



GPCP Profiling Tools and Results

Philippe Canal, Daniel Elvira,
Krzysztof Genser, Soon Yung Jun
(Fermilab)

Geant4 Collaboration Meeting
September 23-27, 2013
Seville, Spain

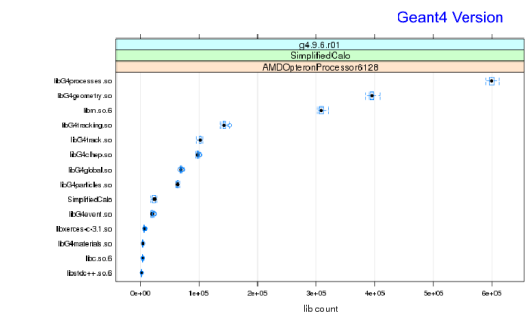

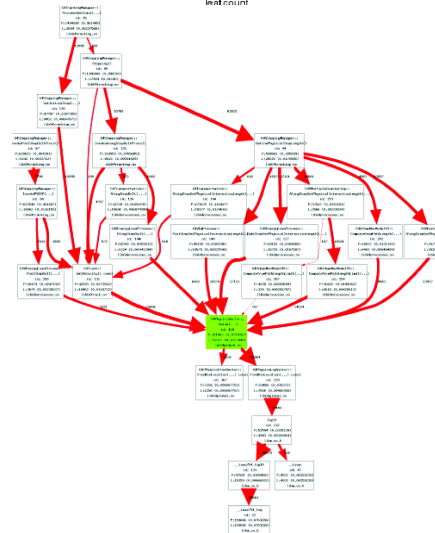


Performance Monitoring

- **Benchmarking for new releases**
 - CPU time per event
 - total memory per run
- **Profiling**
 - fraction of CPU (incl/exclusive)
 - memory footprint (caller/ee)
 - call graph
- **Multi-threaded Geant4 applications**
 - event throughput (scalability)
 - memory reduction



Review of Tools and Metrics

- Current profiling tools
 - FAST (CPU)
 - IgProf (memory)
- Primary metrics
 - time (version, physics)
 - leaf counts (funcs, libs)
 - call path analysis
 - memory footprints on the heap (live, max, total)
- Many other metrics added in 2013



Counter: MEM_LIVE, Sorted by self cost

Rank	Total %	Self	Calls	Symbol name	Callers
26	23.40	8,022,592	9,840	G4AllocatorPool::Grow()	
42	18.24	6,255,000	140,152	_gnu_cxx::new_allocator<double>::allocate(unsigned long, void const*)	
54	12.72	4,362,152	25,130	std::vector<double, std::allocator<double> >::reserve(unsigned long)	
51	11.78	4,038,528	9,708	G4NucleonLevelManager::UseLevelOrMakeLevel(G4NucleonLevel*)	
66	4.79	1,643,440	5	G4PairProduction::InitialiseEnergyLossProcess(G4ParticleDefinition const*, G4ParticleDefi	
23	4.20	1,440,504	351	G4ParticleChange::G4ParticleChange()	
84	3.40	1,164,194	30,717	std::basic_string<char, std::char_traits<char>, std::allocator<char> >::Rep::_S_createUn	
115	1.92	657,376	2	G4PairProduction::InitialiseEnergyLossProcess(G4ParticleDefinition const*, G4ParticleDef	
93	1.57	537,312	40	G4HadronLastCPhysics::ConstructProcess()	
26	1.01	346,664	1,173	G4VaporationDefaultGEMFactory::CreateChannel()	
160	0.94	323,136	612	G4ScatteringCollision::G4ScatteringCollision()	
161	0.94	322,000	35	G4Fancy3DNucleus::G4Fancy3DNucleus()	
168	0.88	301,008	6,271	std::_Rb_tree<GString, std::pair<GString const, G4ParticleDefinition*>, std::_Selectio	
179	0.79	271,584	2,314	std::vector<double, std::allocator<double> >::_M_fill_insert(_gnu_cxx::normal_iterator	

Newly added in 2013 (I)

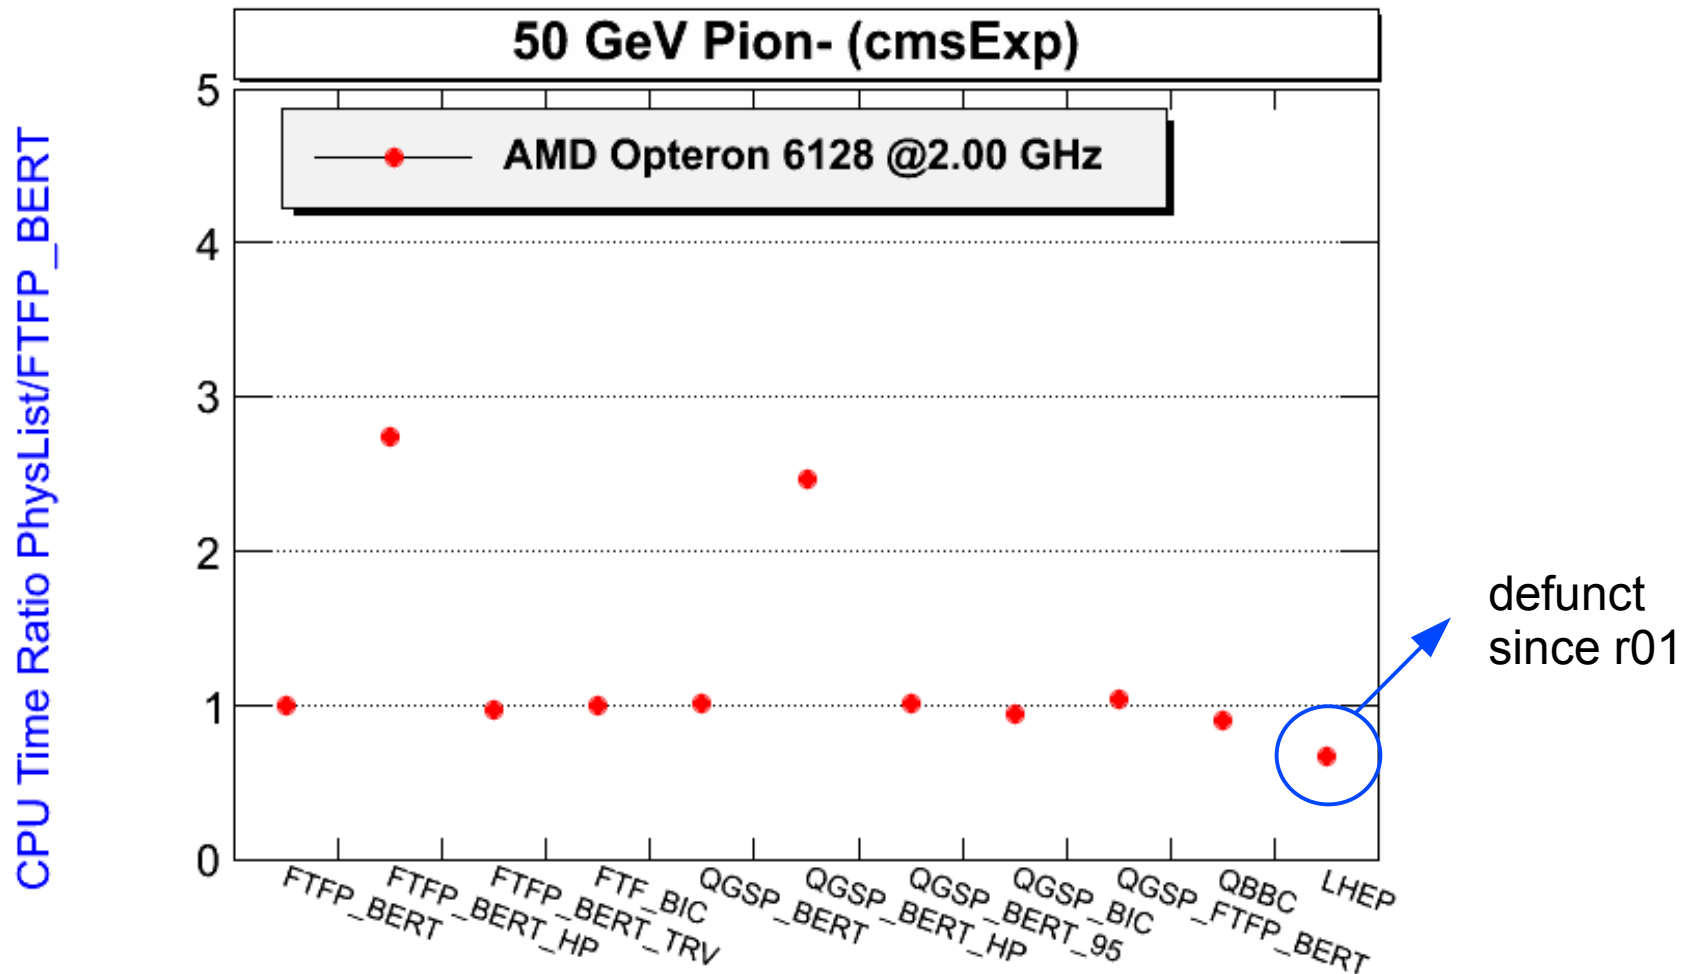
- Number of tracks and steps: a probe whether changes come from physics or geometry: 9.6.r08 H \rightarrow ZZ – proposed by Andrea Dotti and et. al.

Name	Particle	Mean	Sigma	Low	1/4Quad	Median	3/4Quad	High
1	Nstep e-	1.9778e+07	5.6091e+05	1.8313e+07	1.9417e+07	1.9810e+07	2.0209e+07	2.0880e+07
2	Nstep e+	2.5875e+06	1.3751e+05	2.2888e+06	2.4932e+06	2.5858e+06	2.6836e+06	2.8880e+06
3	Nstep gamma	1.7243e+07	4.4169e+05	1.6055e+07	1.6965e+07	1.7277e+07	1.7575e+07	1.8090e+07
4	Nstep N	6.9322e+06	7.3576e+05	5.0607e+06	6.5515e+06	6.9993e+06	7.4629e+06	8.2740e+06
5	Nstep other	4.5789e+05	4.8833e+04	3.2580e+05	4.3315e+05	4.6698e+05	4.9359e+05	5.5189e+05
6	Nstep p	3.2442e+05	3.4520e+04	2.3436e+05	3.0640e+05	3.2843e+05	3.4824e+05	3.8762e+05
7	Nstep pi-	3.9692e+04	4.4008e+03	2.8175e+04	3.6604e+04	4.0667e+04	4.2487e+04	4.8680e+04
8	Nstep pi+	4.1951e+04	4.8321e+03	2.8484e+04	3.9159e+04	4.2472e+04	4.5721e+04	5.1143e+04
9	Ntrack e-	1.4937e+07	3.7792e+05	1.3913e+07	1.4704e+07	1.4963e+07	1.5204e+07	1.5657e+07
10	Ntrack e+	3.4911e+05	1.4789e+04	3.1472e+05	3.3933e+05	3.4963e+05	3.5965e+05	3.7989e+05
11	Ntrack gamma	4.5984e+06	1.3902e+05	4.2423e+06	4.5084e+06	4.6067e+06	4.7058e+06	4.8740e+06
12	Ntrack N	2.2891e+05	2.4509e+04	1.6608e+05	2.1592e+05	2.3134e+05	2.4540e+05	2.7317e+05
13	Ntrack other	3.0140e+05	3.1920e+04	2.1909e+05	2.8400e+05	3.0490e+05	3.2287e+05	3.5925e+05
14	Ntrack p	8.7835e+04	9.3151e+03	6.3879e+04	8.3146e+04	8.8513e+04	9.4045e+04	1.0534e+05
15	Ntrack pi-	2.8052e+03	3.0957e+02	2.0320e+03	2.6150e+03	2.8835e+03	3.0210e+03	3.3890e+03
16	Ntrack pi+	3.1007e+03	3.5035e+02	2.1550e+03	2.8650e+03	3.1190e+03	3.3800e+03	3.7600e+03

10% asymmetry?

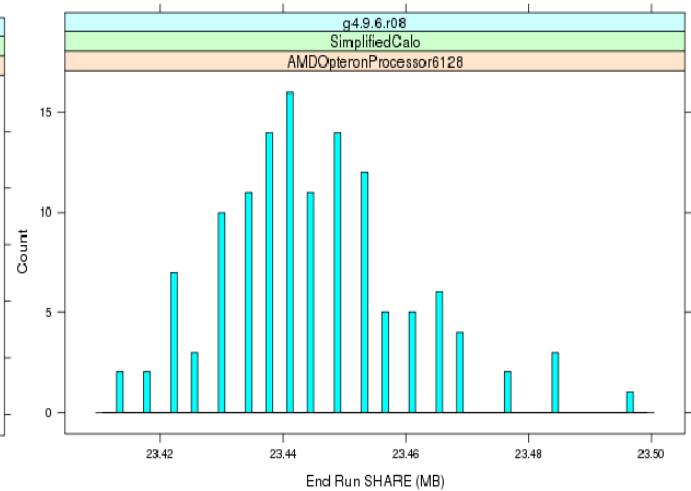
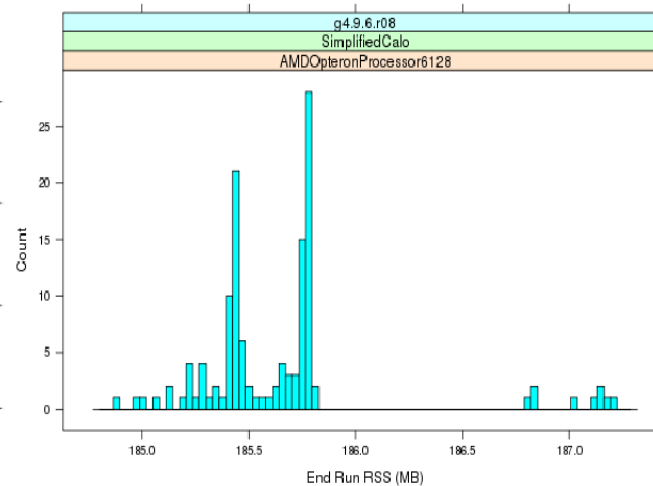

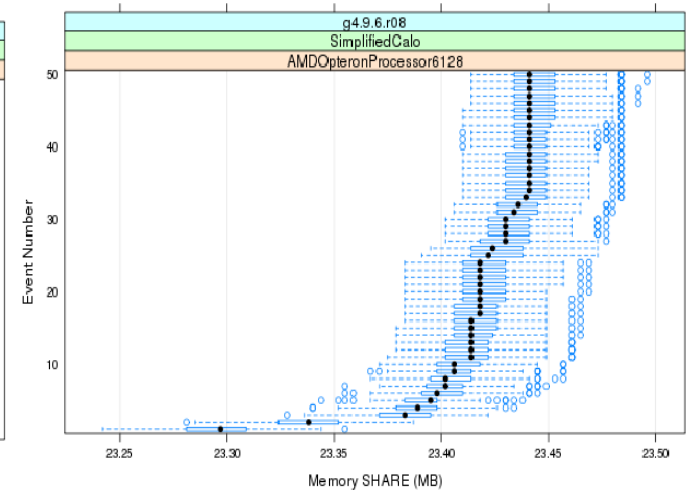
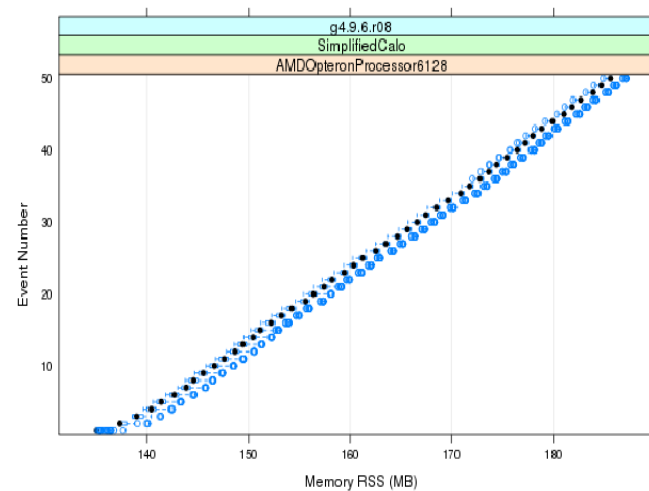

Newly Added (II)

- Performance by Physics List for major releases :
1, 5, 10, 50 GeV pions (cmsExp) – proposed by Alberto Ribon



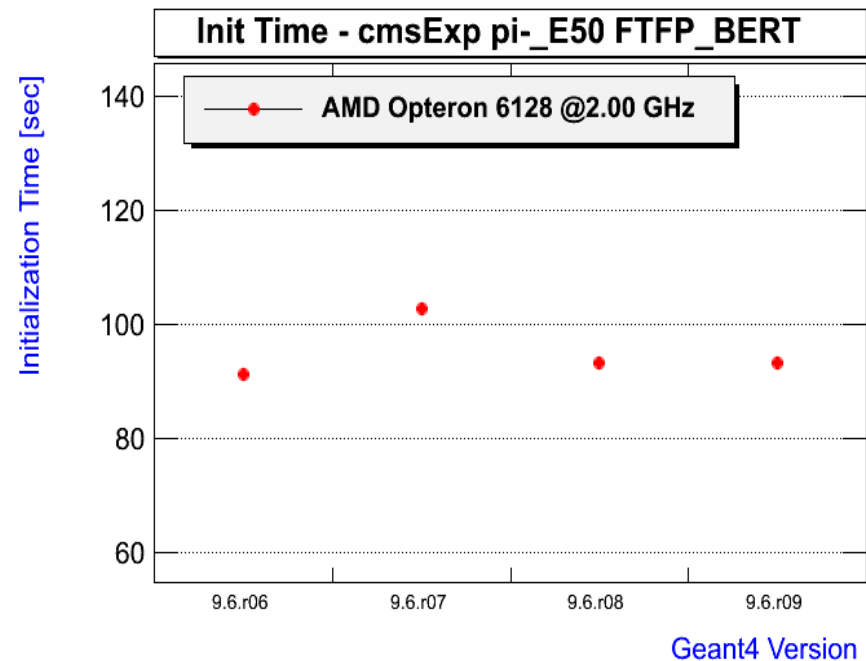
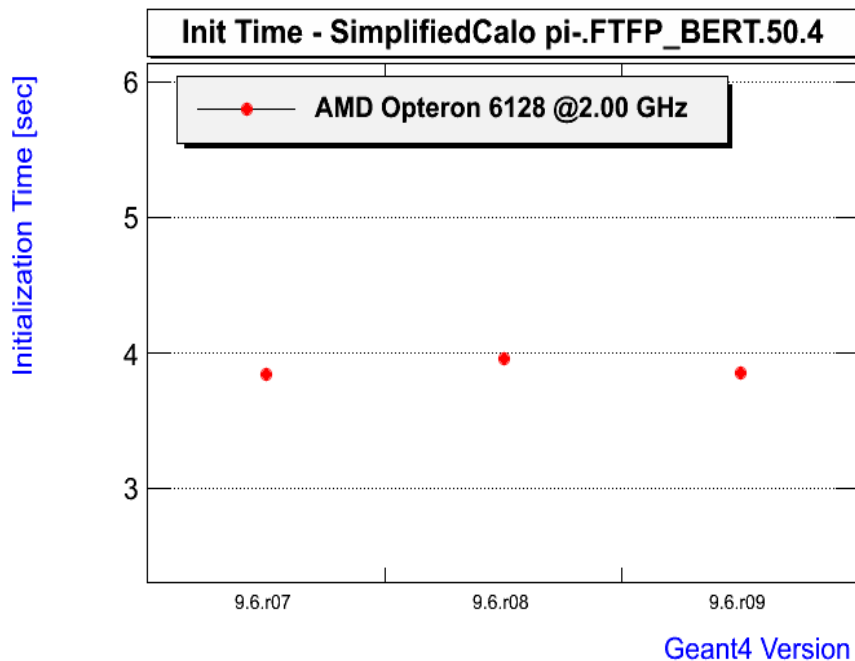
New Added 2013 (III)

- Memory usage: statm (vsize/rss/shared)




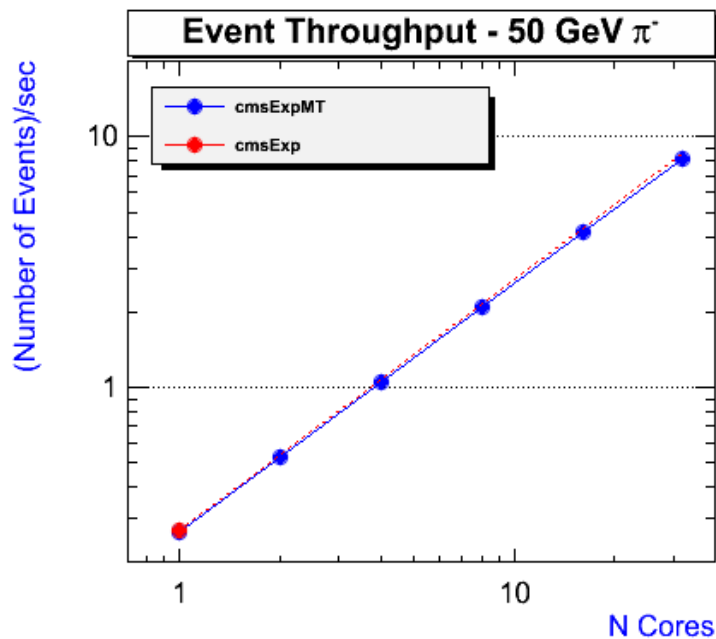


Newly Proposed in 2013 (IV)

- Initialization time – proposed by Makoto
 - measure performance for the initialization and the event loop separately
 - application specific (number of materials, physics list, particle type)
 - initialization time of SimplifiedCalo (~4 sec) and cmsExp (~100 sec)

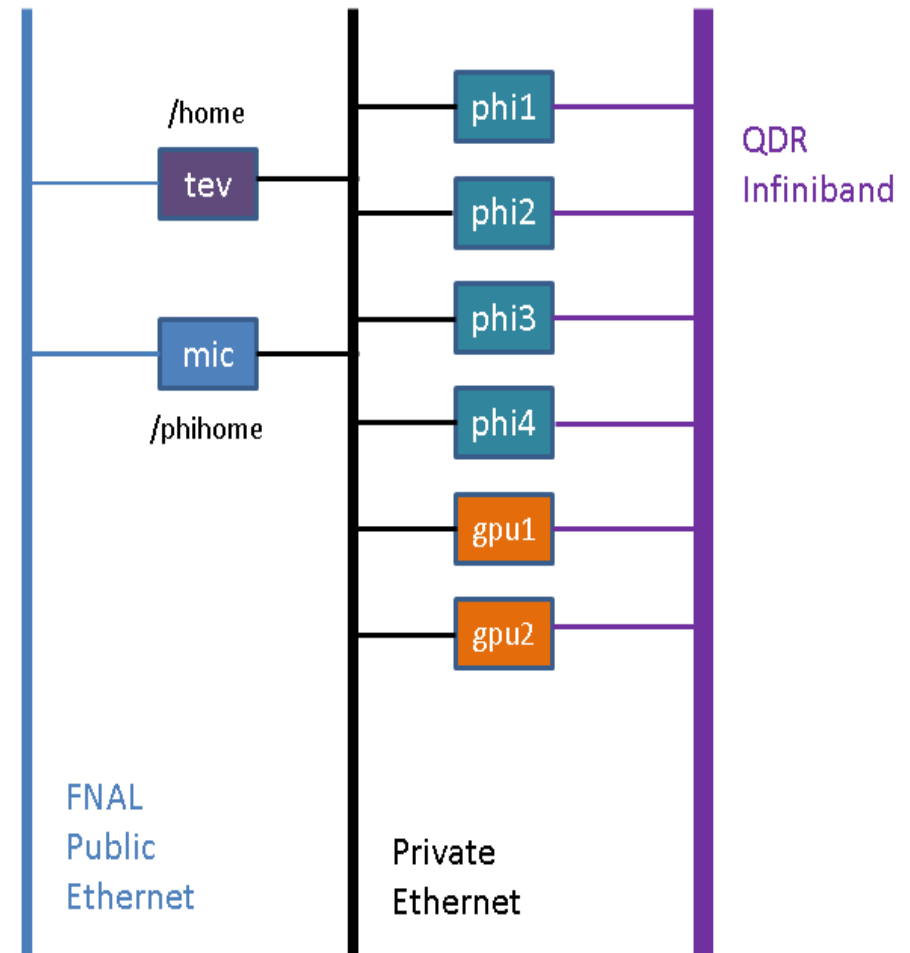


Preliminary Performance of Geant4MT



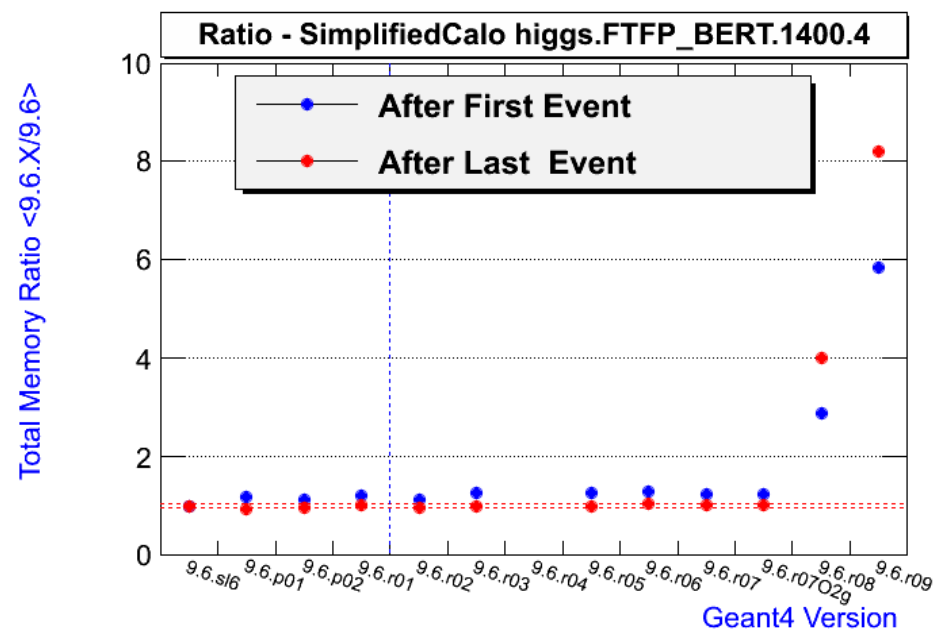
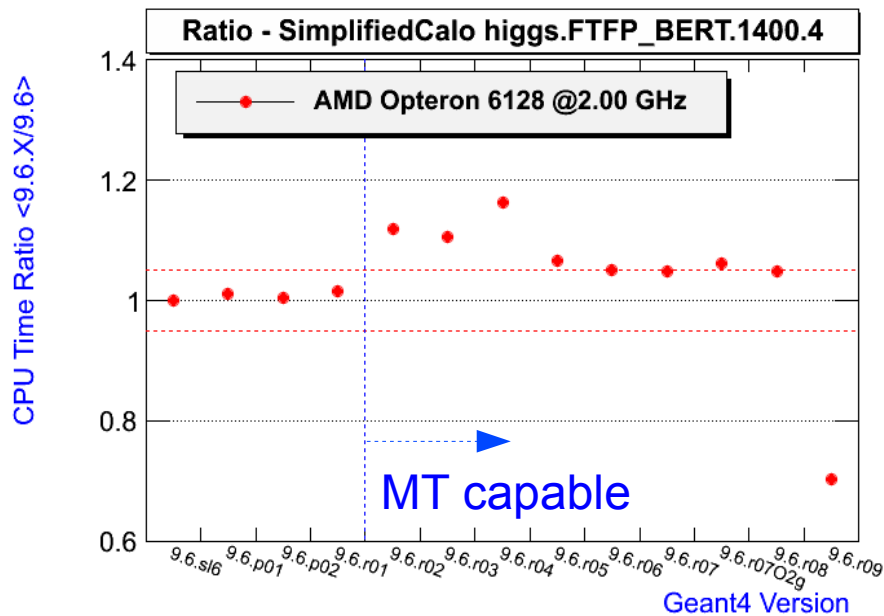
Plan to add Xeon Phi for benchmarking MT

- The Phi cluster at Fermilab
 - 4x12 cores Intel E5-2620 servers
 - 16 Intel Xeon Phi 5110P accelerators (4 per server)
 - QDR Infiniband
- Performance consideration
 - event throughput of MT
 - memory requirement
 - compiler dependency (icc)
 - (TBB application)



Summary: Geant4 9.6 Reference Releases

- Toward multi-threading capability since 9.6.ref02



- Compared to the pre-MT releases (H->ZZ)
 - 5% degradation up to ref-08 (mostly driven by hadron physics)
 - 30% improvement of the CPU performance in ref-09 is mainly driven by applyCuts in G4EmStandardPhysics
 - issues in total memory counts in ref-04, ref-08, ref-09?

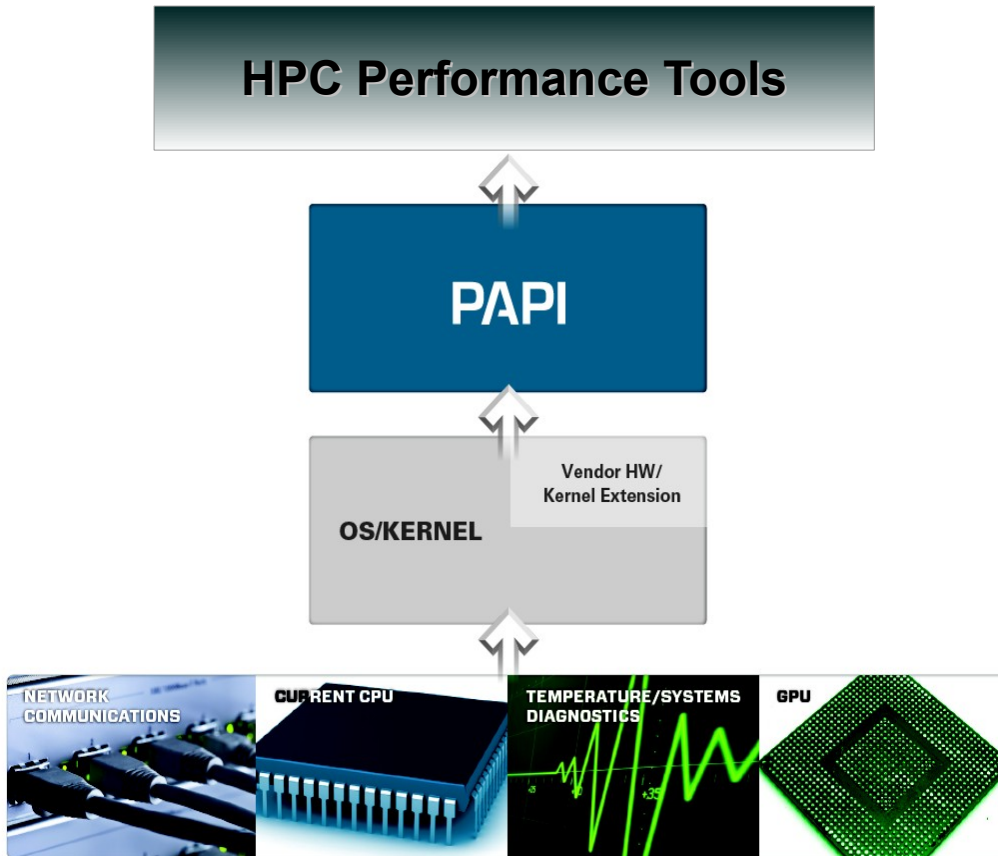
Performance Tools for Geant4-MT

- Requirements
 - Support parallelism and multi-threaded applications
 - Transparent instrumentation (applicable on binaries)
 - Light time overhead and precision of measurement
 - Advanced analysis (tracing, callgraph)
- Metrics for multi-threaded applications
 - speedup (event throughput, scalability)
 - memory (reduction, instruction/data cache miss, TLB miss)
 - I/O (hits)
- Platform dependence: architecture (multicore, MIC) and memory model

Performance Tools Considered

- Community infrastructures: PAPI
- Open source integrated tool kits reviewed
 - HPCToolkit (Rice Univ.)
 - Open|SpeedShop (Krell)
 - TAU (Univ. of Oregon)
- Licensed tools (not reviewed), but possible used for MIC
 - Intel: VTune Amplifier XE, ITAC
 - Allinea (DDD and MAP)

PAPI (Performance API)



- A standard API to access hardware performance counters
- Relation between software performance and processor events
- Event metrics : platform specific metrics, cache hit/miss, TLB miss, Flops, power consumption (MuMMI)

Open|Speedshop

- Comprehensive performance analysis for sequential, multithreaded, and MPI applications
- The base functionality includes
 - sampling experiments
 - support callstack analysis
 - hardware performance counters
 - multi-threaded, MPI profiling and tracing
 - floating point exception analysis
- GUI and CLI (command line instruction)
- Almost ready to support MIC

Default View and Stats Panel

The screenshot displays the OpenSpeedShop application window. At the top, there are menu bars for 'File', 'Tools', and 'Help'. Below the menu bars, there are several toolbars and panels:

- Process Control:** Includes buttons for 'Run', 'Cont', 'Pause', 'Update', and 'Terminate'. A tooltip for the 'Update' button reads 'Update the display with the current information.'
- Status:** A text area showing 'Process Loaded: Click on the "Run" button to begin the experiment.'
- Stats Panel [1] | ManageProcessesPanel [1]:** A toolbar containing various icons for monitoring and control.
- View/Display Choice:** A dropdown menu currently set to 'Functions', with other options being 'Statements' and 'Linked Objects'.
- Executables:** A text field showing '/home/syjun/g4p/test/openss/cmsExpMT/bin/cmsExpMT Host: cluck.fnal.gov Pids: 1 Threads: 33'.
- % of Total Exclusive CPU Time:** A bar chart on the left showing the relative CPU usage of different functions.
- Table:** A table with four columns: 'Exclusive CPU time', 'Inclusive CPU time', '% of Total Exclusive CPU Time', and 'Function (defining location)'. The top row is highlighted in black.
- Command Panel:** A text input field at the bottom with the prompt 'openss>>'.

Red annotations highlight specific areas:

- A red box around the toolbar icons is labeled 'Toolbars'.
- A red arrow points from the 'Update' button to the 'Toolbars' label.
- A red arrow points from the 'Functions' dropdown to the 'Top Functions' label.
- A red circle highlights the table content.

Exclusive CPU time	Inclusive CPU time	% of Total Exclusive CPU Time	Function (defining location)
105.742855	105.742855	11.216171	__tls_get_addr (/lib64/ld-2.12.so)
95.485712	95.485712	10.128193	__ieee754_log (/lib64/libm-2.12.so)
56.685713	56.685713	6.012668	__ieee754_exp (/lib64/libm-2.12.so)
44.771428	62.885713	4.748917	G4HadronCrossSections::CalcScatteringCrossSection
30.914285	156.999997	3.279086	G4CrossSectionDataStore::GetCrossSection (/home/s
20.400000	64.571427	2.163833	G4ElasticHadrNucleusHE::HadrNucDifferCrSec (/hor
16.028571	47.514285	1.700155	G4Navigator::LocateGlobalPointAndSetup (/home/sy
15.400000	15.400000	1.633482	cmsExpMagneticField::GetFieldValue (/home/syjun/g
12.571428	62.085713	1.333455	G4hPairProductionModel::ComputeDMicroscopicCro
11.885714	14.371428	1.260721	G4ProductionCutsTable::ScanAndSetCouple (/home/

Profiling Output (Text-based)

Open|SpeedShop (oss pcsamp)

Geant4.9.6.r07 cmsExpMT

Sample	Energy	Sequential	Multi Thread
Electrons	5 GeV	S1	T1 T2 T4 T8 T16 T32
	50 GeV	S1	T1 T2 T4 T8 T16 T32
Pions	5 GeV	S1	T1 T2 T4 T8 T16 T32
	50 GeV	S1	T1 T2 T4 T8 T16 T32

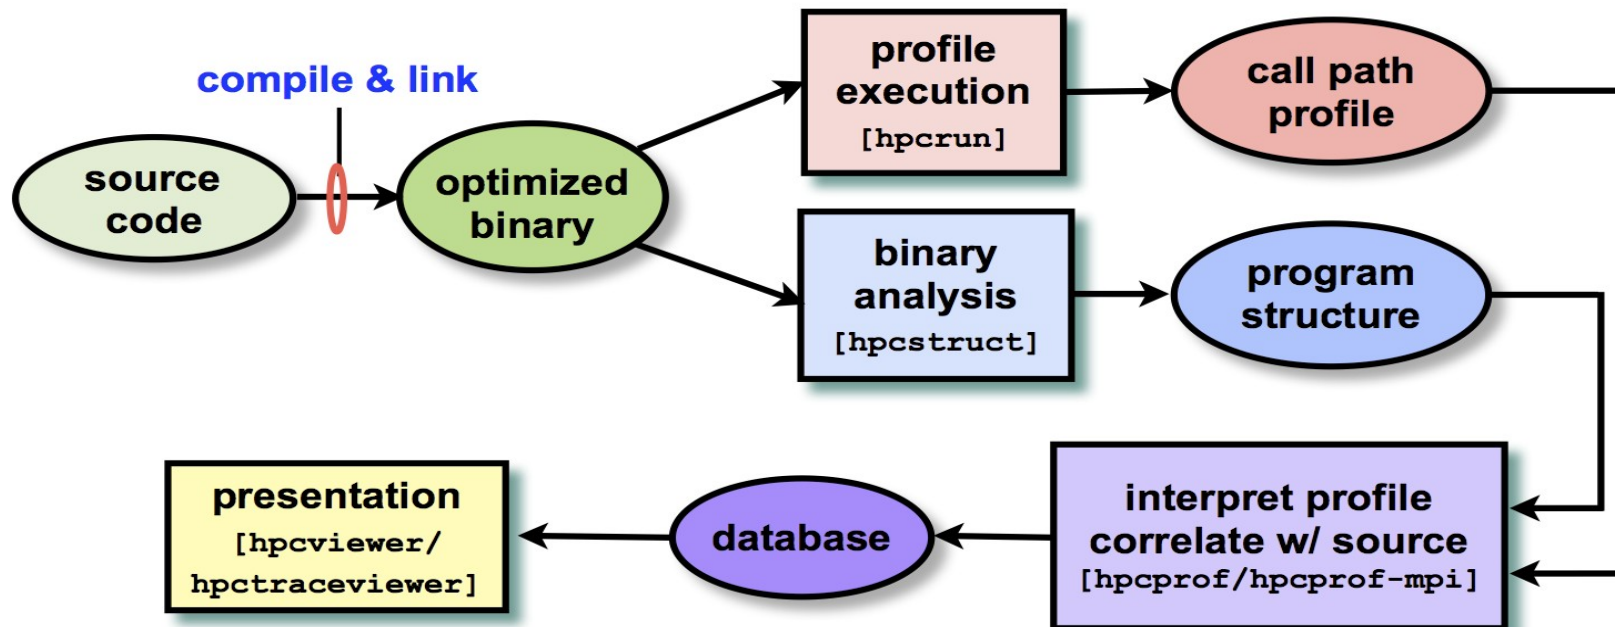
Exclusive CPU time in seconds.	% of CPU Time	Function (defining location)
163.110000	4.156505	__ieee754_log (libm-2.12.so)
140.450000	3.579064	cmsExpMagneticField::GetVolumeBase
115.210000	2.935878	G4Navigator::LocateGlobalPointAndS
100.390000	2.558222	G4CrossSectionDataStore::GetCrossS
85.850000	2.187701	__ieee754_atan2 (libm-2.12.so)
80.550000	2.052642	G4PhysicsVector::Value (libG4globa
78.480000	1.999893	__ieee754_exp (libm-2.12.so)
71.800000	1.829668	G4PhotoNuclearCrossSection::GetIso
68.650000	1.749397	G4SteppingManager::DefinePhysicalS
59.510000	1.516484	CLHEP::RanecuEngine::flat (libG4cl
58.190000	1.482846	G4Transportation::AlongStepGetPhys
54.980000	1.401046	G4HadronCrossSections::CalcScatter
54.580000	1.390853	G4Navigator::ComputeStep (libG4geo
50.420000	1.284845	G4SteppingManager::Stepping (libG4
48.690000	1.240759	G4CrossSectionDataStore::GetIsoCro
48.510000	1.236172	G4VoxelNavigation::ComputeStep (li
47.810000	1.218334	G4BGGNucleonInelasticXS::CoulombFa
46.130000	1.175523	G4PhysicsLogVector::FindBinLocatio
44.570000	1.135770	G4NavigationLevel::~~G4NavigationLe
43.410000	1.106210	G4CrossSectionDataStore::GetCrossS
39.380000	1.003514	G4Mag_UsualEqRhs::EvaluateRhsGiven
39.380000	1.003514	G4SteppingManager::InvokePSDIP (li
39.310000	1.001730	G4ClassicalRK4::DumbStepper (libG4
38.760000	0.987715	G4EnhancedVecAllocator<G4Navigatio

Exclusive CPU time in seconds.	% of CPU Time	Function (defining location)
167.360000	4.208250	__ieee754_log (libm-2.12.so)
139.140000	3.498661	cmsExpMagneticField::GetVolum
108.970000	2.740039	G4Navigator::LocateGlobalPoin
95.920000	2.411899	G4CrossSectionDataStore::GetC
86.690000	2.179811	G4PhysicsVector::Value (libG4
82.830000	2.082752	__ieee754_exp (libm-2.12.so)
80.040000	2.012598	__ieee754_atan2 (libm-2.12.so)
67.350000	1.693509	G4SteppingManager::DefinePhys
65.110000	1.637184	G4PhotoNuclearCrossSection::G
60.280000	1.515734	CLHEP::RanecuEngine::flat (li
57.070000	1.435019	G4HadronCrossSections::CalcSc
56.200000	1.413143	G4CrossSectionDataStore::GetI
54.600000	1.372911	G4Navigator::ComputeStep (lib
53.950000	1.356567	G4VoxelNavigation::ComputeSte
48.770000	1.226317	G4SteppingManager::Stepping (
46.890000	1.179044	G4PhysicsLogVector::FindBinLo
46.440000	1.167729	G4Transportation::AlongStepGe
46.070000	1.158425	G4CrossSectionDataStore::GetC
45.170000	1.135795	G4BGGNucleonInelasticXS::Coul
42.980000	1.080728	G4SteppingManager::InvokePSDI
40.440000	1.016860	G4NavigationLevel::~~G4Navigat
39.710000	0.998504	G4Mag_UsualEqRhs::EvaluateRhs
39.300000	0.988194	G4EnhancedVecAllocator<G4Navi
38.800000	0.975622	G4PropagatorInField::ComputeS

Exclusive CPU time in seconds.	% of CPU Time	Function (defining location)
4445.200000	3.455208	cmsExpMagneticField::GetVolumeBaseBfield (cmsExpMT)
4220.720000	3.280722	G4Navigator::LocateGlobalPointAndSetup (libG4geometr
4173.260000	3.243832	__ieee754_log (libm-2.12.so)
3324.950000	2.584449	G4CrossSectionDataStore::GetCrossSection (libG4proce
2978.090000	2.314839	G4PhysicsVector::Value (libG4global.so)
2641.130000	2.052923	__ieee754_atan2 (libm-2.12.so)
2351.700000	1.827952	__ieee754_exp (libm-2.12.so)
2308.410000	1.794303	G4SteppingManager::DefinePhysicalStepLength (libG4tr
2036.690000	1.583098	G4PhotoNuclearCrossSection::GetIsoCrossSection (libC
2001.680000	1.555885	CLHEP::RanecuEngine::flat (libG4clhep.so)
1787.340000	1.389281	G4VoxelNavigation::ComputeStep (libG4geometry.so)
1734.460000	1.348178	G4CrossSectionDataStore::GetIsoCrossSection (libG4pr
1706.780000	1.326663	G4HadronCrossSections::CalcScatteringCrossSections (
1639.070000	1.274032	G4Transportation::AlongStepGetPhysicalInteractionLen
1622.740000	1.261339	G4Navigator::ComputeStep (libG4geometry.so)
1550.650000	1.205304	G4CrossSectionDataStore::GetCrossSection (libG4proce
1436.450000	1.116538	G4BGGNucleonInelasticXS::CoulombFactor (libG4process
1435.940000	1.116141	G4SteppingManager::Stepping (libG4tracking.so)
1432.300000	1.113312	G4SteppingManager::InvokePSDIP (libG4tracking.so)
1423.790000	1.106697	G4PhysicsLogVector::FindBinLocation (libG4global.so)
1317.140000	1.023799	G4Mag_UsualEqRhs::EvaluateRhsGivenB (libG4geometry.s
1257.190000	0.977201	G4EnhancedVecAllocator<G4NavigationLevel>::allocate
1240.180000	0.963979	G4NavigationLevel::~~G4NavigationLevel (libG4geometry
1200.970000	0.933502	G4PropagatorInField::ComputeStep (libG4geometry.so)

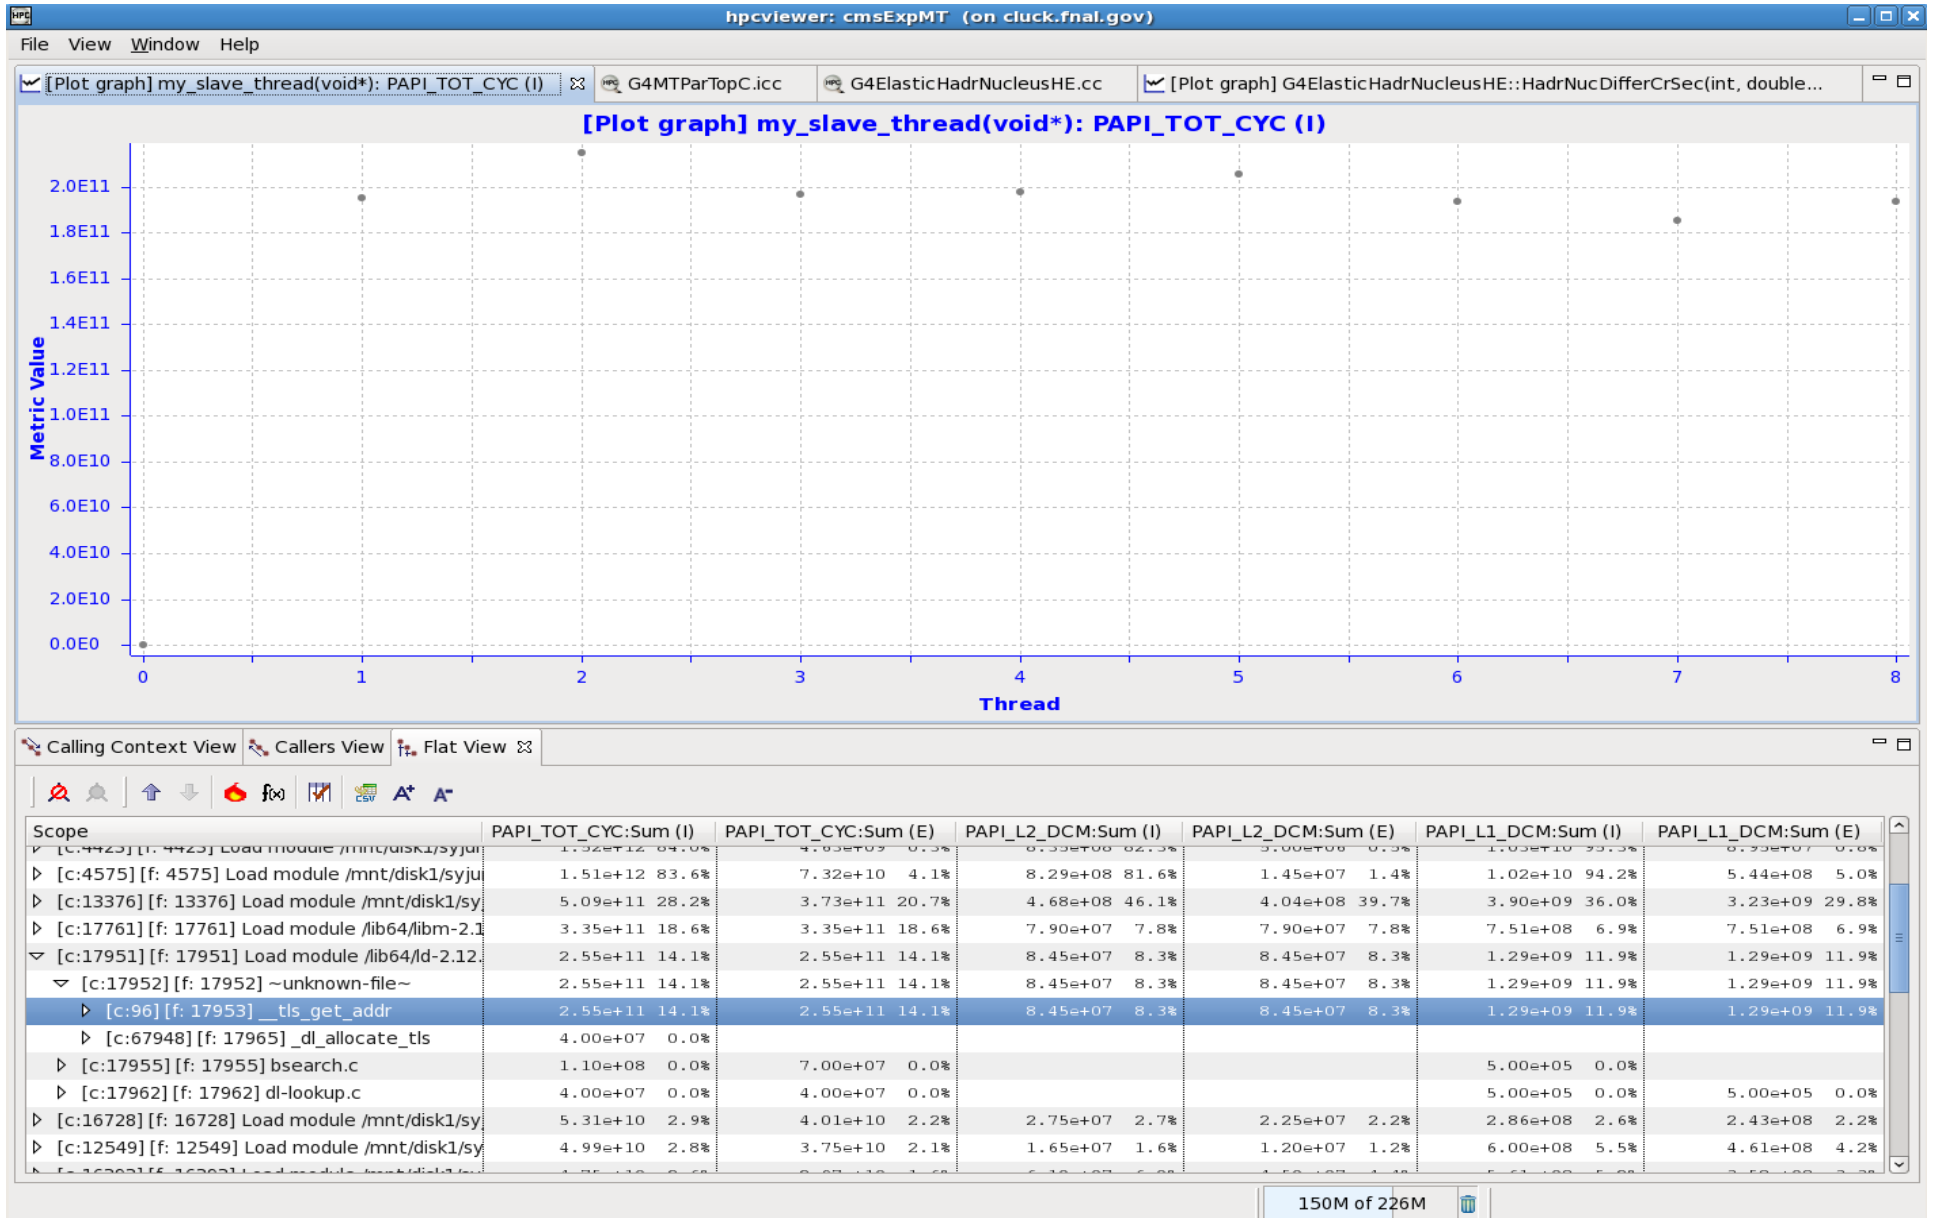
HPCToolkit

- Overview of HPCToolkit tool's work flow (from manual)



- Code centric view, GUI and text-base flat profile
- Supporting performance analysis of heterogeneous architecture

HPCToolkit: hpcviewer



TAU (Tuning Application Utilities)

- Dynamic, compiler based, source based Instrumentation
- Analysis tools
 - ParaProf
 - PerfExplorer
 - Tracer (Jumpshot, vampir)
- Various built-in graphical presentations
- Advantage/disadvantage: compiler/source-based instrumentation
- Direct contact to experts/developers (B. Norris @Oregon)

Summary

- Geant4 performance monitoring has been continuously deployed during the 9.6 release period
- Preliminary performance measurements for multi-threaded Geant4 applications have been tested
- GPCP task force will continue to look at various tools and libraries to improve performance profiling and analysis

Your inputs are always welcome!