Geant4 in ATLAS Report

Geant4 Technical Forum

21st January 2021

Marilena Bandieramonte (University of Pittsburgh)

On behalf of the ATLAS collaboration

Computing complexity challenge

- The upgrade to the **HL-LHC for Run 4** produces a step change for **ATLAS**.
- The beam intensity will rise substantially, giving bunch crossings where the number of discrete proton-proton interactions (pileup) will rise **to about 200**, from about 35 today (2018 and foreseen for 2022)
- Accurate simulations and larger Monte Carlo samples will be needed to achieve the desired precision in physics measurements, while avoiding that simulation dominates the systematic uncertainties
    - ATLAS plans different R&D lines to reduce the need for *detailed full simulation*

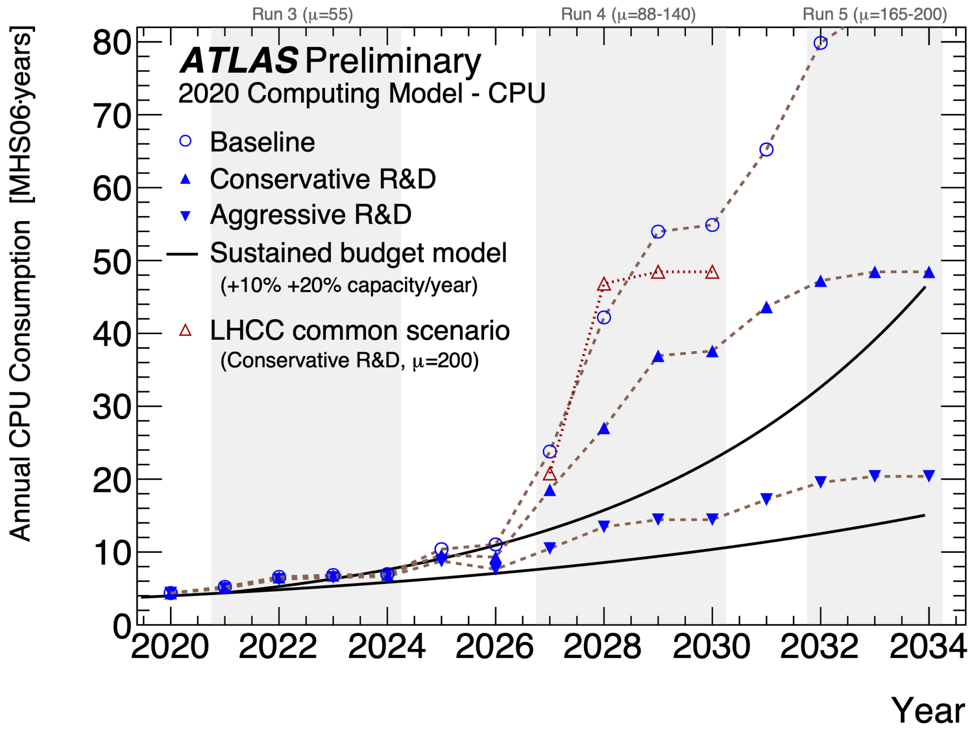

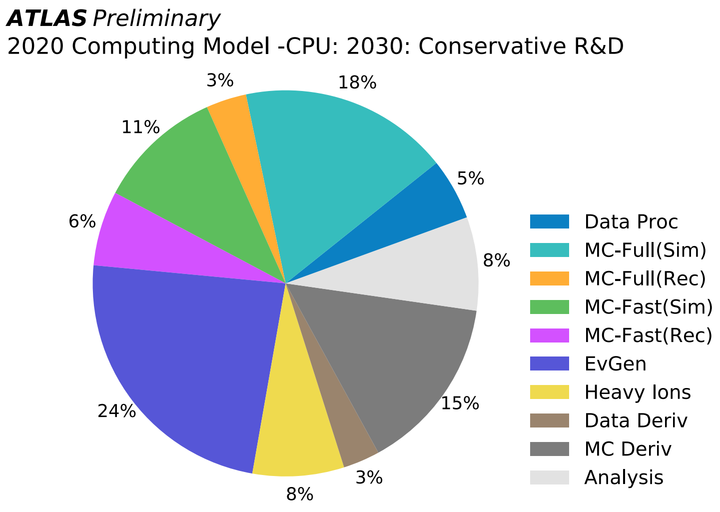

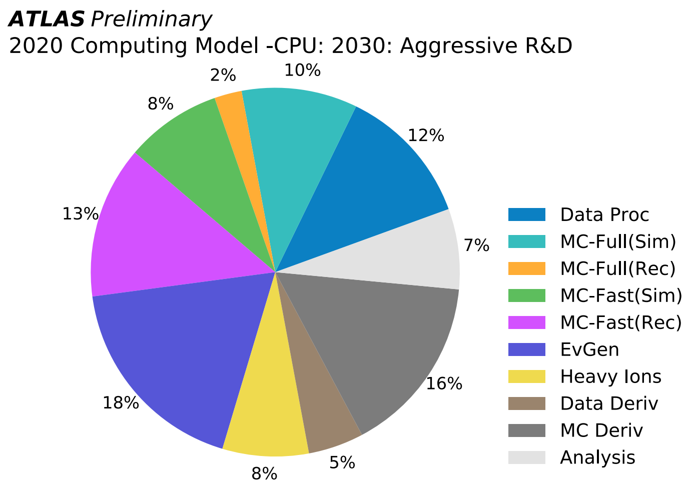

ATLAS Public plots

New Geant4 version for Run3

- We are testing different Geant4 versions to decide which will be the default for RUN3. Likely:
  - Geant4.10.6
  - Geant4.10.7
- **Geant4.10.6.patch03.atlas01** in the process of become the default version in Athena master
- Plan to have a local build with Geant4.10.7 as well
- Birk's constant tuning
  - Effects still to be fully understood (comparison with default MC16 Geant4 version)

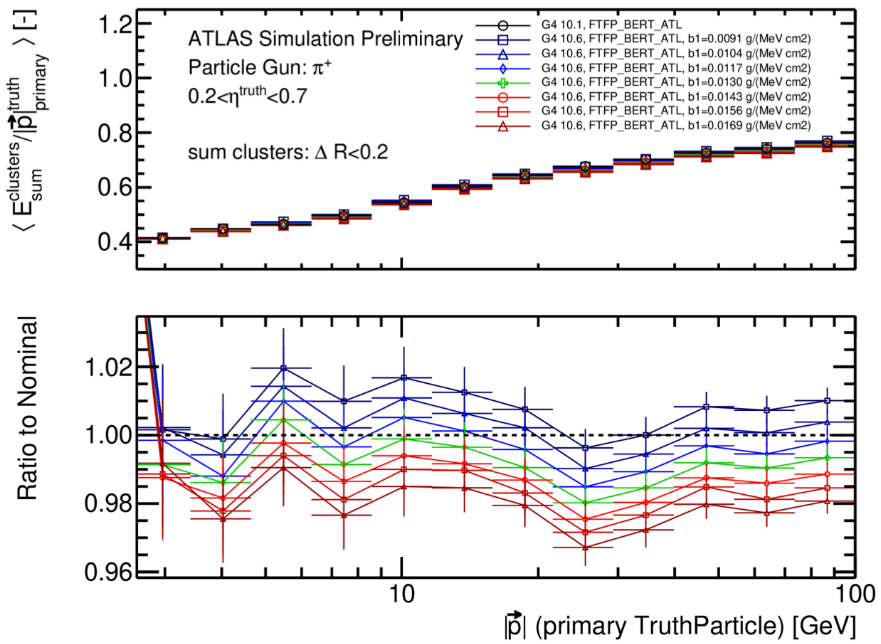

FTFP_BERT_ATL Physics list

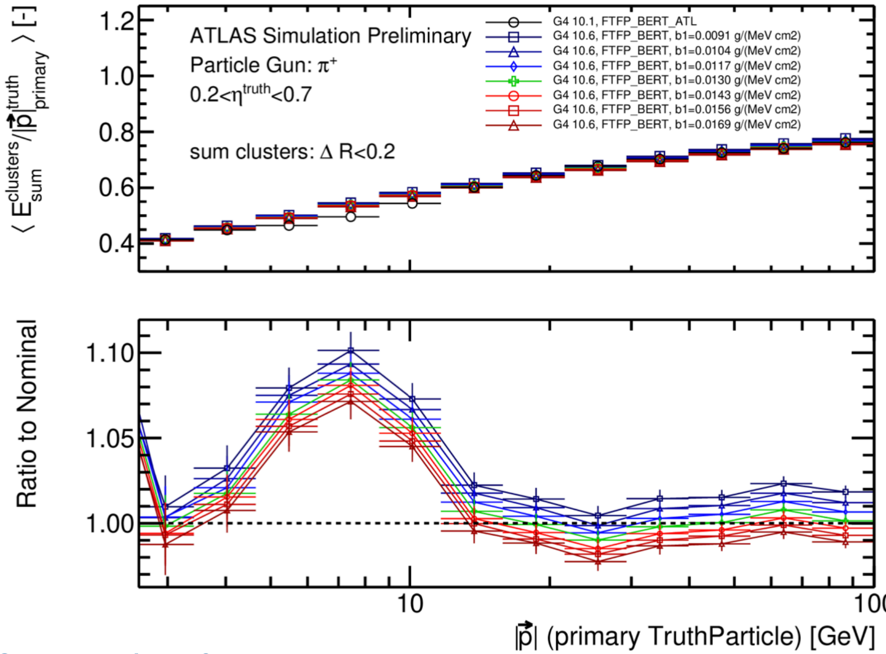

FTFP_BERT Physics List

b1 = 0.0130 g/(MeV.cm2) as default value of Birk's Law in Tile

b1 variation of ±10% ±20% ±30%

central Barrel: 0.2 < ηtruth < 0.7

New Geant4 Optimization Task Force

- The Geant4 (G4) Optimisation Task Force (TF) is responsible for optimising the performance of the ATLAS G4 simulation software:
  - investigating configuration options and simplified geometries and magnetic-field descriptions
  - improving the ATLAS interface code to G4.

- The TF's one-year mandate is to achieve for Run 3 >30% CPU performance improvements compared to the comparable Run-2 simulation

- Kick-off meeting was held on 1 September 2020.

- Some of the activities:
  - Taking advantage of **intrinsic performance optimizations** coming with newer Geant4 versions
  - Optimization with **tuning of G4 parameters** (physics models, physics lists per regions)
  - **Neutron** and **Photon Russian Roulette** + **EM range cuts** (ongoing physics validation)
  - **Geometry** optimisations (new EMEC variants + R&D on ML guided steppers in geometry)
  - **Magnetic field** tailored switch-off
  - **Geant4 linking** as static library (a.k.a. Big library)
  - Explore **machine learning options** especially for **simulations optimization:**
    - Machine learning solutions to optimise the detector simulation and optimally tune/re-weight parameters (i.e. physics models, physics lists per regions, range cuts, magnetic field)

Approximate timeline

- The new G4 Optimization task force main goal is to speedup ATLAS Geant4 simulation by >30%
  - By the start of RUN3 MC campaign (End of Sept 2021)
- Expected speedup from ongoing optimization activities with respective risk-level:

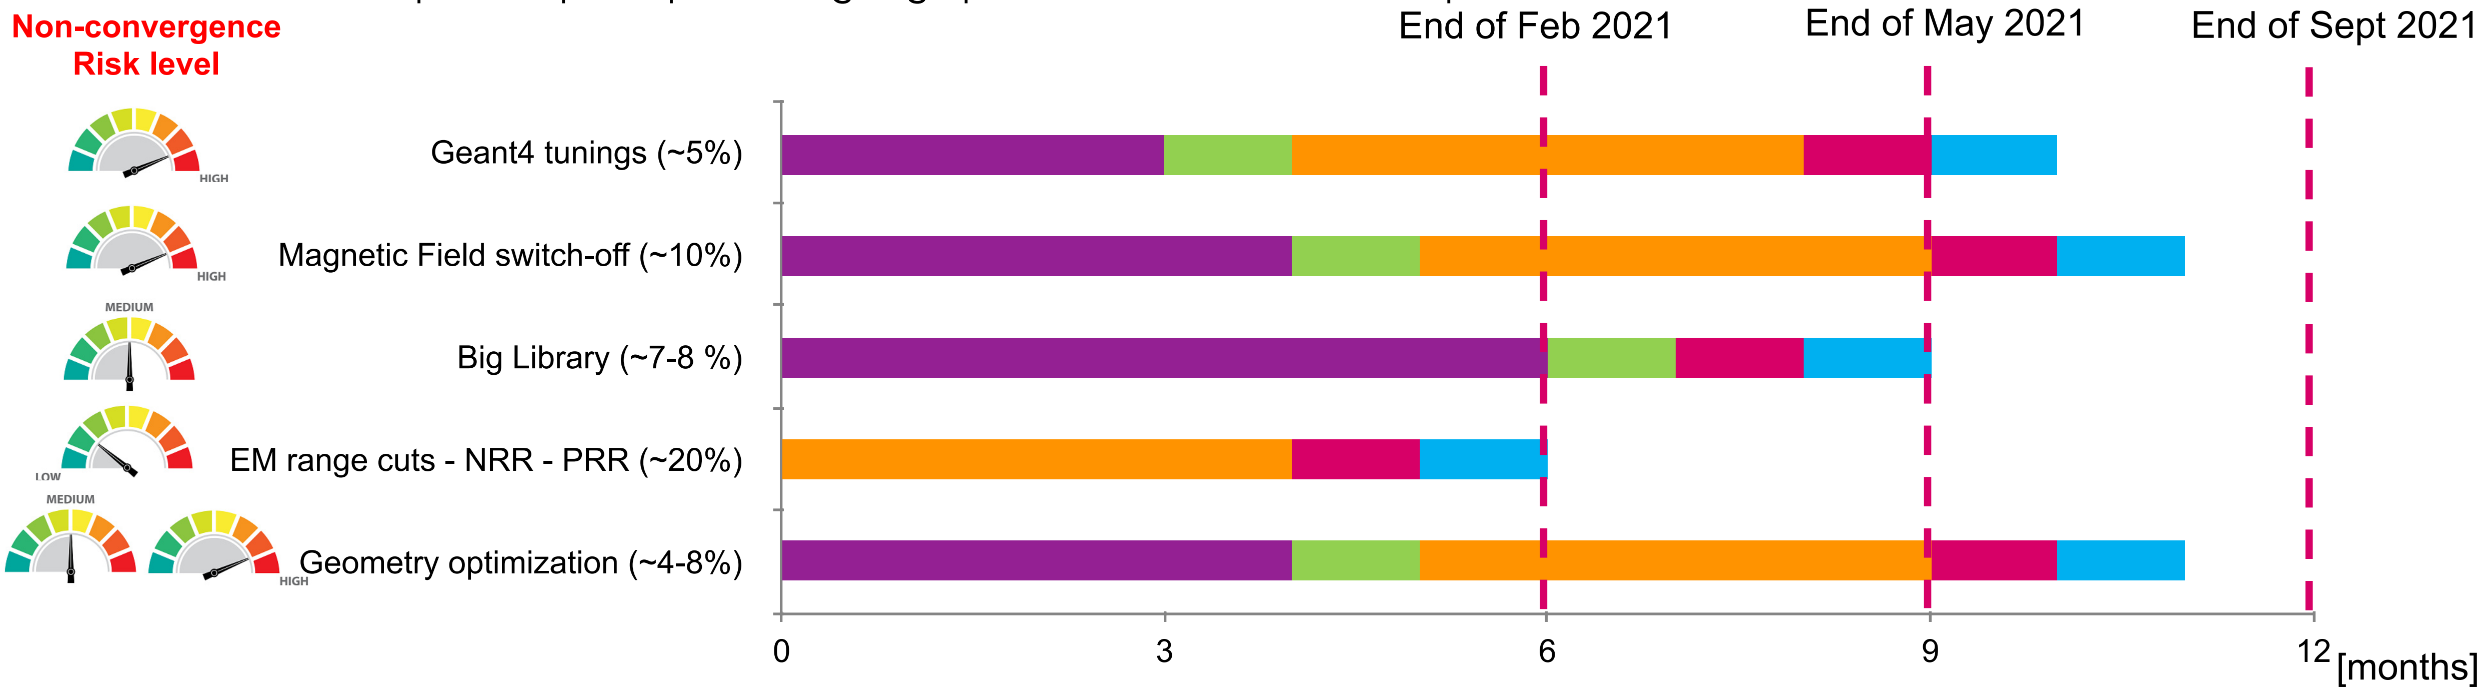

First implementation

Some features/parameters tuning might be easily and quickly tried with FullSimLight

Local validation test

Every technical change/new implementation must be locally tested before undergoing a physics validation campaign

Physics Validation

The physics validation most probably needs more than one iteration (~4 months)

Readiness for production

Once the physics validation is signed off, need to implement the feature and make it tested and ready for production

Final tests in production

Last step will be to run some tests once in production, before starting the MC campaign

FullSimLight^: a light full simulation prototype

Detector construction

The ATLAS detector can be built from the persistified ATLAS geometry*

Or plug your own geometry!

Magnetic Field

The ATLAS magnetic field can be used if the map is available/ or a uniform mag field can be set

Primaries Generator

Primary generators with a particle gun** or with Pythia

NEW!

Scoring

A set of observables***is collected during the simulation per primary-particle

Physics List

The Physics List can be specified as an input argument with a flag (e.g.FTFP_BERT)

^Standalone project, available here
Geometry and Mag Field can be Shared with G4 developers

*Including the new EMEC variants:
- Wheel
- Cone
- Slices

**# Number of primaries per event (default [1, 10])
Primary particle energy (default [1 GeV, 100 GeV])
Primary particle direction (default isotropic distribution)
Primary particle type (currently e-, e+ and gamma, geantino)

***Mean energy deposit
Mean charged and neutral step lengths
Mean number of steps made by charged and neutral particles
Mean number of secondary e-, e+ and gamma particles

Intrinsic Geant4 10.6 performance improvements

EM physics reduced number of calls to calculation of log (particle_kineticEn):

- Run with ***FullSimLight***
  - Geant4.10.6.p00.atlas01/Geant4.10.5.p00.atlas01
  - full ATLAS geometry
  - constant MagField (4 Tesla) or ATLAS MagField (AtlasRK4 with NEW way of setting the stepper)
  - 10k mixed 10 GeV primaries (e+,e-, gamma), isotropic distribution
- Perf stat with 10 runs
  - Speedup with constant magnetic field **~ 7.86%**
  - Speedup with ATLAS magnetic field **~5.04%**

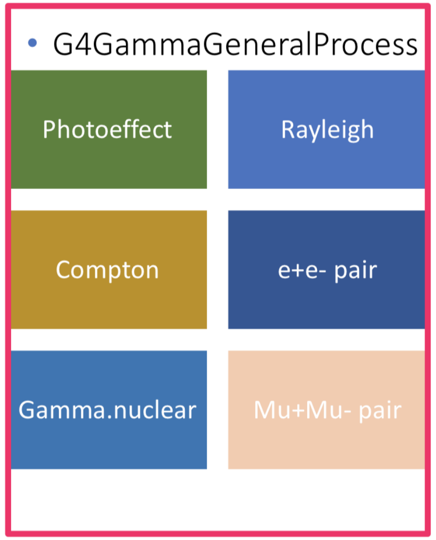

General Gamma Process:

- Run with ***FullSimLight*** + Geant4.10.6.2 + full ATLAS geometry + constant MagField (4 Tesla)
  - 10k mixed 10 GeV primaries (e+,e-, gamma), isotropic distribution
  - Perf stat with 10 runs: Speedup **~6.54%**
- Run with **Athena**, private build with geant4.10.6
  - 100 tt-bar events
  - Perf stat with 10 runs: speedup **~4.35%**

General Gamma Process OFF

```
INFO Statistics for 'evt': (nbr entries = 99)
INFO <cpu>:        ( 390761.616 +/-  14102.411 ) ms
INFO <cpu_user>:   ( 390358.081 +/-  14086.755 ) ms
INFO <cpu_sys>:    (    403.535 +/-     28.740 ) ms
INFO <real>:       ( 390805.980 +/-  14102.045 ) ms
INFO <vmem>:       (   2843.980 +/-      0.413 ) MB
INFO <malloc>:     (      0.000 +/-      0.000 ) MB
INFO <nalloc>:     (      0.000 +/-      0.000 ) calls
INFO <rt>:         ( 390762.326 +/-  14102.472 ) ns
```

General Gamma Process ON

```
INFO Statistics for 'evt': (nbr entries = 99)
INFO <cpu>:        ( 373643.475 +/-  13356.019 ) ms
INFO <cpu_user>:   ( 373380.545 +/-  13352.104 ) ms
INFO <cpu_sys>:    (    262.929 +/-     35.477 ) ms
INFO <real>:       ( 374230.626 +/-  13374.682 ) ms
INFO <vmem>:       (   2848.944 +/-      0.438 ) MB
INFO <malloc>:     (      0.000 +/-      0.000 ) MB
INFO <nalloc>:     (      0.000 +/-      0.000 ) calls
INFO <rt>:         ( 373643.823 +/-  13356.004 ) ns
```

Benchmarks new/old way of setting stepper/driver

- Run with **FullSimLight** + full ATLAS geometry + ATLAS MagField
- Geant4 versions:
  - **Geant4.10.1.patch03.atlas07: MC16 default production version**
  - Geant4.10.4.patch03.atlas01: currently in master
  - Geant4.10.5.patch01.atlas01
  - Geant4.10.6.patch00.atlas01
- NEW and OLD way of setting the AtlasRK4 stepper
- 10k mixed 10 GeV primaries (e+,e-, gamma), isotropic distribution
- Perf stat with 10 runs

Geant4 version	OLD WAY	NEW WAY	Speedup NEW vs OLD	Speedup OLD way	Speedup NEW way
g4.10.1p03.atl07	5669.88+- 50.35 sec				
g4.10.4.p03.atl01	5143.27+-19.52 sec.	5303.06+- 24.06 sec.	+3.10%	10.5 vs 10.4 = +3.04%	10.5 vs 10.4 = -0.7%
g4.10.5.p00.atl01	5299.94 +- 30.54 sec.	5265.06 +- 21.36 sec.	-0.65%	10.6 vs 10.5 = -4.15%	10.6 vs 10.5 = -5.04%
g4.10.6.p00.atl01	5079.57 +- 29.25 sec.	4999.33 +- 23.73 sec.	-1.58%	10.6 vs 10.4 = -1.95%	10.6 vs 10.4 = -5.72%

New way in Geant4.10.6 vs **old way** in Geant4.10.4 : **-2.8%**

New way in Geant4.10.6 vs **old way** in Geant4.10.1.patch03.atlas07 : **-11.82%**

Steppers performances across G4 versions

- Comparison of old/new way of setting different steppers across different Geant4 versions

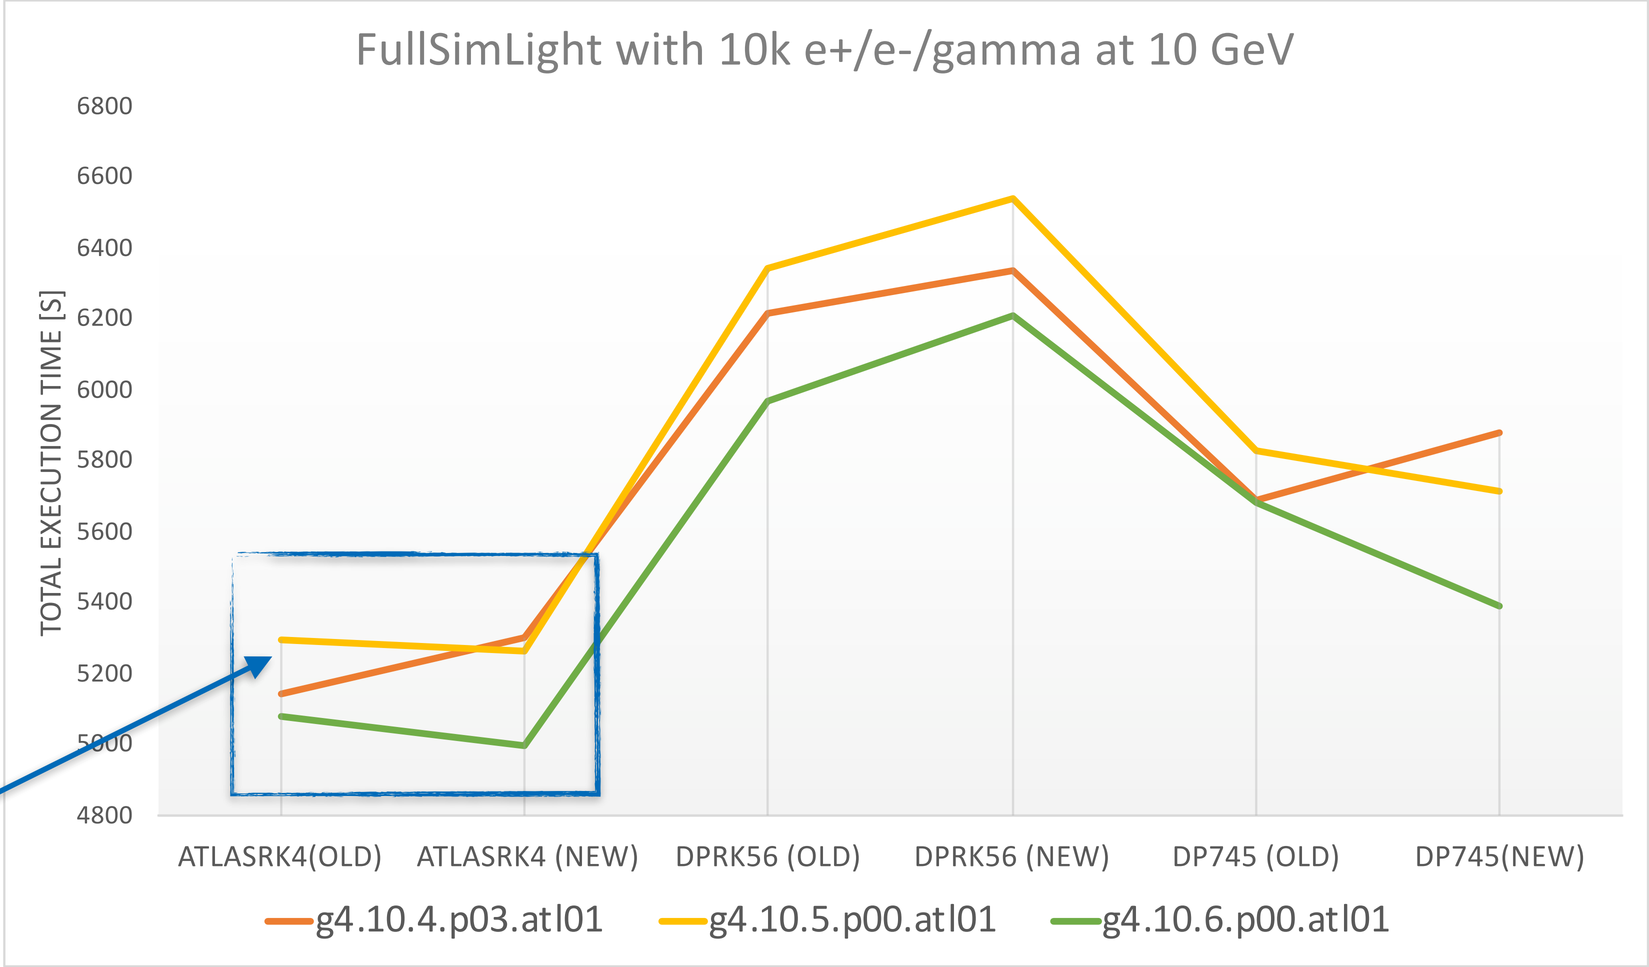

ATLASRK4 stepper seems to outperform the others

Magnetic Field tailored Switch-off

- Speedup on simulation (~10%) observed after switching off the magnetic field in LAr calorimeter.
  - Simply turning the field off in a volume was not successful.
  - Workaround: create a *"transition region"*, where the field smoothly decays until zero value deep inside the calorimeter
  - Problems with G4 exceptions raised by the G4Transport solved
  - So far, there seems to be no gain in terms of CPU time
    - The switch-off not correctly passed through the whole stack

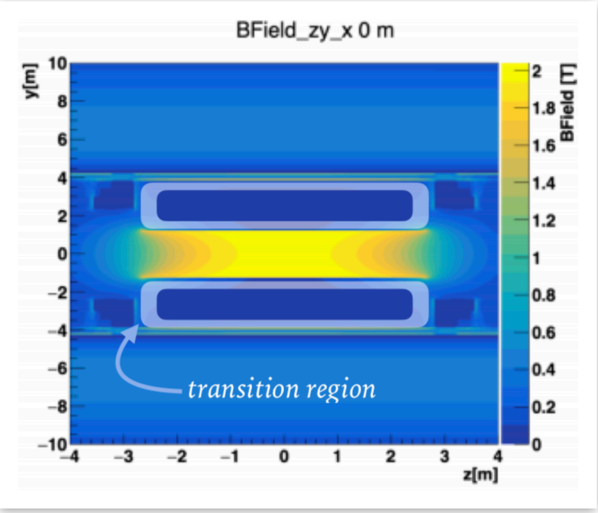

Magnetic Field tailored switch-off [1]

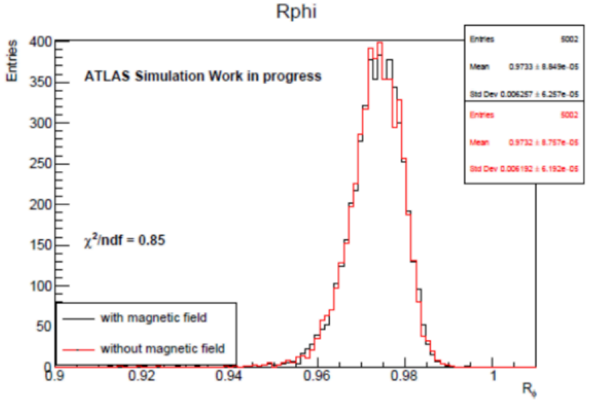

0.2 < |η| < 0.25

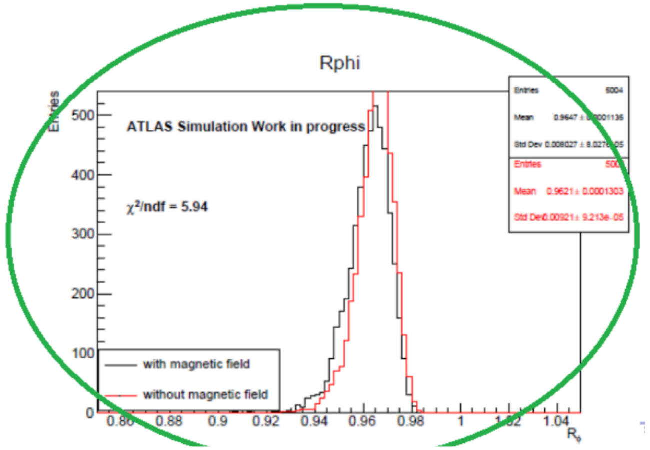

1.2 < |η| < 1.25

- The field in the barrel region is low and doesn't impact much the shower shapes (info):
  - Shower lateral width distributions of different calorimeter layers unchanged.
  - *Rphi* slightly affected in high eta.

Geometry Optimisations: TRT and EMEC

TRT

- Currently the TRT geometry is described using Boolean operations. This approach is not optimal as Boolean operations are slow.
- Idea: describe these volumes using alternative shapes:
  - extruded solid (e.g. BREP);
  - generic trapezoid (e.g. arb8);
  - tessellated solid.
- This can lead to:
  - significant reduction of the amount of code needed to describe the geometry (~ hundreds of lines) -> gain in terms of code maintenance;
  - gain in computational terms (to be investigated).

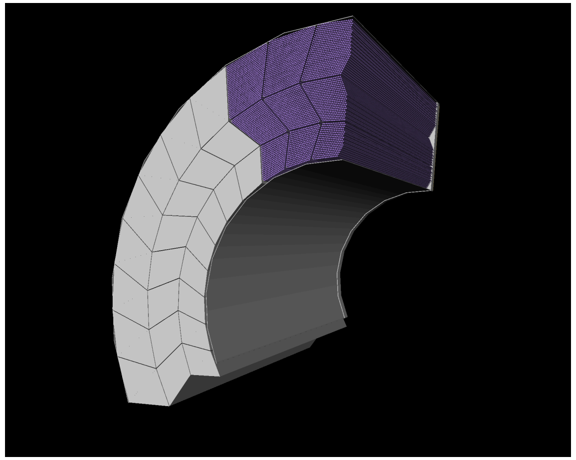

TRT barrel Geometry [2]

EMEC

- EMEC detector is implemented with as a custom shape:
  - Recent optimisation (new variants) gives ~ **5% speedup**
  - NEW Idea: implement it with std Geant4/VecGeom shapes/volumes
    - Twisted Trapezoids
    - Trapezoids
    - Tessellated Solids
  - Potential gain coming from
    - Use of standard shapes
    - Acceleration/vectorization (GPUs/VecGeom)

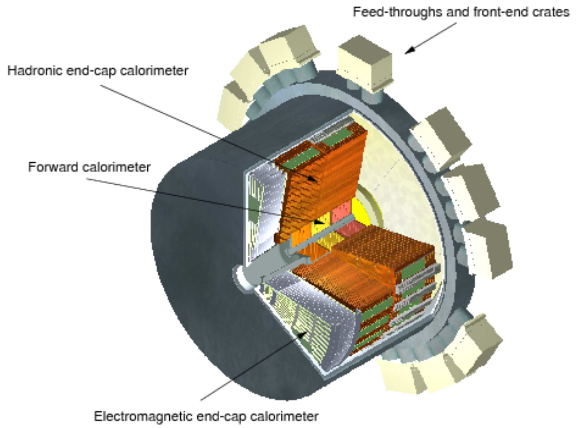

EMEC Geometry

Quasi-stable particles

- Some long-lived b-hadrons may travel far enough to interact with the detector before decaying, so some energy deposits may be missed.
- Observed difference between the generator level decay length and the decay length after G4 simulation for quasi stable particle simulation
- Calculation of the proper lifetime of the pre-defined decay used to be done with:

```
const auto& prodVtx = genpart->production_vertex()->position();
const auto& endVtx = genpart->end_vertex()->position();
const G4LorentzVector lv0( prodVtx.x(), prodVtx.y(), prodVtx.z(), prodVtx.t() );
const G4LorentzVector lv1( endVtx.x(), endVtx.y(), endVtx.z(), endVtx.t() );
double proptime=(lv1-lv0).mag()/Gaudi::Units::c_light;
```

- Along the chain the precision of the **vertex time and position** was reduced to ***float***, which caused the issue for highly boosted vertices
- If instead of using only the vertex position and time, the proper lifetime is calculated from the **3-distance of the vertices and the beta/gamma factors for the tau 4-momentum**, the difference in decay length between the HepMC record and the pre-defined decay in G4 essentially disappears.
- Suggestion/Request to add a new Geant4 method
  - `G4PrimaryParticle::SetProperTimeFromDetectorFrameDecayLength(double GeneratorDecayLength)`

BEFORE

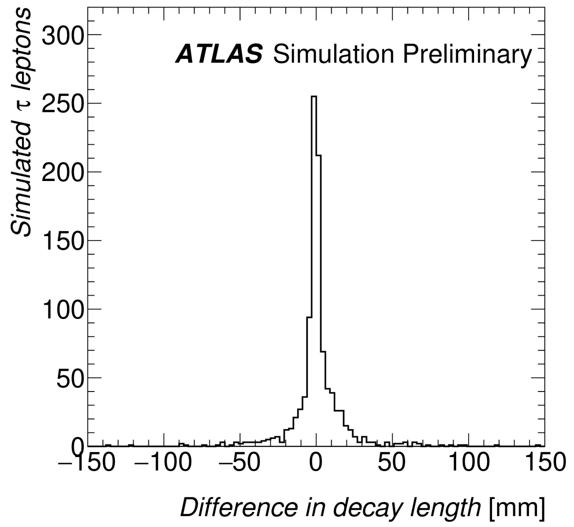

AFTER

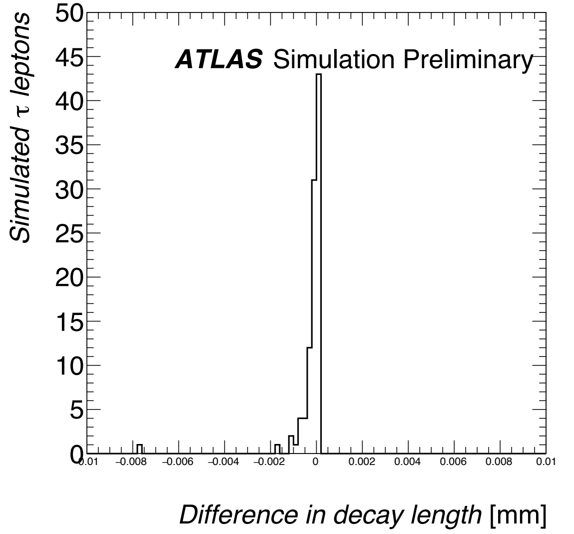

100 Z->tautau events

New ideas (I) : EMEC photon optimisation

- Photons don't interact during transport (no continuous energy deposition) and it seems that they are dominated by the transportation process
- 2 NEW IDEAS
  - Woodcock tracking:
    - Build a parallel geometry that shows all the EMEC as a large **volume** of **lead**
    - If a physics interaction happens for a photon:
      - check in which material the photon actually is
      - only accept this interaction with a probability proportional to the cross section ratio of the photon process in the "real" material compared to in lead
  - Kill photons that undergo to multiple transport steps (deposit energy locally). Similar to range cuts
    - Correct the accuracy with ML correction techniques in a highly parallel way

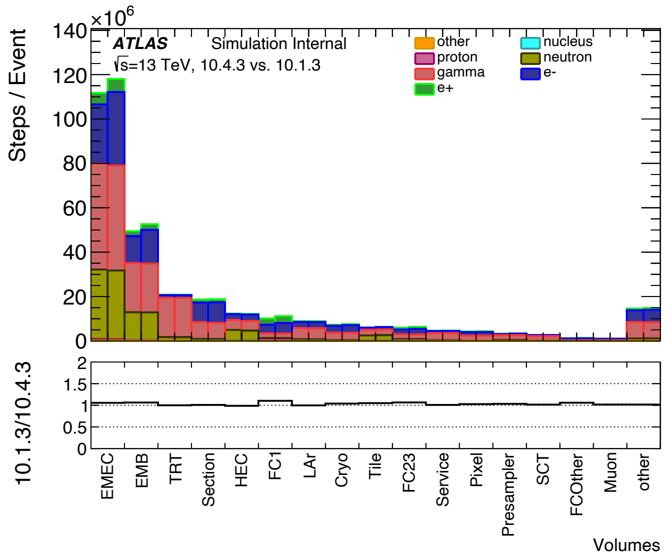

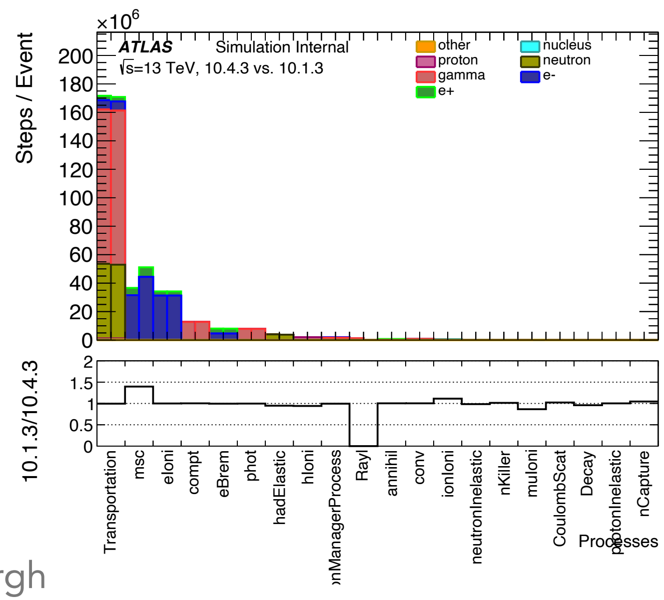

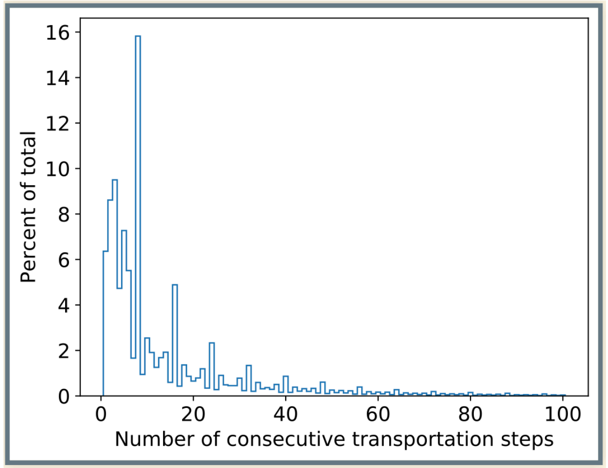

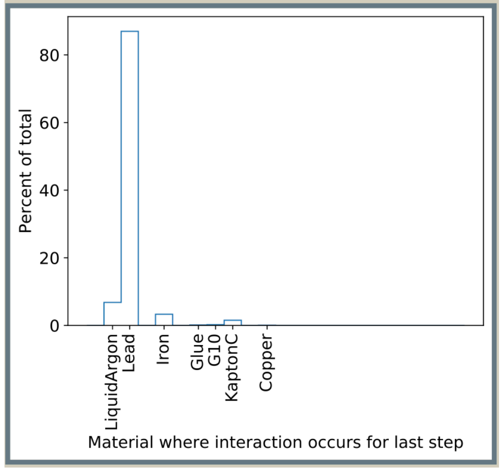

Plots [link, link]

New ideas (II) : new ISF particle filter

- Kill primary particles generating secondaries close to the beam-pipe at 5-6 m
- There is a huge amount of secondaries being created 5-6m away from (0,0,0) close to the beam-pipe
- Many of these secondaries will never cause any energy in the calo or a muon hit
    - the primary particles that caused these interactions could just be dropped directly
- It would be worth generating a large sample of single particles with eta value 4.5-6 and different energies and then map out which eta/energy combinations can produce a relevant signal and then drop the rest directly with a new ISF particle filter.
- We already kill all particles at eta>6
    - Particles at eta>5 and pT < 10 GeV?
    - Or/and particles at eta>4 and pT < 1 GeV?
- **To be investigated**

Conclusions

- **ATLAS** will likely adopt Geant4.10.6/10.7 for RUN3
- Speedups coming from intrinsic **Geant4 improvements:**
  - GammaGeneralProc from Geant4.10.6 (~4.35% speedup): **needs physics validation**
  - Reduced n. of log calls in EM physics (~5.04% speedup)
  - **Important:** New stepper/driver should not penalise CPU performance:
    - **New way** in Geant4.10.6 vs **old way** in Geant4.10.1.patch03.atlas07 : ~11.82% speedup (NB: this includes ~10% performance improvements across different G4 versions)
- **Geant4 simulation related optimisations**
  - **NRR+PRR and EM range cuts** (~20% speedup): **validation ongoing**
  - **Geometry**:
    - Ready to test the new EMEC variants implementations (~5% speedup): **needs physics validation**
    - Early stage for the other Geometry optimisations (TRT, EMEC with GPUs)
  - **Big Library** (~7% speedup):
    - Confirmed ~7% speedup with static linking and full ATLAS geometry
    - Currently running benchmarks in FullSimLight with full Geometry+ATLAS mag field
    - Next steps: implement within Athena
  - **Magnetic Field switch-off, ISF particle killer and Gamma transport in the EMEC:**
    - Very early stages

Thanks for your attention!

Marilena Bandieramonte
marilena.bandieramonte@cern.ch

Backup slides

EM Range cuts + Neutron Russian Roulette

- EM Range cuts + Neutron Russian Roulette
    - Physics Validation done for both separately: ATLPHYSVAL-603, ATLPHYSVAL-604
    - Physics Validation for both together almost done: ATLPHYSVAL-632 – can be used for RUN3
    - Gain of ~20% in total CPU time
- Lessons learned in validation:
    - Very important to suppress random fluctuations
    - Run without pileup to remove random fluctuations in pileup
    - Implement special Geant4 Stacking Action to simulate 'rouletted' particles last
- Photon Russian Roulette:
    - First tests show 3% speedup for a weight of 10 and Eth of 0.5 MeV and 10% speedup for 1 MeV,
    - Going beyond the annihilation peak at 0.511 MeV might have bad effects on physics performance
    - Validation was paused until we resolved issues in random fluctuations
    - Validation continuing now: ATLPHYSVAL-614
-

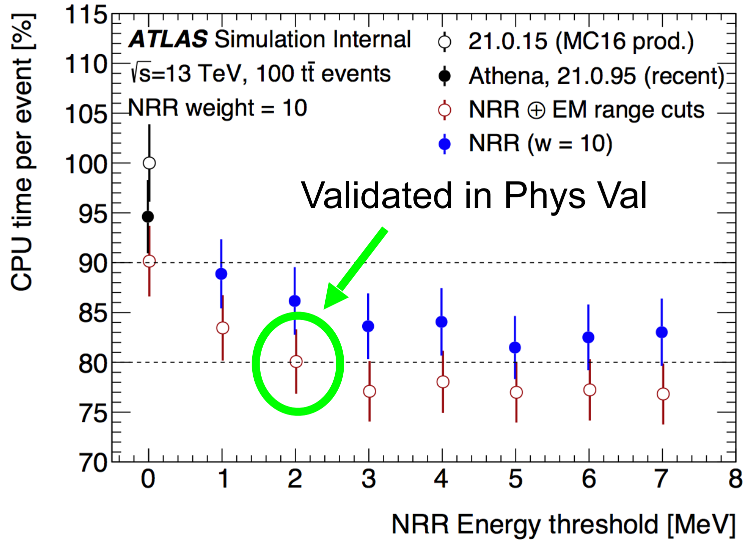

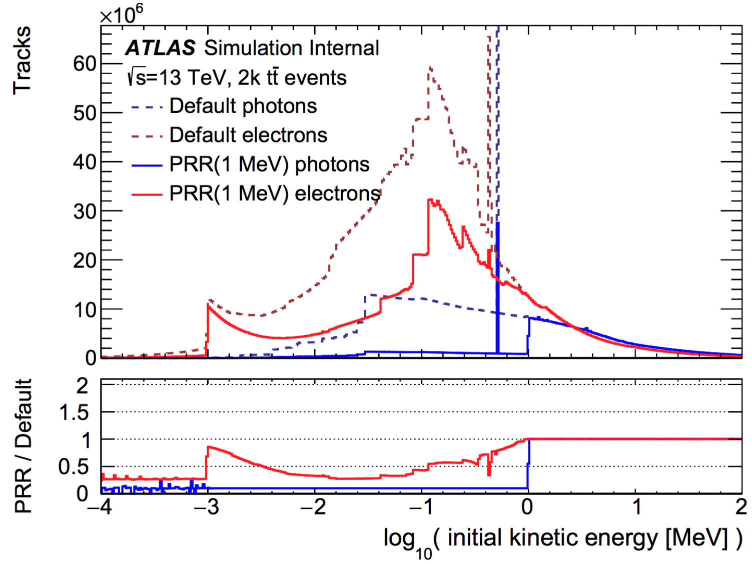

Improved EMEC geometry

Andrei Sukharev [talk]

- LArWheelSolid uses helper G4Polycone for some internal calculations.
  - The idea (issue ATLASSIM-3778): to replace these objects with a
    - simpler (and thus faster) shape (i.e. a cone).
    - To keep local coordinate system (z = 0 at front face) we need a
      - G4ShiftedCone — it was developed from standard G4Cons.
- Recently Implemented custom solid variants (selection is at job options level):
  - Wheel the default LArWheelSolid with G4Polycone
  - Cone improved LArWheelSolid with G4ShiftedCone — outer wheel divided into two conical-shaped sections
  - Slices new LArWheelSliceSolid — each wheel is divided into many thick slices along Z axis
- But also:
  - code cleanup: reorganize cycles, remove unnecessary calculations
  - optimized sincos calculation

- Physics Validation will start when rel22 sim+reco version will be validated(ATLPHYSVAL-651)
- Preliminary Performance test (thanks to Serhan Mete)
  - ISF_FullG4 default vs. new EMEC Slices variant
    - gave ~5-6% improvement on a quiet desktop machine
-

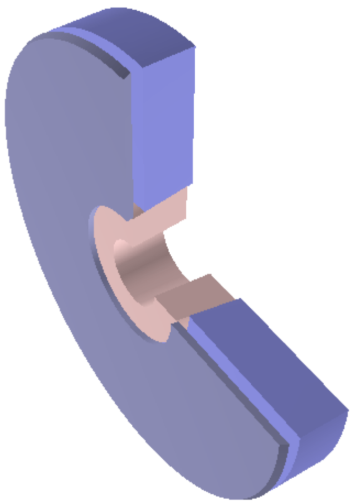

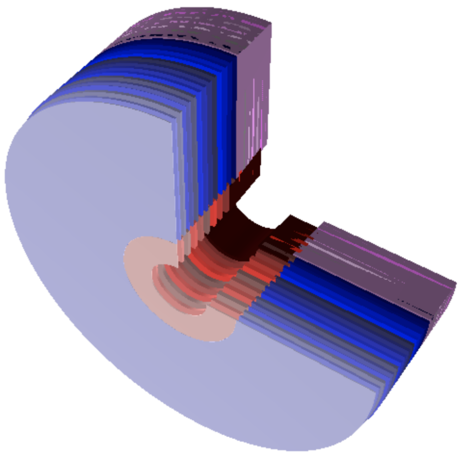

Geant4 as a static/single-dynamic library

Caterina Marcon & Ben Morgan

<u>Talk</u> at the kick-off meeting

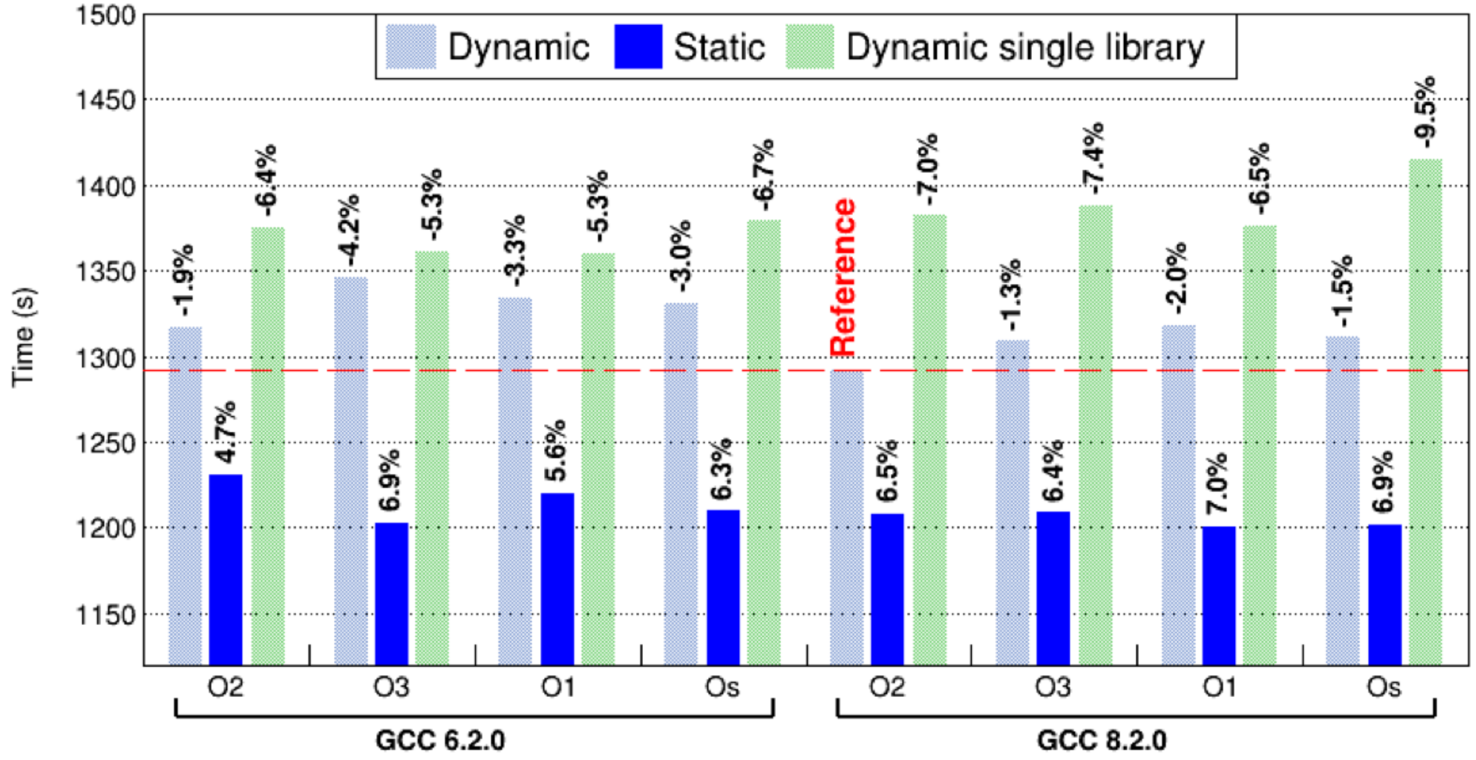

- Three different build types, default dynamic multi-library, dynamic single library and static library, have been used.
- 2500 initial events and 4 threads.
- Full GDML ATLAS model without EMEC.
- The dynamic single library approach, for both GCC versions, increases the execution time up to 10%.
- The dynamic single lib has been compiled with the PIC option enabled; previous studies, using non-PIC objects, found an improvement of 10% event/s -> under investigation.
- The dynamic single library is 998 MB, the default dynamic multi-library is 1007 MB and the static library is 1.8 GB.

Detector simulation requirements: ATLAS

- Very active **Geant4 Full Simulation Optimization** work ongoing
  - tackling RUN3 but also RUN4
- Taking advantage of **intrinsic performance optimizations** coming with newer Geant4 versions:
  - confirmed ~5% speedup coming from new GammaGeneralProcess + ~7% speedup btw Geant4.10.5 and Geant4.10.6 due to logarithmic calls reduction in EM physics
- Optimization with **tuning of G4 parameters** (physics models, physics lists per regions)
- **Neutron** and **Photon Russian Roulette** + **EM range cuts** (ongoing physics validation)
- **Geometry** optimisations (new EMEC variants + R&D on ML guided steppers in geometry)
- **Magnetic field** tailored switch-off
- **Geant4 linking** as static library (a.k.a. Big library)
- Explore **machine learning options** especially for **simulations optimization:**
  - Machine learning solutions to optimise the detector simulation and optimally tune/re-weight parameters (i.e. physics models, physics lists per regions, range cuts, magnetic field)

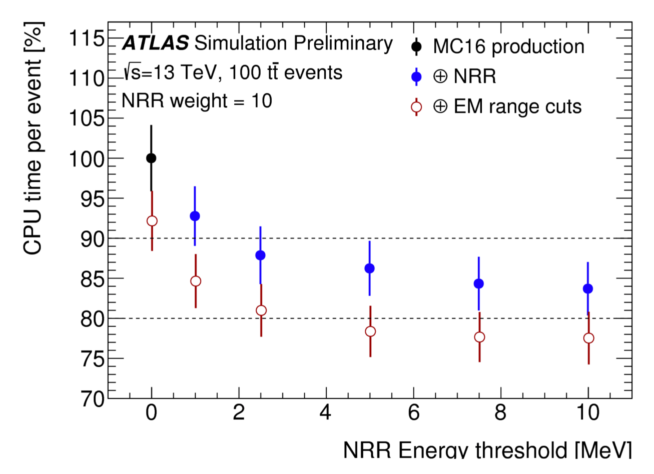

CPU time with NRR + EM range cuts [1]

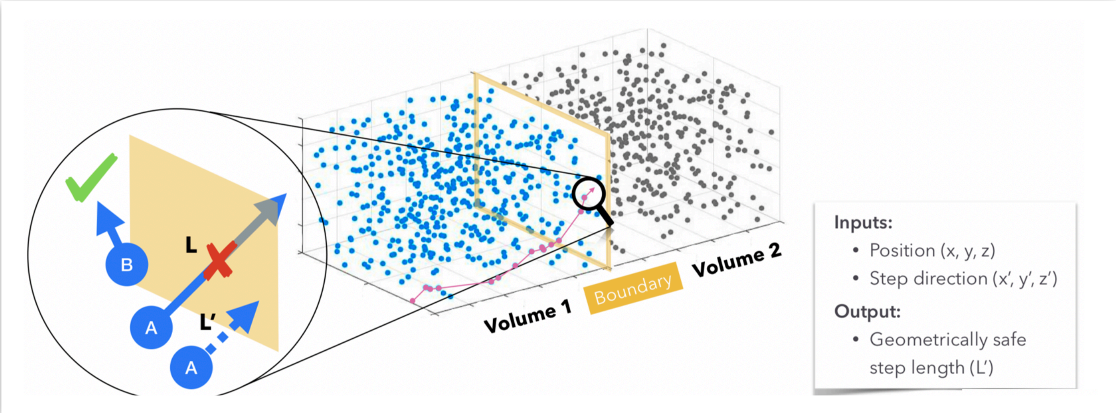

Machine Learning approach to Geant4 steppers in complex geometries [2]

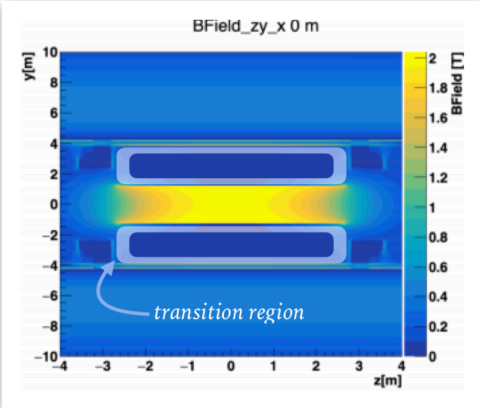

Magnetic Field tailored switch-off [3]

Benchmarks different steppers

- Run with **FullSimLight** + full ATLAS geometry + ATLAS MagField
- Geant4 versions:
  - Geant4.10.4.patch03.atlas01: currently in master
  - Geant4.10.5.patch01.atlas01
  - Geant4.10.6.patch00.atlas01
- NEW and OLD way of setting the **AtlasRK4**, **DormandPrince745** and **DormandPrinceRK56** steppers
- 10k mixed 10 GeV primaries (e+,e-, gamma), isotropic distribution
- Perf stat with 10 runs

Geant4 version	ATLASRK4(OLD)	DPRK56 (OLD)	DPK56/ATLASRK4 OLD way	ATLASRK4 (NEW)	DPRK56 (NEW)	DPK56/ATLASRK4 NEW way
g4.10.4.p03.atl01	5143.27+-19.52 sec.	6216.55+- 68.22 sec.	+20.86%	5303.06+- 24.06 sec.	6340.25 +- 153.74	
g4.10.5.p00.atl01	5299.94 +- 30.54 sec.	6345.10 +- 82.74 sec.	+19.72%	5265.06 +- 21.36 sec.	6544.45 +- 169.34	
g4.10.6.p00.atl01	5079.57 +- 29.25 sec.	5973.12 +- 108.6 sec.	+17.59%	4999.33 +- 23.73 sec.	6210.21+-116	

Geant4 version	ATLASRK4(OLD)	DP745 (OLD)	DP745/ATLASRK4 (OLD)	ATLASRK4 (NEW)	DP745(NEW)	DP754/ATLASRK4 (NEW)
g4.10.4.p03.atl01	5143.27+-19.52 sec.	5692.87 +- 131.20 sec	+10.68%	5303.06+- 24.06 sec.	5882.78+- 131.93	+10.93%
g4.10.5.p00.atl01	5299.94 +- 30.54 sec.	5834.20 +- 150.63 sec	+10.08%	5265.06 +- 21.36 sec.	5714.14 +- 128.57	+8.53%
g4.10.6.p00.atl01	5079.57 +- 29.25 sec.	5686.46 +- 143.76 sec	+11.94%	4999.33 +- 23.73 sec.	5391.98 +- 118.81	+7.85%

Using the InterpolationDriver

New Geant4 version for Run3

- We are testing different Geant4 versions to decide which will be the default for RUN3. Likely:
  - Geant4.10.6
  - Geant4.10.7
- **Geant4.10.6.patch03.atlas01** in the process of become the default version in Athena master
- Plan to have a local build with Geant4.10.7 as well
- Birk's constant tuning
  - Effects still to be fully understood (comparison with default MC16 Geant4 version)

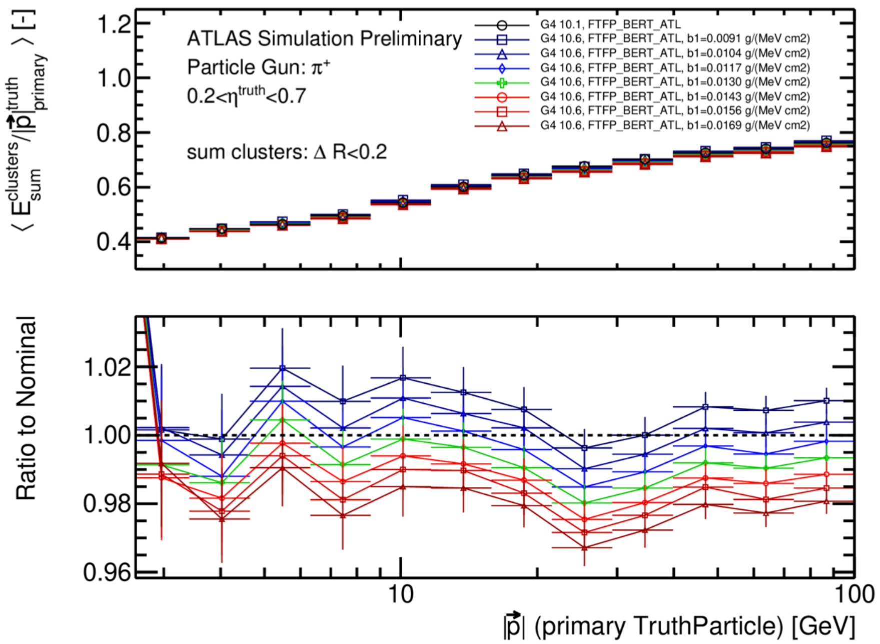

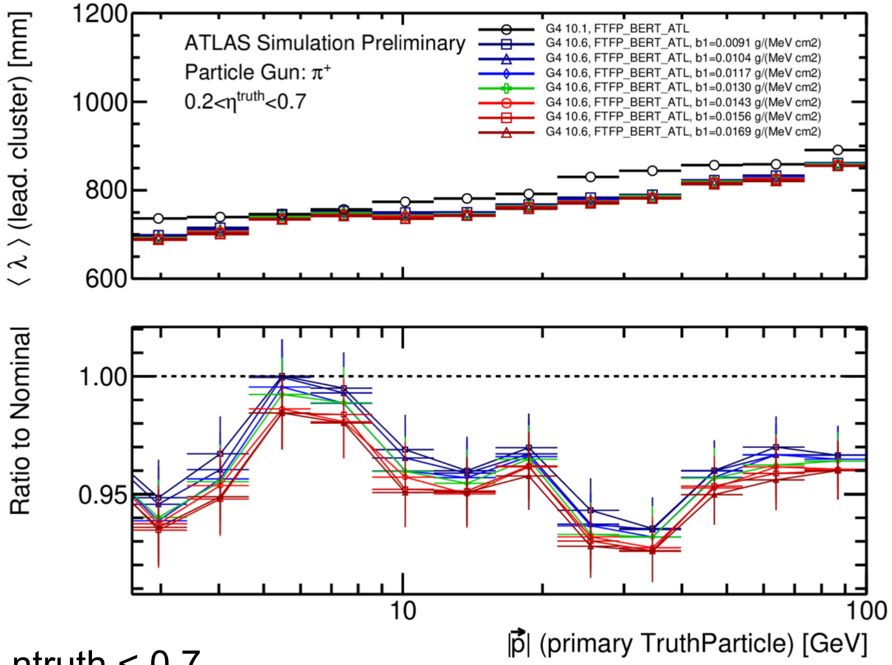

central Barrel: 0.2 < ηtruth < 0.7

Mean E/P

Mean lambda of leading cluster (not the sum of clusters)