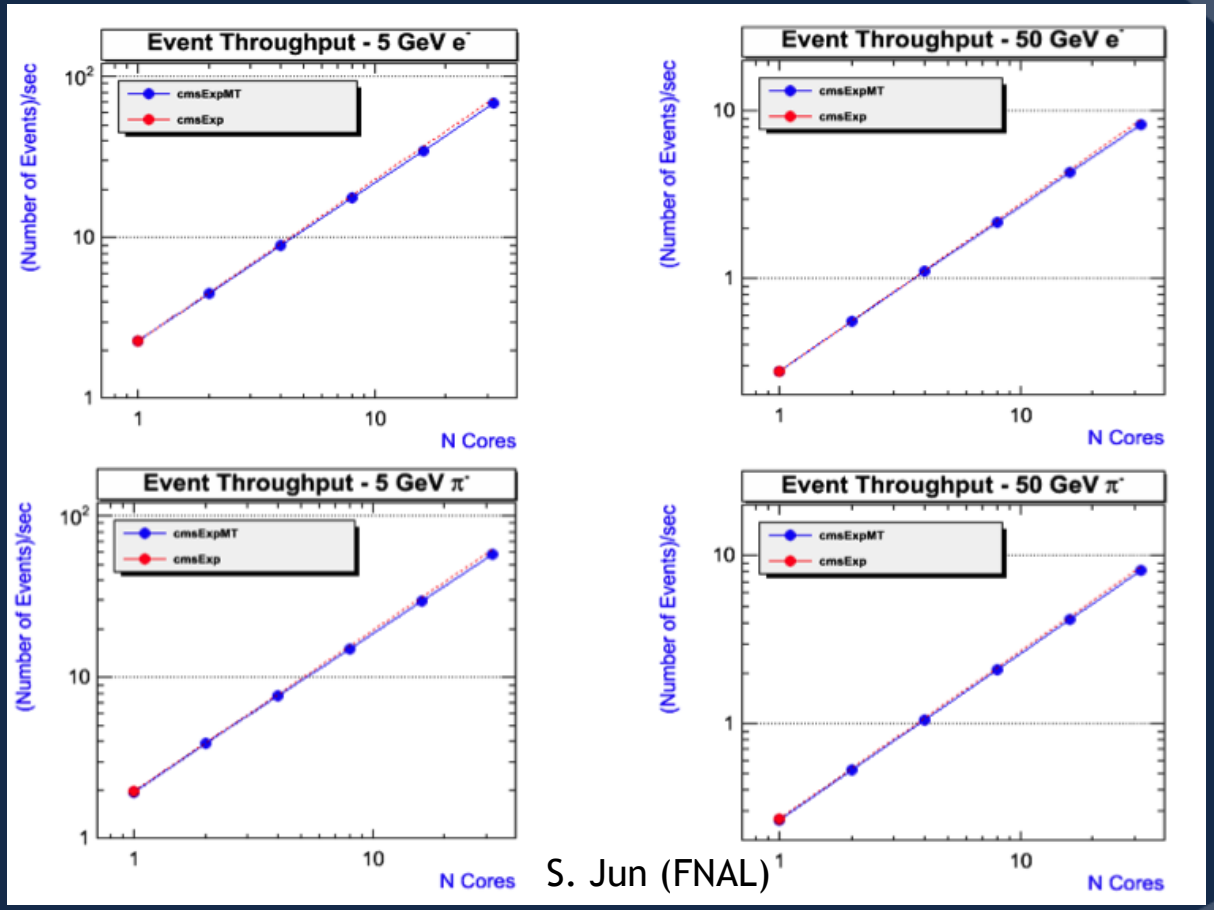

S. Jun (FNAL)

Profiling & Monitoring Results

G4MT efficiency tests: CPU time/event/core versus # of cores

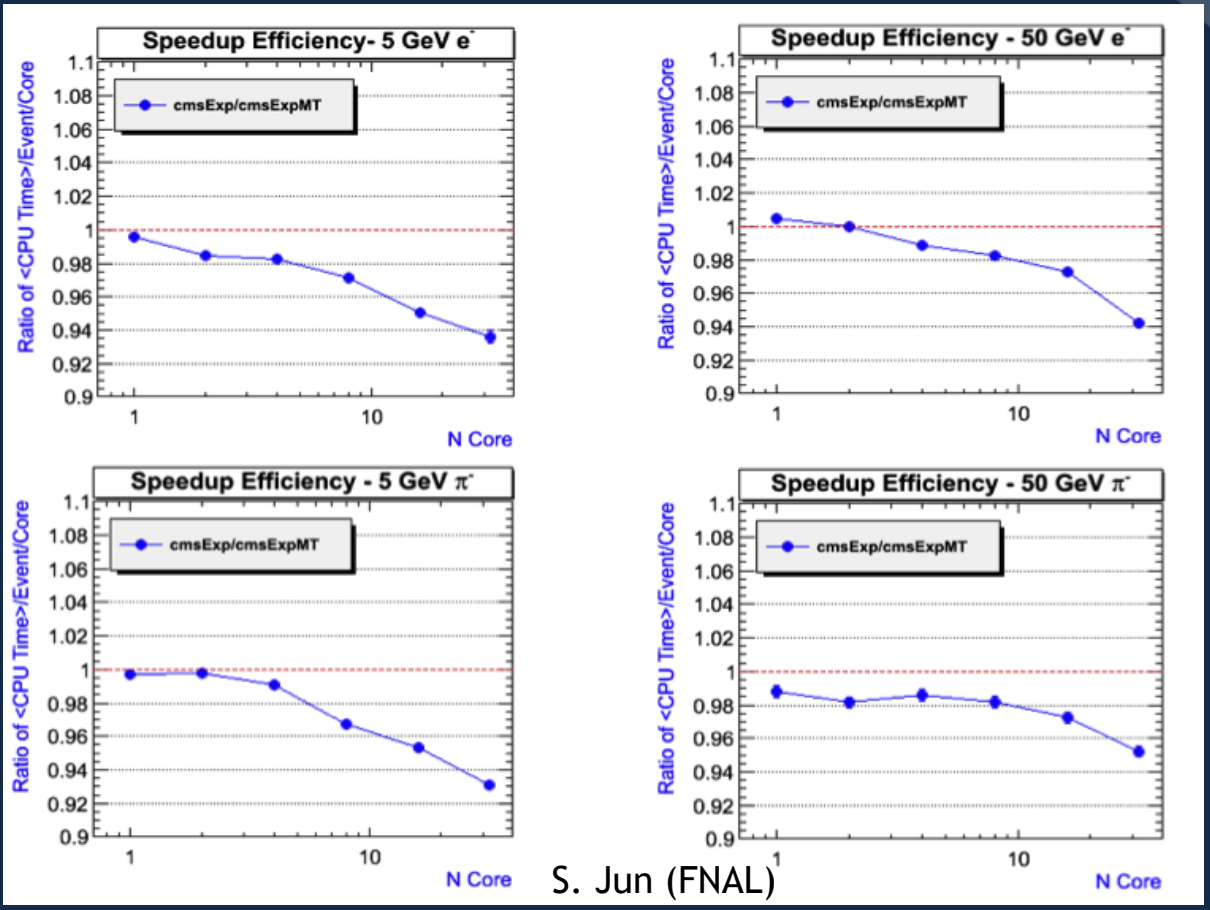

S. Jun (FNAL)

Profiling & Monitoring Results

G4MT memory reduction tests: mem (RSS-shared) vs. # of cores

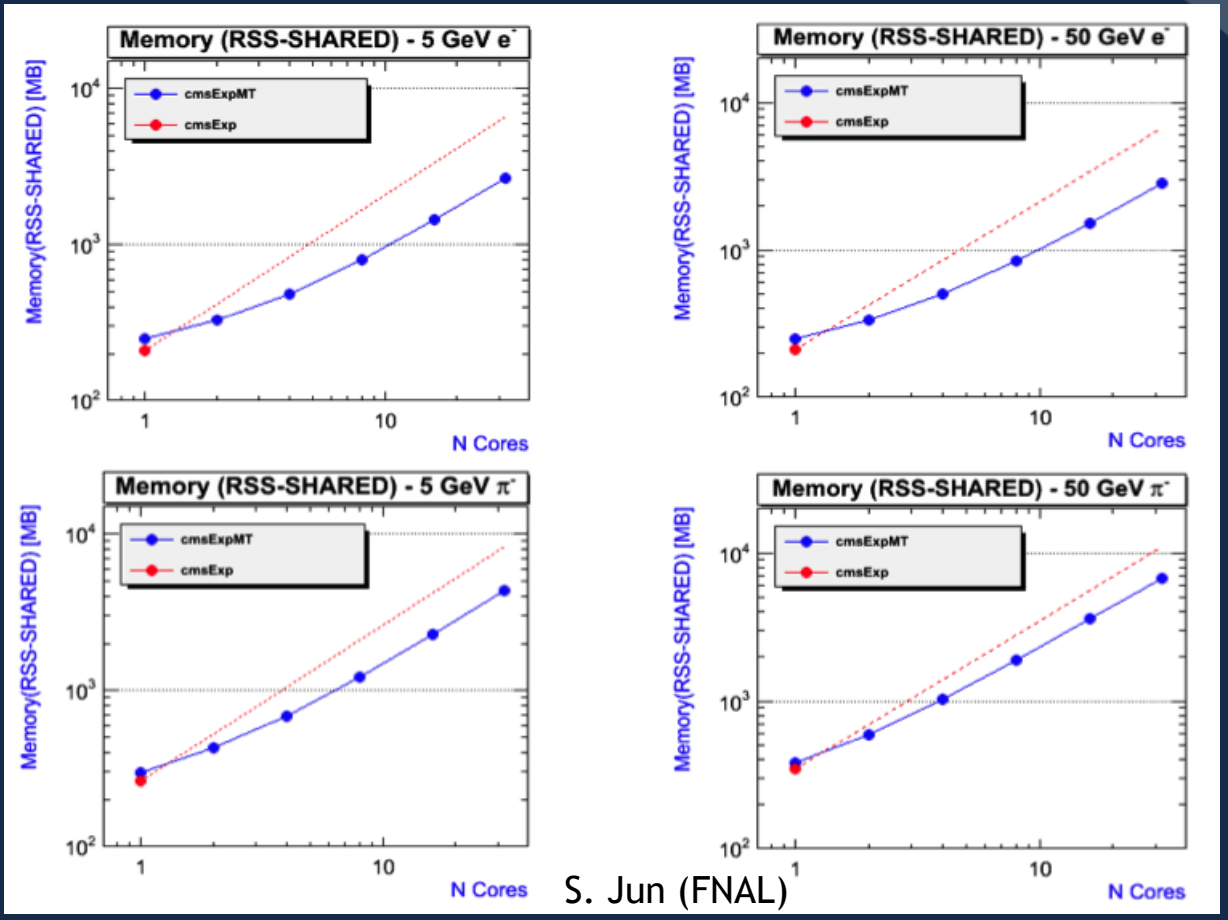

S. Jun (FNAL)

EM Code Review

The G4CPT assembled an interdisciplinary team with the (partial) support of the US DOE (HEP & ASCR-Scientific Computing)

J. Apostolakis (CERN), A. Dotti (SLAC), K. Genser (FNAL), S. Jun (FNAL), B. Norris (former ANL now U of Oregon)

Review of performance aspects of EM code G4 classes considering

- Correctness, performance, maintainability, adaptability
- Multi-threading aspects
- Potential issues related to parallelization, migration to GPU
- Issues related to potential migration to C++11

EM Code Review

SimplifiedCalo Application is used to study the effects of proposed modifications. Areas covered so far:

- G4PhysicsVector and underlying classes, especially (Compute)Value. (1.5% performance improvement.)
- G4Physics2DVector, changed type of underlying container to float. (~12% improvement in GPU, ~13% degradation in CPU.)
- Use random number generator array API in G4UniversalFluctuations. (10% faster in unit tests.)
- Started looking at G4VEmProcess
- Plan to review computing intensive EM functions, especially in classes derived from G4VContinuousDiscreteProcess

A written report will be produced within the next couple of months with recommendations

G4.10 (MT) and the Experiments

Successful testing of a new G4 product with novel capabilities needs engagement of the major HEP experiments

- Commissioning involves testing new features, Computing Performance benchmarking, and Physics Validation
- Critical to establish close communication with ATLAS and CMS through their commissioning process (not the only customers but the ones with tightest timelines and largest resources)
- Our G4CPT will focus only on the CP benchmarking process and leave the rest to our G4 MT and physics validation task teams.

G4.10 (MT) CP in ATLAS and CMS

ATLAS

- Migration to G4.6 has been time consuming, no concrete plans for G4.10 migration testing but they will probably setup a test build soon after G4.10 release

CMS

- CMSSW_7_0_0 with multi-threaded framework in Nov. 2013 will include a G4.10 release for validation
- Plan a multi-step validation of the CMS simulation application to check “physics changes independently from threading support”
- CMS will start first with physics validation and then move to Computing Performance benchmarking

The G4CP task will offer support and be in close communication with V. Ivantchenko, G. Cosmo, J. Apostolakis, A. Dotti through this process

Summary

The G4CP task is delivering all items in the to do list elaborated a couple of years ago from Collaboration input

- Profiling and CPU/mem benchmarking of every reference and public release documented in web page
- Protocol updated to incorporate new observables suggested by various collaborators
- Protocol includes MT capabilities testing as of April 2013
- Code improvement and optimization: EM code review
- Communication of CP results and problems through email to Testing & QA, working group leaders, SB mailing list

Comments and suggestions are always welcome