Geant4 Optimization Task Force Group Report

Mustafa Schmidt
on behalf of the ATLAS Simulation Group

February 2, 2023

BERGISCHE
UNIVERSITÄT
WUPPERTAL

Outline

Overview over validated and ongoing tasks:

Overview

Validated & Ongoing Tasks

Performance Comparison

We highly acknowledge the work of the Geant4 collaboration

G4 Optimization Task Force

Simulation Group – conveners : M. Bandieramonte, T. Lari

Full Simulation Convener : W. Hopkins	Fast Calorimeter Simulation (FCS/FastCaloGAN) Conveners : M.F. Giannelli, C. Young	Fast Chain Conveners : M. Javurkova, L. Mijovic	Simulation Quality Group Convener : B. Wynne

- Full Simulation Group evolution of the former "G4 Optimization Task Force"
- Close collaboration with G4 developers (especially CERN-EP-SFT) is implied
- Recently validated tasks:
  - **Woodcock Tracking**: Major impact on speed
  - Voxelization Optimization: Optimizing CPU time and memory
  - Monopole Bug: Crash for monopoles in G4
- Ongoing tasks: EMEC from standard Geant4/VecGeom Shapes, GMClash, ISF Particle Killer, Advanced Compiler Optimisation, In-Field Parameter Tuning

Woodcock Tracking

- Goal: Reduce simulation steps + CPU time without approximations
  - Very powerful in highly granular detectors (e.g. the EMEC) where geometric boundaries limit steps rather than interactions
  - Performs tracking in one material without boundaries (e.g. Pb - the densest)
  - Interaction probability proportional to the cross-section ratio between real material and Pb
  - Boundaries can be ignored ⇒ reduction of simulation steps

Investigated by Mihaly Novak & John Apostolakis

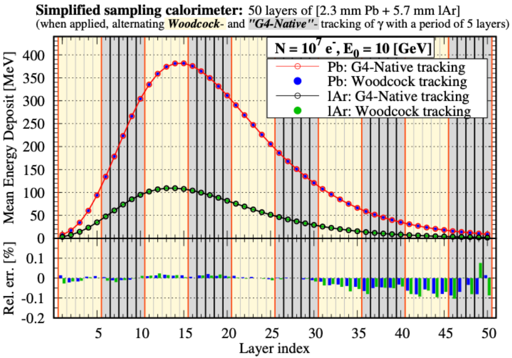

Woodcock Tracking

- Implemented as a modification to the G4GammaGeneralProcess - GammaGeneralProcess also reduces the cost of a step
- Physics validation completed
- Total speedup for Woodcock Tracking from Run3Opt
- Speedup from $\approx 33\%$ to $\approx 48\%$ (50% gain from Woodcock)

BNL Cluster, 1000 Jobs 100 tt-bar events/job	**Walltime,s**	**sigma**	**Speedup CPU time**	**Job ID**
Athena 22.0.47 (baseline)	288	23.2	-	27857092
Athena 22.0.67 "out the box" (BigLibrary only)	274	21.9	-4.9%	31366397
Athena 22.0.67 "out the box" (BigLibrary only) +Run3Opt	193	16.2	-33%	31366406
Athena 22.0.93 private build (BigLibrary + VecGeom)	257	20.9	-10.8%	31548908
Athena 22.0.93 private build (BigLibrary + VecGeom)+ **+WoodcockTrack&G4GammaGeneralProc**	198	20.2	-31.2%	31548954
Athena 22.0.93 private build (BigLibrary + VecGeom)+ **+WoodcockTrack&G4GammaGeneralProc+Run3optimizations**	**150**	**15.9**	**-47.9%**	31544224

BERGISCHE UNIVERSITÄT WUPPERTAL

Voxel Density Optimization

Investigated by Mustafa Schmidt

- Goal: Find the optimal values of voxel density in Run-3 geometry for optimization of CPU time and memory consumption
  - Size/Granularity of the voxels can be tuned by users
  - Voxel density - member variable in Geant4 logical volume class
  - Changes studied in FullSimLight and Athena
  - Improvement in memory consumption for geometry optimizations

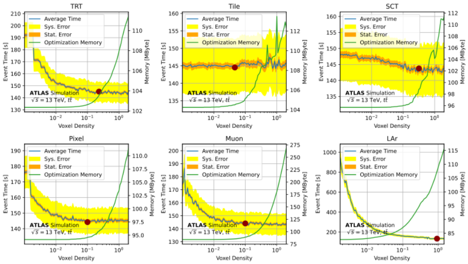

- Physics validation completed
- Follow-up study is ongoing (increase of CPU observed)

Fixing Monopole Bug

Investigated by John Chapman, John Apostolakis, and Mustafa Schmidt

- Goal: Fix problems appearing during highly ionizing particles (HIP) trigger simulations
  - Problems related to the transportation of monopoles
  - Monopole tracking was not possible anymore because of using more than one field manager
  - All jobs failed
- Problem Resolved: ATLASSIM-5960, G4Extensions Monopole Repo

EMEC Geant4 Shapes

Under Investigation by Akanksha Vishwakarma and Evgueni Tcherniaev

- Work was paused for some time and Evgueni took over around November 2022
- Goal: Define geometry of EMEC with Standard G4 shapes to speed up simulation as well as allow usage in other architecture (GPU prototype currently being developed by Benjamin Morgan)
  - EMEC: Computationally dominant part of full G4 simulation
  - Custom solid is not portable to GPU - many ongoing efforts to provide GPU accelerated particle transport in the community (HPC)
  - No accordion shape available within G4 standard geometry shapes $\Rightarrow$ Defined custom solid geometry according to "neutral fibre"
  - Use trapezoids (standard G4 shapes) for the EMEC geometry description.

EMEC Geant4 Shapes

Comparison between G4Trap and G4TwistedTrap (Benchmark Run: 10000 events with 10 GeV)

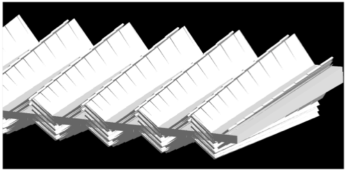

G4Trap: More replicas than twistedTrap to address the twist

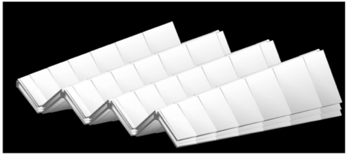

G4TwistedTrap

Particles	Time (s)	CPU usage (%)
e-	209.36	99.3
e+	209.46	99.3
gamma	197.60	96.3
geantino	0.402	99.4

Particles	Time (s)	CPU usage (%)
e-	3882.28	98.4
e+	4027.08	97.9
gamma	3170.14	98.3

Very preliminary benchmarks (proper material & angle change missing) show preferences to G4Trap

GMClash

Goal: Improve gmclash performances with faster random generator and changing navigation algorithm

- Under Investigation by Evgueni Tcherniaev
- gmclash: a MC-based package for detecting anomalies (clashes) in geometry trees
- Study comparing new version of gmclash with the old version, running with Geant4.10.06 and Geant4.11.0

	Current gmclash (N. of points = 1.000)		New gmclash (N. of points = 1.000)	
	Geant4.10.06 M / S / E / I	Geant4.11.00 M / S / E / I	Geant4.10.06 M / S / E / I	Geant4.11.00 M / S / E / I
ATLAS-R3S-2021-03-00-00.db	1468/399/0/136 Tot = 2003 CPU = 17':18''	1485/407/0/136 Tot = 2028 CPU = 13':20''	1501/357/1/136 Tot = 1995 CPU = 11':43''	1465/356/1/136 Tot = 1958 CPU = 9':15''
ATLAS-R3S-2021-03-00-00.gdml	3991/1872/0/120 Tot = 5983 CPU = 16':20''	3986/1879/0/120 Tot = 5985 CPU = 12':28''	2857/1809/2/120 Tot = 4788 CPU = 6':48''	2860/1824/3/120 Tot = 4807 CPU = 4':45''

- Generating random points (default 1k) on surfaces of volumes → position of points checked relative to other volume
- Types of Clashes: Clashes with a mother volume (M), Clashes with sister volume (S), Clashes with entirely encapsulated volume (E), Clashes with invalid solid (I)
- Confirmed problem with how GDML dumps specific volumes (especially displaced solids)
- Request: Support for displaced solids in Geant4 GDML implementation

Link: https://geomodel.web.cern.ch/home/fullsimlight/gmclash/

ISF Particle Killer

- Goal: Kill primary particles generating secondaries close to the beam pipe at 5 – 6 m
  - Huge number of secondaries being produced 5–6 m away from IP (0,0,0) with small r (close to beam pipe)
  - Many of these secondaries will not cause any energy deposits in the calorimeters or a muon hit
  - Primary particles causing interactions could be dropped directly
- Approach:
  1. generate a large sample of single particles with $4.5 < |\eta| < 6$ with different energies
  2. Map out η and E combinations producing relevant signal
  3. Drop others directly with new particle killer

Under Investigation by Dongwon Kim and Michael Duehrssen

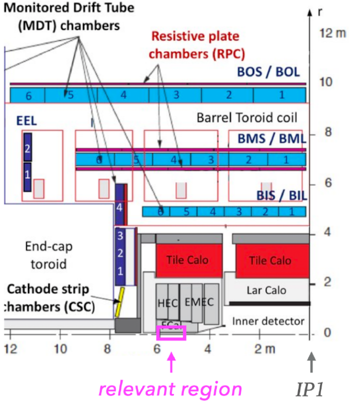

In-Field Parameter Tuning

Under Investigation by Nitika Sangwan

- Goal: Find the optimal values of the in-field tracking parameters for physics performance and CPU savings
- Lists of tuning parameters + descriptions for tracking in a magnetic field
  - DeltaIntersection: accuracy of intersection with boundary volume
  - DeltaOneStep: accuracy for endpoint of 'ordinary' integration step
  - DeltaChord: approximation of curve with linear sections
  - MaxStep: maximum step length
  - Tuning can be done for different detector regions and various particle energies

Advanced Compiler Optimization

Under Investigation by Benjamin Morgan & Caterina Macron

- Goal: Speed up simulation with link time and profile-guided optimization (LTO & PGO)
- Advanced compiler optimizations can lead to non-negligible speed-up factors
- CMS report: $\approx$ 10% speed in their software (https://indico.cern.ch/event/394788/contributions/2357347/attachments/1368686/2074705/slides.pdf)
- Two approaches to reducing application runtime, both relying on the compiler smarter usage of compiler $\rightarrow$ more throughput $\rightarrow$ efficient use of computing resources
- LTO: instrument compilation units with metadata
  - Consults to optimize when building shared objects
  - Expands scope of inter-procedural optimizations to encompass all objects visible at link time
- PGO: lets you optimize a whole executable (optimizer uses data from test runs)
  - build instrumented binaries, produce a profile for the application, rebuild from sources and profile
  - inlining, block ordering, register allocation, conditional branch optimisation, virtual call speculations, etc

Performance Comparison

- MC21
  - Geant4 version for mc21: Geant4 10.6.patch03.atlas03
  - mc21: $\geq$ 30% CPU speedup w.r.t. mc16
- MC23
  - Geant4 version for mc23: Geant4 10.6.patch03.atlas04
  - mc23: $\approx$ 2 times faster than release 22 (without optimization)!
  - Colored boxes in table not using exactly same job settings in plots but convey same message
  - Observation in CPU time speed-up!

BNL Cluster, 1000 Jobs 100 tt-bar events/job	Walltime,s	sigma	Speedup CPU time	Job ID
Athena 22.0.47 (baseline)	288	23.2	-	27857092
Athena 22.0.67 "out the box" (BigLibrary only)	274	21.9	-4.9%	31366397
Athena 22.0.67 "out the box" (BigLibrary only) +Run3Opt	193	16.2	-33%	31366406
Athena 22.0.93 private build (BigLibrary + VecGeom)	257	20.9	-10.8%	31548908
Athena 22.0.93 private build (BigLibrary + VecGeom)+ **+WoodcockTrack&G4GammaGeneralProc**	198	20.2	-31.2%	31548954
Athena 22.0.93 private build (BigLibrary + VecGeom)+ **+WoodcockTrack&G4GammaGeneralProc+Run3optimizations**	**150**	**15.9**	**-47.9%**	31544224

Plot by Vangelis Kourlitis and Marilena Bandieramonte

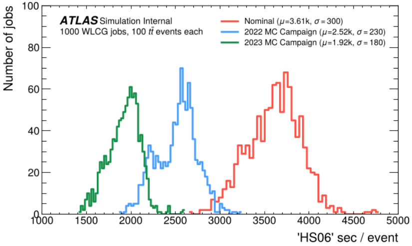

HEC Simulations

Investigated by Lorenzo Pezzoti

- Geant4 validation project exploiting test beams ongoing including ATLAS HEC+TileCal
- HEC simulation first one studied and successfully ported into Geant4 standalone simulation within the Geant Val testing suite
- Most observables (shower length, energy resolution, shower shape) in good agreement with test beam data
- Reproduction of ATLAS results (some cross-validation still possible and under disposal) for their mentorship

Other studies (Sven Menke): Neutron cross section changed significantly from Geant4 10.1 to 10.7 (most likely not improvement)

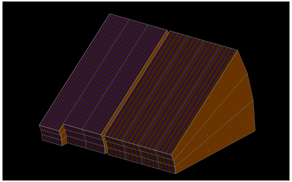

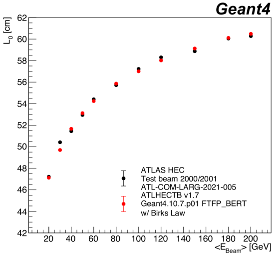

Link: `https://www.mdpi.com/2410-390X/6/3/41`

Investigated by Lorenzo Pezzoti

- Developed new Geant4-based simulation of the 2017 ATLAS TileCal test-beam
- Features all main ingredients for a realistic simulation (RO cell description, Birks' Law, U-shape correction, PMT emulation, and energy calibration) without any ATHENA dependency
- Result: Jet energy resolution improved since 10.4 (better than data)
- Currently Geant4 reproduces π/e results with great accuracy but investigations needed for better description of the response fluctuations

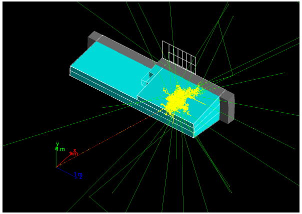

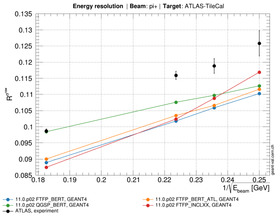

Link: https://www.mdpi.com/2410-390X/6/3/41

- Significant difference between data and MC simulations when modeling shower shapes in calorimeters
- Geant4-10.6 moved away from data compared to Geant4-10.1
- Solution for ATLAS: Switching off Urban-MSC correction for positrons in new release
- Default EMZ physics list providing results closer to data (but much slower)
- Most likely no further improvement in Geant4 physics possible (More studies in ATLAS required)

Under Investigation by Mihaly Novak

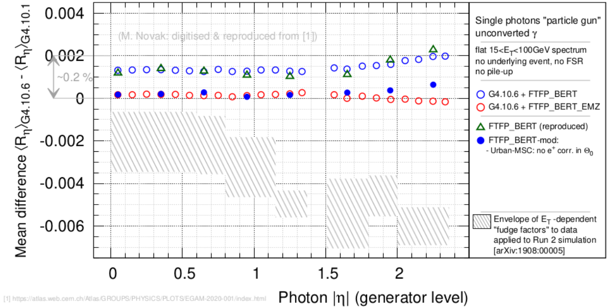

Thank you very much for your attention!