Geant4 developments in multi-threading, reproducibility, and physics

J. Apostolakis[*], M. Asai[+], G. Cosmo[*], A. Dotti[+], G. Folger[*], M. Gayer[*], V. Ivantchenko[¥*], M. Kelsey[+], T. Koi[+], T. Nikitina[*], W. Pokorski[*], A. Ribon[*], V. Uzinskiy[#], M. Verderi[§] and D.H. Wright[+]

Institutions: * CERN, ¥Ecoanalytica, # JINR, § LLNR, + SLAC

DRAFT * ACAT 2014 * Version: 2.7b **SFT practice talk** (last change 01:50 29/8/2014**)**

Overview

- Highlights of latest developments
- New Major release 10.0: first with Multi-Threading
- Strong reproducibility of events
- Improvements in Hadronic Physics
  - Improvement of cross sections and physics models
- Performance monitoring and improvement

Context

- Geant4 9.4 (Dec 2010) - baseline
  - Used in production in ATLAS, CMS for (most of) LHC Run 1
- Geant4 9.6 (Dec 2012) is 'long-term support' version
  - Patches for LHC experiment production into 2015+
- Geant4 10.0 (Dec 2013) was first **major** release in 6.5 years (since rel. 9.0 in June 2007)
  - Add big feature (MT) - plus reorganisations
  - Removed obsolete code/features, & some physics models
- Recent development release 10.1-beta (June 2014)
  - For the upcoming minor release 10.1 (Dec 2014)

MultiThreading - Intro

- Geant4 10.0: First release with Multi-Threading
  - Events are **independent**: each event can be simulated separately
  - Chosen event-level parallelism as current target
- Key goals: more cores within a memory 'budget'
  - Scaling to use more CPU cores
  - Reduce extra amount of memory per worker thread
  - Use caches (L2/L3) better by sharing 'constants'/tables
  - Small effort to port an application (compared to the effort to develop it.)
- First experiences
- Recent improvements

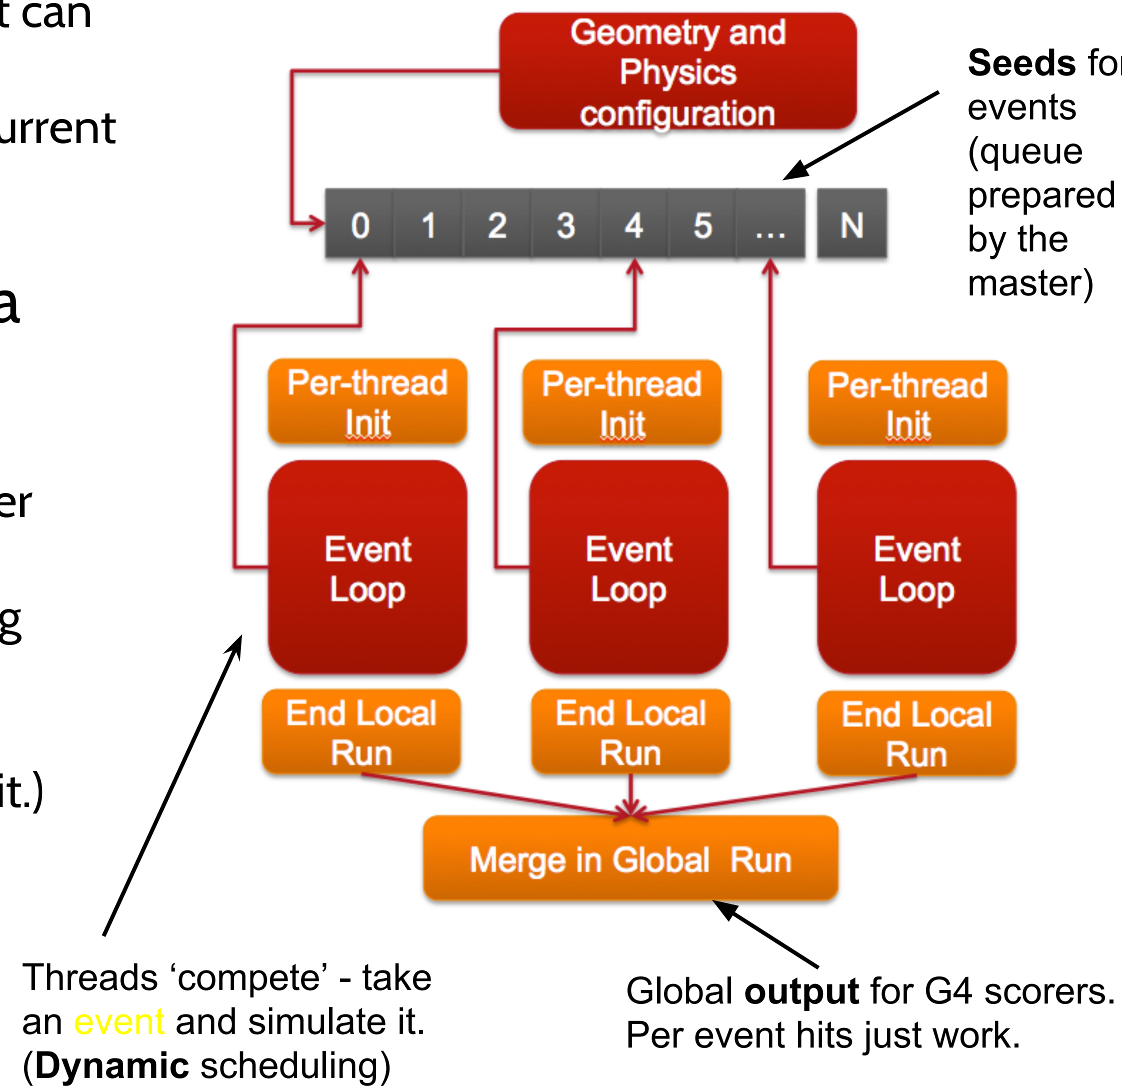

Shared Vs Thread-local

- To reduce memory footprint, threads must share part of their memory

- General rule in G4: ***threads are allowed to share whatever is invariant during the event loop*** (e.g. threads do not change those objects while processing events - they are treated as “read-only”)

The ‘constant’ parts which we identified are:

- Geometry definition
- Electromagnetic physics tables

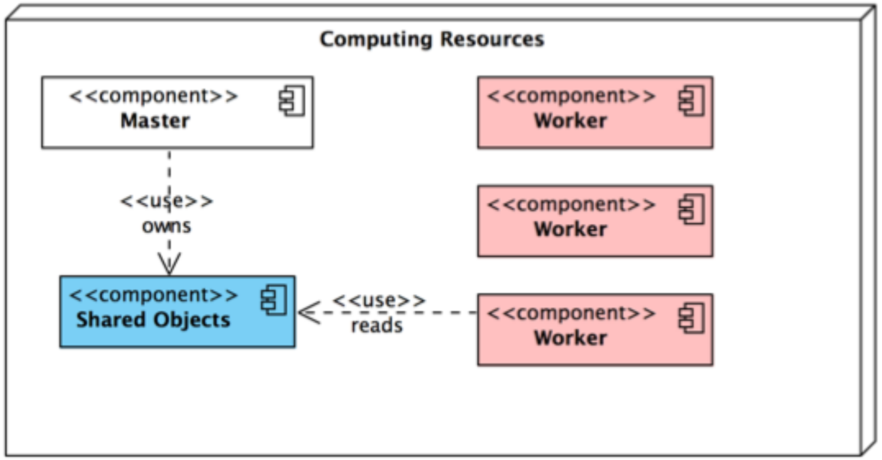

A few other processes have large tables

➢ a few are now shared, others under investigation

- Each additional thread: only 20-40MB

CPU / Memory performance

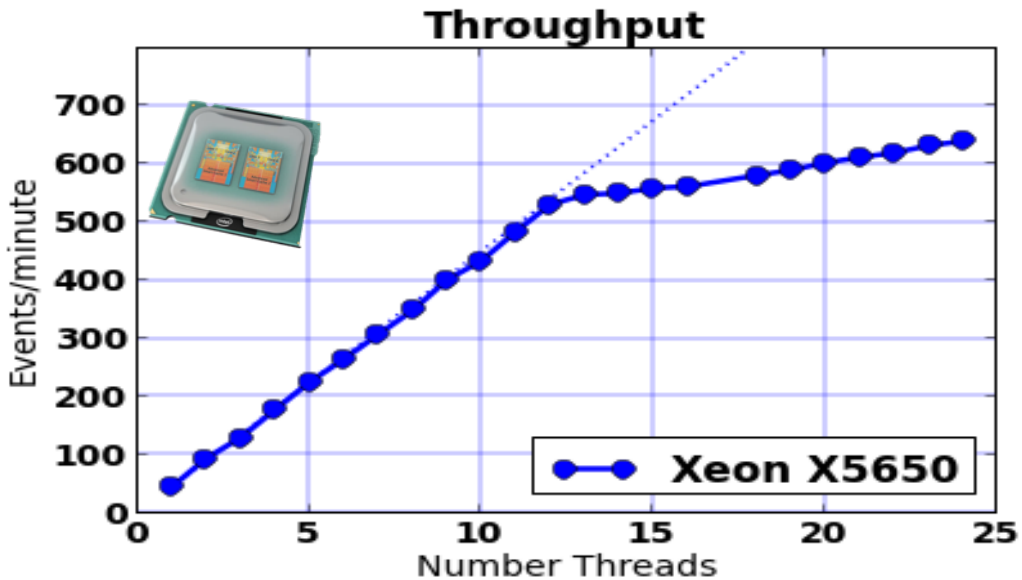

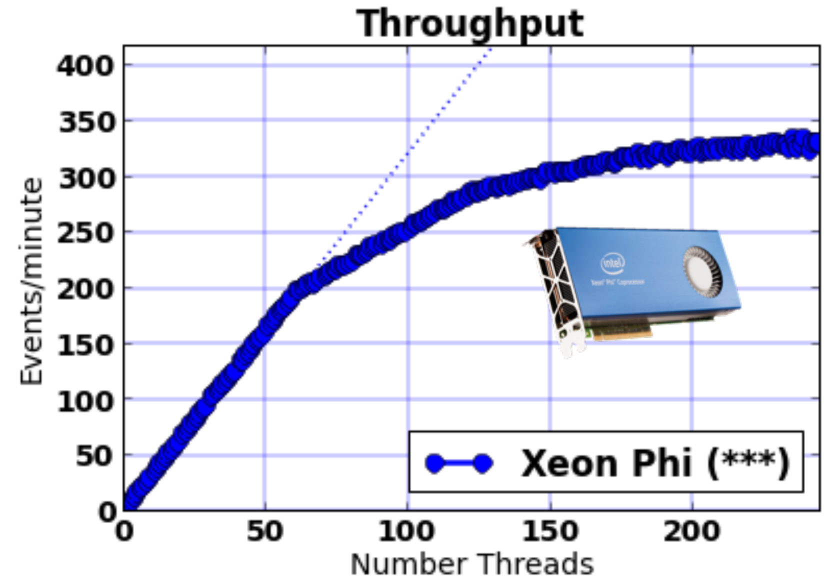

Geant4 10.0p02 on 12-core Xeon X5650

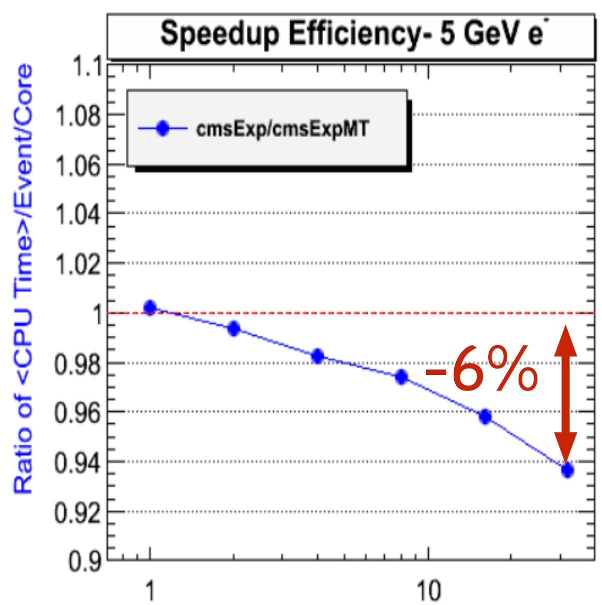

Obtained with HEP geometry

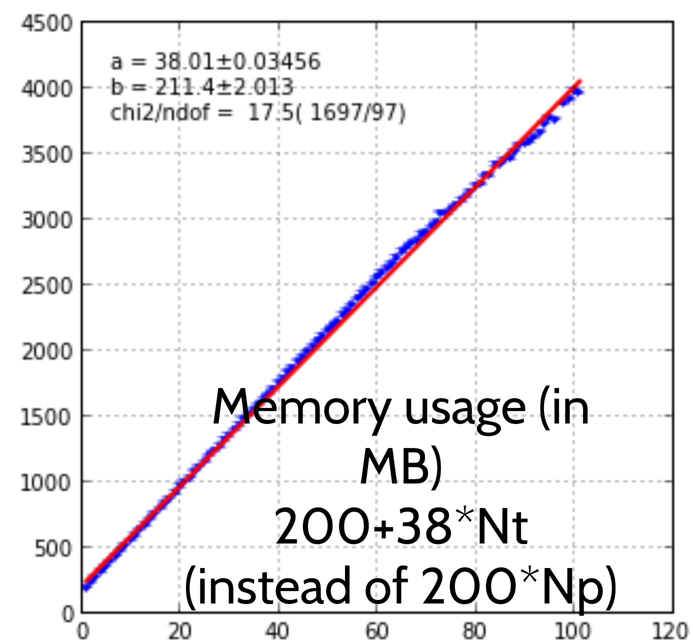

Multi-threading: recent

- Control of per-thread memory improved
  - Ensuring more ‘constants’ (tables) are shared
  - reduce memory reported ‘lost’ at the end of a job - using new pattern for per-thread statics/globals.
- Extended to other parallelism ‘frameworks’
  - MPI: having 2+ MPI ‘jobs’, each 2+ threads
  - TBB: Changing workers to be task-based
- Enabled by adhering to POSIX standards

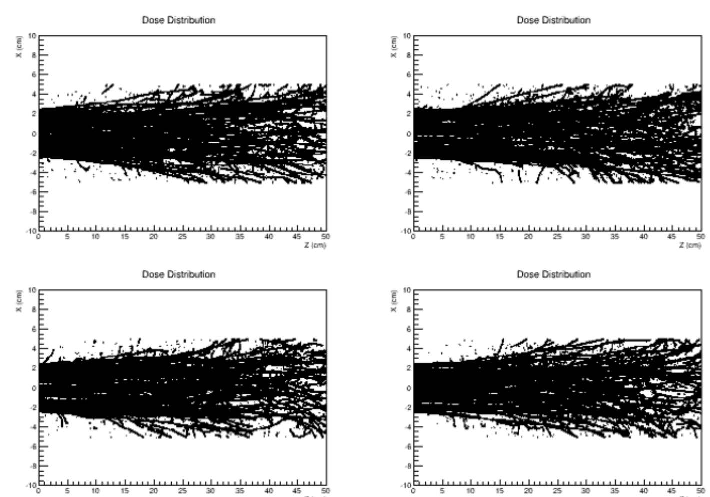

In progress - for version 10.1

- Further **reduce thread memory** footprint to fit O(100) threads on O (GB) RAM (e.g. typical for accelerators)
  - Already identified several areas for improvements
- **Improve control** of thread-local memory objects
  - create 'workspace' that owns thread-local objects, allowing it to be passed between threads (an inactive and a newly activated one.)
  - Needed to reduce memory use on task-based frameworks (e.g. TBB)
- **Improve visualization** with MT
- Allow threads to **join/leave workers pool** at any time during job
  - Simplify integration with TBB based applications, increase flexibility
- Review key CPU intensive algorithms to improve performance

Reproducibility

Requirement(s)
Testing for reproducibility
Finding source of problems and fixing

What is Reproducibility?

- ‘Strong’ reproducibility= same result for an event, independent of its ‘location’ in a job (first, middle or last) or whichever the series of previous events.
- Required for HEP experiment applications
  - To enable debugging of problems/issues
  - To allow repeatable results in MT mode - since the order of events depends on the number of workers (and relative performance/load.)
- It is a strong sign of the quality of implementation of the simulation

Testing reproducibility

- Tests devised to check reproducibility
  - Run an application with many events
  - Rerun each event in a separate job with same RNG seed/state and input tracks
- Check the results of each event
  - check the number of RNG seeds
  - compare observables(E-deposit, secondaries)

Problems and fixes

- First deployment of tests found several *discrepancies* (events with repetition errors)
  - Careful investigation needed to find source of each discrepancy
  - Typical error: caching a value of a cross-section using a different input (Energy within bin, or one isotope of an element)
- Expanded, regular testing introduced
  - It has found both new and existing errors.

Hadronic Physics - 1

What are (production) Physics Lists ?
Improving modeling of hadronic showers
Evolution of (production) physics lists

Hadronic Physics

- Goal: **replace** 'weak' physics models and plug the **gaps** between the strong 'theoretical' / 'phenomenological' models
- Starting point: QGSP_BERT production physics list for HEP experiments (2005-2012)
  - BERT(E<9.9) - 'stretched' cascade beyond 3-5 GeV
  - LEP (9.5 < E < 25 GeV) - missed ~10% E, obsolete parameterised models (now retired) - weak
  - QGS (E>12GeV) + Precompound/evaporation
- Developed Fritiof/FTF model in past 5 years
  - to plug gap below 5 GeV (and up to 12-15 GeV
  - provide alternative/improved model at high E

Starting point: Physics lists 9.4

QGSP_BERT (Geant4 9.4, Dec 2010)

EM Standard ('hard-wired')

- Urban93 Multiple Scattering
- 'Standard' Bremstrahlung

Hadronic final state models

- Elastic
  - CHIPS (H), Coherent (other)
- Inelastic (p,n,π) a
  - Bertini cascade: 0 - ***9.9*** GeV
  - ***Parameterised* (LEP)**: 9.5 - 25 GeV
  - ***QGS model***: 12 GeV - 1TeV
  - Precompound, evaporation, fission
  - Anti-baryons: LEP
- Quasi-elastic: CHIPS
- Stopping particles - CHIPS
- γ-nuclear, e-nuclear: CHIPS
- n-capture: LEP

FTFP_BERT (Geant4 9.4, Dec 2010)

EM ('hard-wired')

- Urban93 Multiple Scattering
- 'Standard' Bremstrahlung

Hadronic final state models

- Elastic
  - CHIPS (H), Coherent (other)
- Inelastic (p,n,π)
  - Bertini cascade: 0 - 5 GeV (revised)
  - ***Fritiof FTF model***: 4GeV - 1TeV
  - Precompound, evaporation, fission
  - Anti-baryons: LEP
- Quasi-elastic: FTF
- Stopping particles: CHIPS
- γ-nuclear, e-nuclear: CHIPS
- n-capture: LEP

Hadronic Improvement

- Introduced isomer, and create in key models
- Many improvements in string, cascade models

- New Production - inelastic physics models:
  - FTFP_BERT=BERT(E<5) + FTF(E>4)/Preco-Evap
  - Replaces QGSP_BERT= BERT(<9.9)+LEP(9.5-25)+QGS(>12)/Preco-Evap

Fritiof (FTF) improvement

- Add/tune Reggeon cascade de-excitation (G4 v9.3)
- Light anti-ion interactions (9.5) - $\bar{p}$, $\bar{n}$, $\bar{d}$, $\bar{t}$, $\bar{\alpha}$, $\overline{He^3}$
- Extension to nucleus-nucleus (10.0)
  - Interfaced to Binary Cascade, Pre-Compound
  - Support energy from 3-4 A*GeV to RHIC
- Recent / ongoing:
  - Correction in string decay
  - Tuning of parameters continues to describe h+p & h+A interactions - with NA49 data
  - more validation of nucleus-nucleus interaction

Hadronic Physics - II

Cascades
Pre-equilibrium and equilibrium processes
Introduction of isomers
Interactions of stopped particles

Cascades

- Treat projectiles 0.1-3 GeV (extended to 5-10 GeV)
- Nucleus = 1-6 shells (density) or individual nucleons
- Collisions create ‘holes’, secondaries => tracked
  - Bertini-like: shells, ($E<10$GeV) wide use “BERT”
  - Binary QMD-like: static nucleons, pot. V, $E< 3$ (π: 1.5)
  - INCL++ Liege cascade: mature ‘first-principles’ cascade
- Remnant de-excited by pre-compound, evaporation, fission, ..

Cascades: Bertini

(Comparison to data used to generate these fits.)

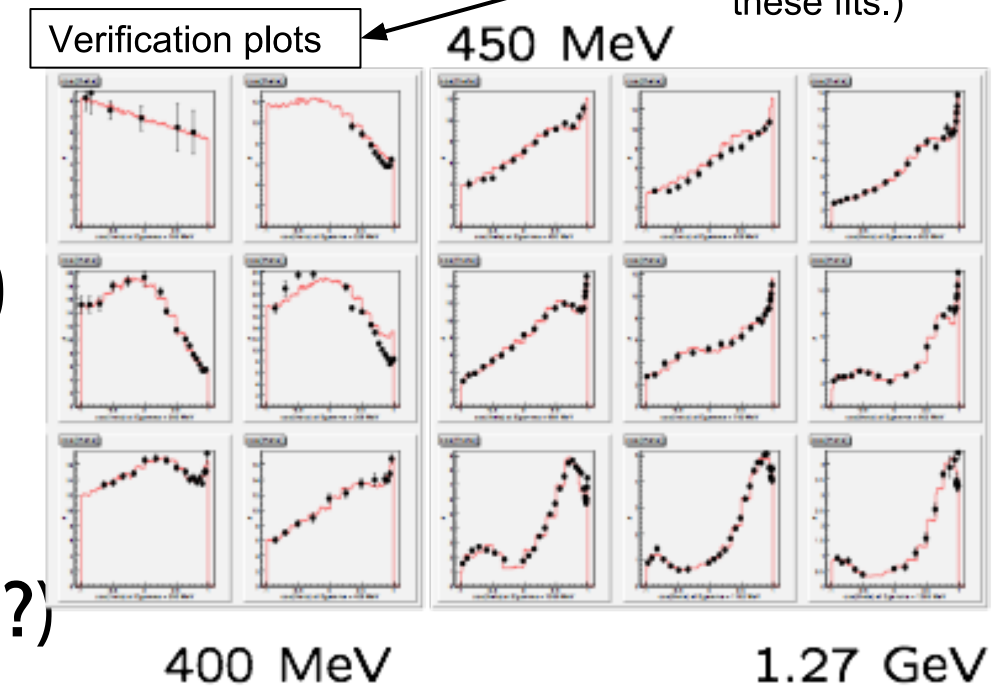

- New angular distributions for two-body final states
  - γ p, N-N elastic + pi-N (10.1)
- Added γ- and electro-nuclear interactions (9.6)
  - replace CHIPS model (10.0?)
- Improved multi-body final state (N>5)
  - Kopylov N-body phase space generation (optional)
- Leptons as projectiles
  - Stopping/capture of muons generate cascade
- Improvement of selection of first interaction

Cascades: INCL++

- INCL++ = Liege cascade INCL recast in C++
  - Full implementation of models from INCL 4.6
- Extended validity up to 10-15 GeV (G4 10.1-β)
- Improved cross-sections for small clusters
- Restored ABLA model for de-excitation

Isomers

- **Isomer levels** existed for all nuclei already
  - now enabled consistent handling in framework
- New: on-demand creation of isomers with short lifetime
  - Default: only states with lifetime > 1 ns are created
- Relevant processes create/handle isomers
  - Cascades and decay create excited residuals
  - Capture, absorption, quasi-elastic can excite nucleus
- Improvement of radioactive decay / database

Improved Radioactive decay

- Fix: entire decay chains occur with their correct timing
  - in the past if a short-lived nuclide was reached, all subsequent decays would happen instantaneously
- Database synchronization
  - New radioactive decay database (4.1)
  - New photon evaporation database (3.1)
  - Synchronized to have identical energy levels and lifetimes down to a lifetime of 1 ns
- Model adapted to new isomer creation scheme (previous slide)

Physics List Revision/Improvement

Old recommended LHC production physics list:
QGSP_BERT (ver 9.4, Dec2010 - patch4 Apr 2012)

EM

- Urban93 Multiple Scattering
- Bremstrahlung

Hadronic

- Elastic
  - CHIPS (H), Coherent (other)
- Inelastic (p,n,π)
  - Bertini cascade: 0 - 9.9 GeV
  - Parameterised (LEP): 9.5 - 25 GeV
  - String *QGS model* : 12GeV - 1TeV
  - Precompound & Evaporation
  - Anti-baryons: LEP
- Quasi-elastic model: CHIPS
- Stopping particles - CHIPS
- γ-nuclear, e-nuclear: CHIPS

New recommended LHC production physics list:
FTFP_BERT (Geant4 ver 10.0, Dec 2013)

EM

- Urban(96) Multiple Scattering
- Improved Bremstrahlung

Hadronic

- Elastic
  - CHIPS (H), Coherent (other)
- Inelastic (p,n,π)
  - Bertini cascade: 0 - 5 GeV
  - String *Fritiof/FTF model* 4GeV - 1TeV
    - Reggeon cascade (in FTF)
  - Improved Precompound (G4 v9.5)
  - Anti-baryons: FTF (new)
- Quasi-elastic model: FTF
- Stopping particles: Bertini (new)
- γ-nuclear, e-nuclear: Bertini (new)

Kernel

Geometry

Event biasing

Geometry

- New USolids shapes library
  - Implementations of all 'shapes' (start: G4, Root)
  - Goal: faster, robust and (next) vectorized
  - Future single library for Geant4, Root +
  - New: fast tesselated solid, multi-union (10.1β)
- See also talk by S.Wenzel on VecGeom, the Vector extension of USolids

Event biasing

- **PLACEHOLDER / DRAFT slide**
- New design for event biasing
    - Physics biasing
    - Angle biasing
    - Geometry/important biasing
- Flexible coupling of different types of biasing

Performance

Monitoring

Improvement

Absolute throughput (sequential)

Heavy developments: FTF becomes competitive with QGS

Fast implementations of log/exp

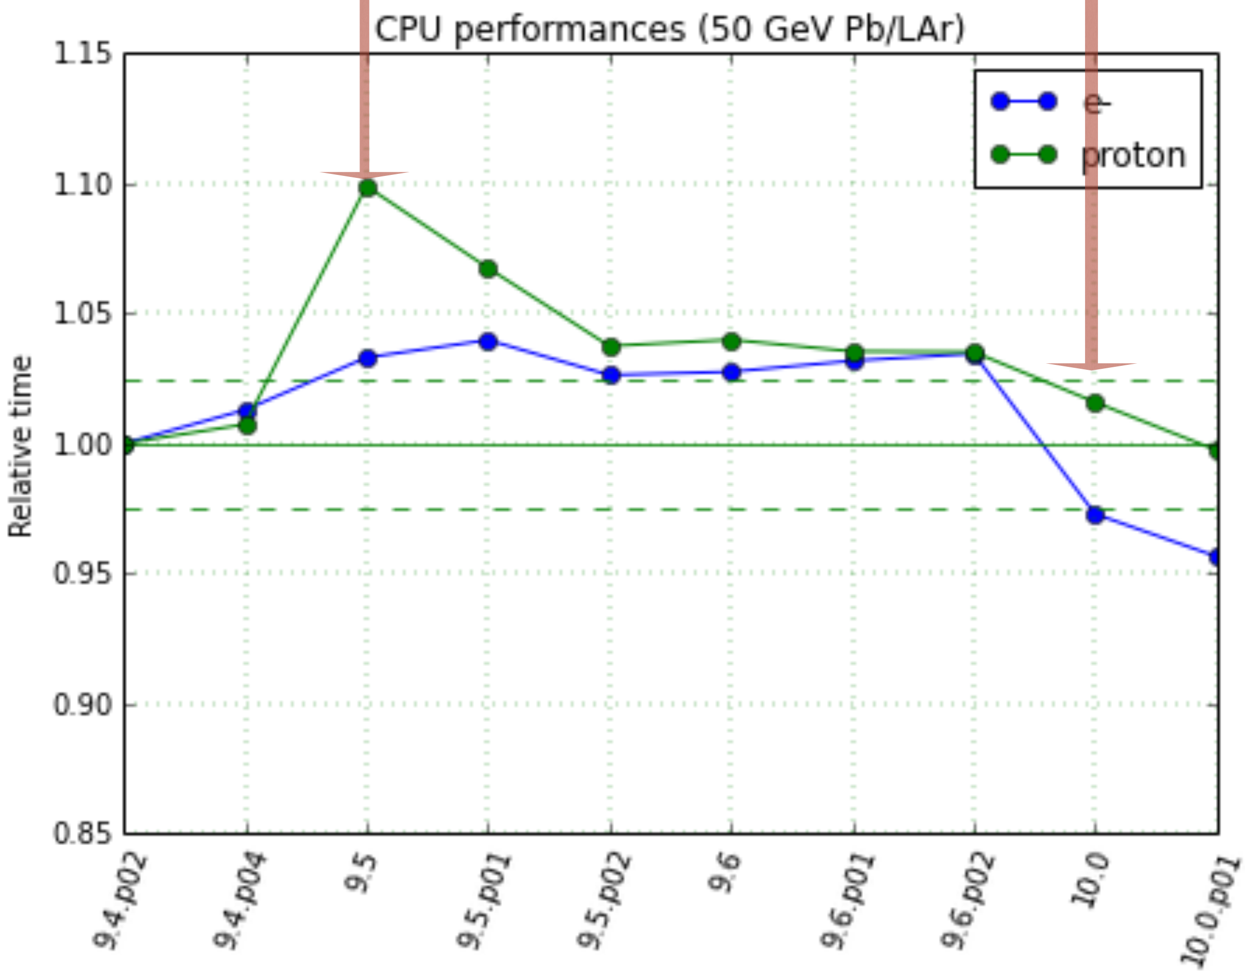

Improvements for MT brought CPU benefits also to sequential

Obtained with simple 'slab' geometry

Improving Performance

- Continuous monitoring of CPU time of benchmarks
  - Regular (monthly+) running
  - Simpler (slabs) and complex geometries
  -
- Repeated campaigns for CPU improvement
  - better models => increase time => review & improve
- A number of recent changes (gains of 3-5% each)
  - Used fast implementation of log/exp (from VDT)
  - Refactored eA, gA cross-sections (from per isotope to per element)
- Reviewed key EM process/model classes
  - small improvements from code refinements

Summary: 1.Multi-threading

- Achieved large memory re-use / savings:
  - Shared geometry, EM physics tables
  - Extra 38MB/thread (vs 200MB) in our test
- CPU Performance
  - Sequential performance mainained/improved
  - Excellent scaling to 40-60 threads (in the regime of 1 thread/core.)
- Simple porting of most applications
  - Some changes in 'user action' classes needed
-

Summary: 2.Physics

- AB
  - a
  - b
- C
  - c
  - cd
- D
  - e
  - f
- G