LHCb workloads in the HEPscore benchmark

Andrea Valassi (CERN IT)

HEPscore workshop, Monday 19th September 2022
https://indico.cern.ch/event/1170924/contributions/4954486/

Many thanks to C. Bozzi, G. Corti and the whole LHCb Simulation team!
Many thanks to R. Sobie and V. Innocente for the studies cited here!

LHCb Grid workloads *(slide from Feb 2021)*

- MCSimulation is by far the largest LHCb consumer of CPU for Grid jobs
  - MCFastSimulation is a similar workload involving “ReDecay”
  - Around ~90% overall at the end of 2020, expect this throughout Run3

- GEN-SIM jobs: event generation (e.g. Pythia) + detector simulation (Geant4)
  - SIM is ≫ GEN: in other words, *almost all of LHCb Grid CPU is Geant4*
  - Today, single-threaded, single-process application, Gauss (will change)

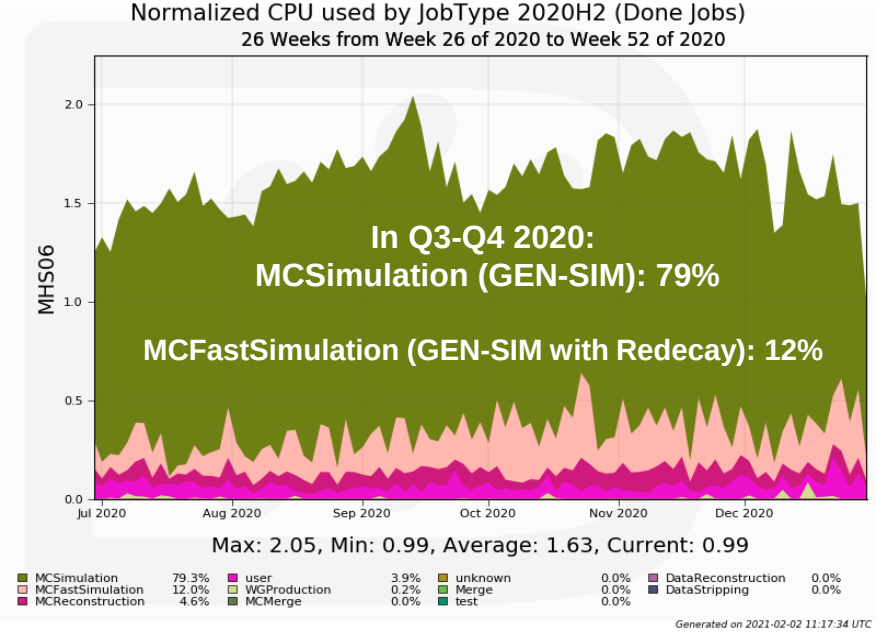

LHCb “v2021” SIM workload for HEPscore

- HEP-SCORE candidate uses workload agreed with LHCb Simulation team in 2021
  - See the benchmark script for details

- Software details: this uses an LHCb "SIM10" software release from 2021
  - See a typical logfile for details
  - Gauss v55r1 on x86_64-centos7-gcc9-opt, single process, single threaded
  - GEN: Pythia8, inclusive B-physics events (options/10000000.py)
  - Pileup configuration nu=3.8, minimum bias is generated for each event (no I/O)
  - No spillover from adjacent beam crossings (“off-time pileup”: prev/next event)
  - SIM: Geant4-10-06-patch02 (29 May 2020)
  - AppConfig v3r404, DecFiles v31r7, CondDB sim-20210617-vc-md100
  - Default setup processes 5 reproducible events, ~10 minutes on a typical CPU

- The application runs GEN+SIM (starts from random seeds), separate scores exist
  - *The recommended SIM score uses SIM-only throughput for n-1 events*
  - *Skip readback of magnetic field and other initialization tasks during the 1st event*

- This is representative of what LHCb runs in production jobs on the Grid
  - "SIM9" is the default for Run2, but "SIM10" productions already exist for Run3

Aside: LHCb SIM differs from other SIMs

- Context (extensively described in this workshop! also see backup slides)
  - (1) Analysis of *correlation of experiment workloads with one another* [R. Sobie]
    - Goal: speed up HEPscore - can we remove some workloads from the "mix"?
    - *Finding: large deviations (20% to 40%) between LHCb SIM and other SIMs*
  - (2) Analysis of *software profiles of experiment workloads* [V. Innocente]
    - *Finding: ~50% of CPU time in LHCb SIM is in G4LogicalBorderSurface::GetSurface!*
      - This is very different from the SIM profiles in other experiments
      - Vincenzo's observation: workload uses G4 10.6, performance was improved in G4 10.7

- LHCb reply on these observations
  - Thanks a lot, extremely useful studies! Will certainly try G4 10.7 soon
  - Large use of G4LogicalBorderSurface is due to RICH detector physics (mirrors)
    - RICH and CALO have always been the two largest consumers of CPU in LHCb SIM
    - RICH was optimized in SIM9
    - CALO (via Geant4) was optimized in SIM10 (for info, SIM10 is 40% faster than SIM9)
  - *It is important to keep LHCb in the HEPscore mix precisely because it differs*

- More generally (personal observation):
  - Different workloads stress the CPUs in different ways and may scale differently
    - (This will be a theme in my talk tomorrow about Madgraph on SIMD CPUs)
  - If a workload scales differently, it is not necessarily "wrong", it is just different

Expected evolution of LHCb SIM

- The "v2021" LHCb workload is now ~as-good-as-it-gets representative of Run3
  - "SIM10": as of today, there is no more recent SIM version in production
  - Several changes are on the way, however...
- LHCb has moved from gcc9 to gcc11 as default compiler since 2021
  - But this change by itself does not justify building a new container today
- *LHCb will move from single-threaded Gauss to multi-threaded Gauss-on-Gaussino*
  - *Timescale: fully replace SIM10 by end 2023, not before*
- LHCb will move from G4 10.6 to a more recent version of G4 such as G4 10.7
  - Timescale: not before the migration to Gauss-on-Gaussino is completed
- More exotic changes on the way: ARM (and Power9)
  - x86 is still the only CPU architecture used in production (and for Grid jobs)
  - An old ARM build exists, and the effort to port the code to ARM is restarting
  - No builds on Power9 yet - the interest and motivation are there, not the effort...
  - Note: GPUs in LHCb are used for HLT but not in Grid/SIM (no Geant4 on GPU)

Other questions from the TF

As per the guidelines: https://new.cernbox.cern.ch/s/AMLHAgWKPMWjNjp

- Experience in integrating LHCb workloads in the HEP standalone containers?
  - Integration was easy for LHCb (AV co-developed the HEP-workloads framework)
- About the HEPscore "mix"
  - It seems reasonable to keep the benchmark mix unchanged for 3-5 years
  - LHCb SIM should be part of the "mix" because it is different from other SIMs
  - Averaging technique: no strong opinion (give more weight to SIMs, if anything)
- Any other comments
  - Thanks for the good work! The studies by Randy and Vincenzo were very useful

Backup slides

For reference – previous talks

- In LHCb
  - LHCb Simulation Meeting, Wed 24th Aug 2022, https://indico.cern.ch/event/1187106
  - LHCb Simulation Meeting, Tue 9th Feb 2021, https://indico.cern.ch/event/1004421
  - LHCb Computing Management, Thu 28th Jan 2021, https://indico.cern.ch/event/984385
  - LHCb Week, Computing Parallel, Thu 04th Mar 2019, https://indico.cern.ch/event/802489
- In the benchmarking WG and TF
  - (HEPscore Task Force, Wed 16th Jun 2021, https://indico.cern.ch/event/1030676)
  - (HEPscore Task Force, Wed 5th May 2021, https://indico.cern.ch/event/1030672)
  - HEPscore Task Force, Wed 3rd Feb 2021, https://indico.cern.ch/event/992804
  - (HEPiX benchmarking WG, Fri 3rd Jul 2020, https://indico.cern.ch/event/872154)
  - Benchmarking pre-GDB, Tue 8th Oct 2019, https://indico.cern.ch/event/739897
  - System performance WG, Wed 6th Mar 2019, https://indico.cern.ch/event/772026
  - HEPiX benchmarking WG, Fri 1st Feb 2019, https://indico.cern.ch/event/782598

Backup slides

(from LHCb Simulation Meeting, Wed 24th Aug 2022)
https://indico.cern.ch/event/1187106

LHCb SIM is different (1) – throughput scaling

- Results presented by Randy Sobie to the WG in May 2022
  - https://indico.cern.ch/event/1104203
- Study correlation of SIM throughputs of LHCb/ATLAS/CMS SIM on different CPUs
  - ATLAS and CMS are in reasonable agreement with one another
  - LHCb behaves quite differently! Why?...

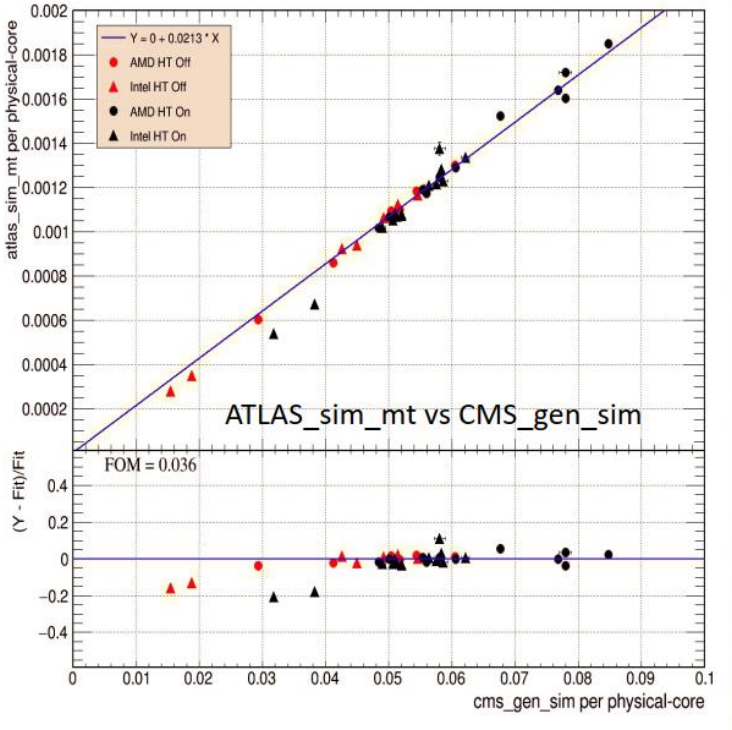

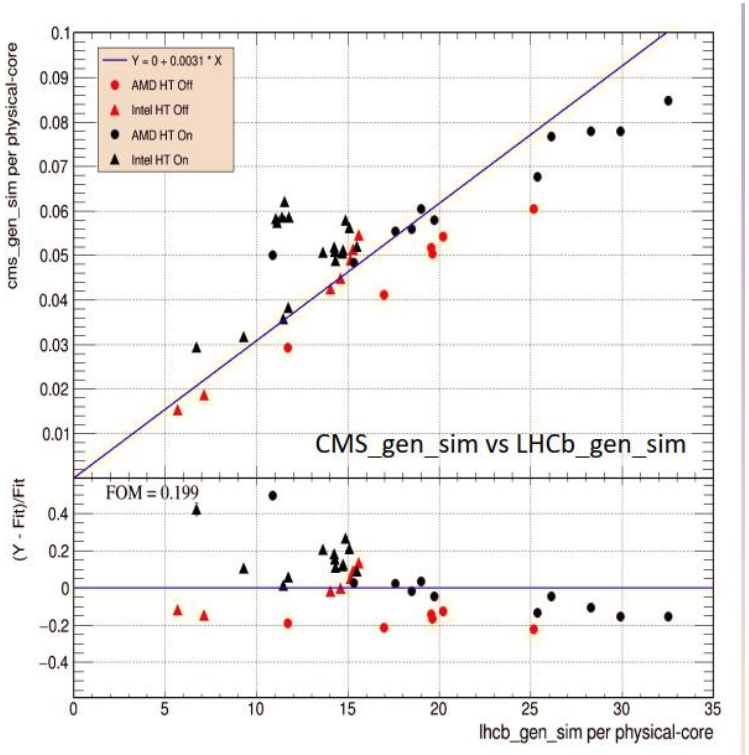

LHCb SIM is different (2) – software profiling

- Results presented by Vincenzo Innocente in email threads in July 2022
  - He will make a detailed presentation at the workshop in September

- CPU time profiling in LHCb/ATLAS/CMS/ALICE SIM on different CPUs
  - LHCb behaves very differently! (See details on the next three slides)
  - **Almost 50% of CPU time is in G4LogicalBorderSurface::GetSurface**
  - Why? Vincenzo's suggestion: *"the code in question (G4LogicalBorderSurface) has changed between G4 10.6 and 10.7: it is now a map and the search should be faster (in may opinion a hashmap would be even better)"*

- Questions:
  - Was it known in LHCb that 40-50% of SIM is in G4LogicalBorderSurface?
  - Which official Geant4 version does Geant4-10-06-patch02 correspond to?
  - Have you tried Geant4 10.7 and did you observe different performance?

SIM software profiling – LHCb

- Results presented by Vincenzo Innocente in email threads in July 2022
  - He will make a detailed presentation at the workshop in September

LHCb 20 events

>>> 40.69%	**python**	**libG4geometry.so**	**[.] G4LogicalBorderSurface::GetSurface**
>>> 5.18%	python	libCLHEP-2.4.4.0.so	[.] CLHEP::RanluxEngine::flat
>>> 2.34%	python	libCLHEP-2.4.4.0.so	[.] CLHEP::RanluxEngine::flatArray
>>> 1.63%	python	libxerces-c-3.2.so	[.] xercesc_3_2::DOMDeepNodeListImpl::nextMatchingElementAfter
>>> 1.51%	python	libG4geometry.so	[.] G4Navigator::LocateGlobalPointAndSetup
>>> 0.97%	python	libG4processes.so	[.] G4VEmProcess::PostStepGetPhysicalInteractionLength
>>> 0.95%	python	libG4tracking.so	[.] G4SteppingManager::DefinePhysicalStepLength
>>> 0.94%	python	libG4geometry.so	[.] G4VoxelNavigation::ComputeStep
>>> 0.80%	python	libG4tracking.so	[.] G4SteppingManager::InvokePSDIP
>>> 0.78%	python	libG4processes.so	[.] G4VProcess::ResetNumberOfInteractionLengthLeft
>>> 0.75%	python	libDetDescLib.so	[.] LHCb::MagneticFieldGrid::fieldVectorLinearInterpolation
>>> 0.75%	python	libG4global.so	[.] G4PhysicsVector::Value
>>> 0.68%	python	libstdc++.so.6.0.27	[.] __cxxabiv1::__vmi_class_type_info::__do_dyncast
>>> 0.67%	python	libG4geometry.so	[.] G4SubtractionSolid::Inside
>>> 0.60%	python	libm-2.17.so	[.] __ieee754_log_avx
>>> 0.54%	python	libG4processes.so	[.] G4ProductionCutsTable::ScanAndSetCouple
>>> 0.52%	python	libGaussTools.so	[.] virtual thunk to GiGaStepActionSequence::UserSteppingAction(G4Step const*)
>>> 0.51%	python	libG4processes.so	[.] G4UniversalFluctuation::SampleFluctuations

SIM software profiling – ATLAS

- Results presented by Vincenzo Innocente in email threads in July 2022
  - He will make a detailed presentation at the workshop in September

ATLAS 10 events

```
>>  9.01%  athena.py     libGeoSpecialShapes.so          [.] LArWheelCalculator_Impl::DistanceCalculatorSaggingOff::DistanceToTheNeutralFibre
>>  3.32%  athena.py     libimf.so                       [.] __libm_sincos_e7
>>  2.92%  athena.py     libG4processes.so               [.] G4VEmProcess::PostStepGetPhysicalInteractionLength
>>  2.12%  athena.py     libG4geometry.so                [.] G4Navigator::LocateGlobalPointAndSetup
>>  1.97%  athena.py     libGeoSpecialShapes.so          [.] LArWheelCalculator::parameterized_sincos
>>  1.85%  athena.py     libG4processes.so               [.] G4UniversalFluctuation::SampleFluctuations
>>  1.61%  athena.py     libimf.so                       [.] __libm_atan2_l9
>>  1.51%  athena.py     libG4processes.so               [.] G4UrbanMscModel::SampleCosineTheta
>>  1.43%  athena.py     libG4geometry.so                [.] G4PolyconeSide::Inside
>>  1.43%  athena.py     libG4geometry.so                [.] G4VoxelNavigation::ComputeStep
>>  1.40%  athena.py     libG4tracking.so                [.] G4SteppingManager::DefinePhysicalStepLength
>>  1.22%  athena.py     libMagFieldElements.so          [.] BFieldCache::getB
>>  1.16%  athena.py     libGeo2G4Lib.so                 [.] LArWheelSolid::search_for_nearest_point
>>  1.14%  athena.py     libG4geometry.so                [.] G4Navigator::LocateGlobalPointWithinVolume
>>  1.11%  athena.py     libG4geometry.so                [.] G4AtlasRK4::Stepper
>>  1.03%  athena.py     libG4tracking.so                [.] G4SteppingManager::Stepping
>>  1.01%  athena.py     libG4geometry.so                [.] G4VCSGfaceted::Inside
>>  1.01%  athena.py     libG4geometry.so                [.] G4Navigator::ComputeStep
>>  1.01%  athena.py     libG4geometry.so                [.] G4NavigationLevel::operator=
>>  1.01%  athena.py     libG4geometry.so                [.] G4TouchableHistory::GetVolume
>>  1.00%  athena.py     libG4tracking.so                [.] G4SteppingManager::InvokePSDIP
>>  0.94%  athena.py     libG4processes.so               [.] G4Transportation::PostStepDoIt
```

SIM software profiling – ALICE

- Results presented by Vincenzo Innocente in email threads in July 2022
  - He will make a detailed presentation at the workshop in September

ALICE 10 events

```
>  2.91%  o2-sim-device-r  libg4root.so                    [.] TG4RootDetectorConstruction::GetG4VPhysicalVolume
>   2.74%  o2-sim-digitize  libO2TPCSimulation.so                [.] o2::tpc::Digitizer::process
>   2.19%  o2-sim-digitize  libO2TRDSimulation.so                [.] o2::trd::Digitizer::convertHits
>   1.86%  o2-sim-device-r  libO2Field.so                      [.] o2::math_utils::Chebyshev3D::Eval
>   1.72%  o2-sim-device-r  libG4processes.so                    [.] G4RToEConvForGamma::ComputeValue
>   1.57%  o2-sim-digitize  libO2TRDSimulation.so                [.] o2::trd::SimParam::timeResponse
>   1.51%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoSubtraction::Contains
>   1.45%  o2-sim-device-r  libG4processes.so                    [.] G4GEMProbability::CalcProbability
>   1.13%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoUnion::Contains
>   0.99%  o2-sim-digitize  libO2TPCSimulation.so                [.] o2::tpc::SAMPAProcessing::getShapedSignal
>   0.89%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoNavigator::Safety
>   0.88%  o2-sim-device-r  libHist.so.6.24.06                 [.] TGraph::Eval
>   0.78%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoTranslation::MasterToLocal
>   0.77%  o2-sim-digitize  libO2TRDSimulation.so                [.] o2::trd::SimParam::crossTalk
>   0.72%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoIdentity::MasterToLocal
>   0.72%  o2-sim-device-r  libg4root.so                       [.] TG4RootNavigator::SynchronizeHistory
>   0.67%  o2-sim-device-r  libBase.so.18.4.7                    [.] FairMCApplication::Stepping
>   0.66%  o2-sim-device-r  libGeom.so.6.24.06                   [.] TGeoVoxelFinder::GetNextCandidates
>   0.62%  o2-sim-device-r  libG4processes.so                    [.] G4RToEConvForElectron::ComputeValue
>   0.61%  o2-sim-device-r  libG4tracking.so                   [.] G4SteppingManager::DefinePhysicalStepLength
>   0.61%  o2-sim-device-r  libG4geometry.so                     [.] G4Region::BelongsTo
>   0.61%  o2-sim-device-r  libO2SimulationDataFormat.so                [.] o2::data::Stack::ReorderKine
>   0.58%  o2-sim-device-r  libg4root.so                       [.] TG4RootDetectorConstruction::GetNode
>
```