



CMS SIMULATION STATUS

S. Banerjee, D. Elvira, FNAL

V. Ivanchenko CERN & Tomsk State University, Russia

D. Piparo, CERN

For CMS Simulation group

24th Geant4 Collaboration Workshop

Jefferson Lab, Virginia (US), 23-27 September 2019



Outline

- Status of CMS full simulation for Run-2
- CMS plan to use Geant4 10.6 for 2021 production
- List of problems of Geant4 10.5ref08 for CMS
- Results on CPU/memory performance studies

Results reported today are very fresh, CMS allows us to show these results having in mind CMS plan to use 10.6 for production in 2021

Status of CMS full simulation for Run-2



- CMS used Geant4 10.0p02 for 2015-2016 Run-2 simulation
 - *Sequential Geant4 in production*
 - *QGSP_FTFP_BERT_EML Physics List*
- Geant4 10.2p02 is used for 2017 Run-2 simulation production
 - *Minor fixes from Geant4 10.2p03 are added*
 - *FTFP configuration from Geant4 10.1 is used*
 - *FTFP_BERT_EMM PhysicsList*
 - *Multithreaded Geant4*
- Geant4 10.4 is adopted as a production version for 2018 MC production for Run-2 and will be kept as a legacy MC for Run-2
 - *Relevant known fixes are added on top of 10.4 as CMS private patches*
 - *The MT mode, production platforms (gcc7.3):*
 - *slc7-amd64-gcc700*
- **Total number of simulated events for Run-2 ~60B**



CMS plan to use Geant4 10.6 for 2021 production

- Geant4 10.6 is planned for 2021 MC production for Run-3
 - *The MT mode, production platform of today slc7-amd64-gcc820*
- In order to have early test of Geant4 10.6 we integrated reference versions of Geant4 in a special CMSSW branch or test them privately
 - *10.6beta, 10.5ref07, 10.5ref08*
 - *Physics results of comparisons with test-beam data and collision events are reported at parallel session by S.Banerjee*
- In this talk we discuss
 - *CPU performance for Geant4 10.6beta, 10.5ref08 versus 10.4p03*
 - *Problems observed with Geant4 10.6beta, 10.5ref08*
- We use 3 different CMS geometries:
 - *2018 – Run-2*
 - *2023D17 – Phase-2 geometry with HGCal sensors as Polyhedrons*
 - *2023D28 – Phase-2 geometry with HGCal sensors as ExtrudedSolids and other modifications in Tracker and HGCal*

CPU performance study



- CPU performance study was carried out with cmsdev31 PC
 - 16 AMD cores
 - *slc7-amd64-gcc820*
 - 8-threads runs
 - 2 jobs submitted in parallel in order to emulate grid production and reduce instability
 - Average throughput is shown in plots
 - We do not exclude initialization both from profiling and from CPU analysis
- Two versions of CMSSW considered
 - Legacy Run-2 (*Geant4 10.4p03*)
 - Special branch (*Geant4 10.5ref08*)
- Physics List: **FTFP_BERT_EMM**
 - Based on *FTFP_BERT*
 - Special CMS configuration of EM physics optimizing simulation for Run-2
- We observe 3 types of frequent warnings

Warnings of type #1: Possible overlaps

```
----- WWW ----- G4Exception-START ----- WWW -----  
*** G4Exception : Transport-001-ExcessSteps  
    issued by : G4Transportation::AlongStepDoIt  
Transportation is killing track that is looping or stuck.  
Track is e+ and has 117.895 MeV energy ( pre-Step = 117.895 )  
momentum = (14.9192,-21.6865,-115.442) mag= 118.405  
position = (21.2547,57.2604,-4037.58) is in volume 'BeamVacuum11',  
its material is 'Vacuum' with density = 9.99998e-17 g/cm^3  
Total number of Steps by this track: 168  
Length of this step = 2221.46 mm  
Number of propagation trials = 1 ( vs maximum = 10 for 'important' particles )  
    ( Number of *calls* of Transport/AlongStepDoIt = 6215895 )  
*** This is just a warning message. ***  
----- WWW ----- G4Exception-END ----- WWW -----
```

Warnings of type #2: Tracks getting killed

```
----- WWW ----- G4Exception-START ----- WWW -----  
*** G4Exception : GeomNav1002  
    issued by : G4Navigator::ComputeStep()  
Stuck Track: potential geometry or navigation problem.  
Track stuck, not moving for 10 steps.  
Current phys volume: 'PixelForwardOuterDiskInnerRing'  
- at position : (7.603902648951626, -87.89809353179001, 291.2999999992056)  
  in direction: (0.01154616486386289, -0.2770520649808579, 0.9607855324508061)  
  (local position: (-7.603902648951649, 87.89809353179, -10.50000000079433))  
  (local direction: (-0.01154616486386255, 0.2770520649808583, 0.9607855324508062)).  
Previous phys volume: 'PixelForwardDiskZplus'  
  
Likely geometry overlap - else navigation problem !  
*** Trying to get *unstuck* using a push - expanding step to 1e-07 (mm) ...  
    Potential overlap in geometry !  
*** This is just a warning message. ***  
----- WWW ----- G4Exception-END ----- WWW -----
```

- Also rare problem at a surface of a big Polycon inside another Polycon
- JIRA ticket for VecGeom is set: <https://sft.its.cern.ch/jira/browse/VECGEOM-537>
- Still in progress – not easy to debug and to reproduce

Warnings of type #3: Too excited fragment

```
----- WWW ----- G4Exception-START ----- WWW -----  
*** G4Exception : had0034  
      issued by : G4ExcitationHandler::BreakItUp()  
High excitation Fragment Z= 1 A= 5 Eex/A(MeV)= 104.381  
*** This is just a warning message. ***  
----- WWW ----- G4Exception-END ----- WWW -----
```

- Too many in 10.5ref08 – should users know about it ?



Run Summary (FTFP_BERT_EMM in 10.4.p03)



	Geant4+	Native	10.4.p03	Geant4+	VecGeom	10.4.p03
	CPU	RSS	Warning	CPU	RSS	Warning
MB (2018)	7.682 s	0.65 GB	0 0 0	0.887	0.68 GB	790 0 0
ttbar (2018)	43.434 s	0.63 GB	0 0 0	0.901	0.70 GB	1402 0 0
MB (2023D17)	9.848 s	0.71 GB	0 0 0	0.895	0.70 GB	0 0 0
ttbar (2023D17)	73.637 s	0.75 GB	0 0 0	0.895	0.74 GB	0 0 0
MB (2023D28)	10.229 s	0.68 GB	2 0 0	0.968	0.66 GB	0 0 0
ttbar (2023D28)	75.997 s	0.69 GB	0 0 0	0.903	0.68 GB	0 0 0

- 3000 events in each sample
- The overlap warnings happen more often in versions using VecGeom (v0.5.0) and these are very few for Native geometry
- Switched on the overlap tools - no overlap observed in tracker region nor in parts of calorimeter (except the ones in EB)



Run Summary in 10.5.ref08



	Geant4+	Native	10.5.ref08	Geant4+	VecGeom	10.5.ref08
	CPU	RSS	Warning	CPU	RSS	Warning
MB (2018)	7.885 s	0.65 GB	1 1305 317	0.924	0.66 GB	183 1304 306
ttbar (2018)	42.112 s	0.68 GB	2 1765 610	0.937	0.71 GB	381 1797 642
MB (2023D17)	10.054 s	0.71 GB	0 225 342	0.855	0.69 GB	0 229 335
ttbar (2023D17)	72.635 s	0.75 GB	1 318 604	0.893	0.74 GB	3 343 587
MB (2023D28)	10.108 s	0.70 GB	0 224 320	0.895	0.64 GB	0 222 312
ttbar (2023D28)	78.309 s	0.68 GB	2 362 614	0.886	0.68 GB	0 412 621

- All killed tracks in these versions are e- or e+ and they happen in Air or Vacuum. These warnings happen more often in the current geometry.
- The overlap warnings happen more often in versions using VecGeom (v1.1.4) and these are very few for Native geometry



CPU for alternative Physics Lists

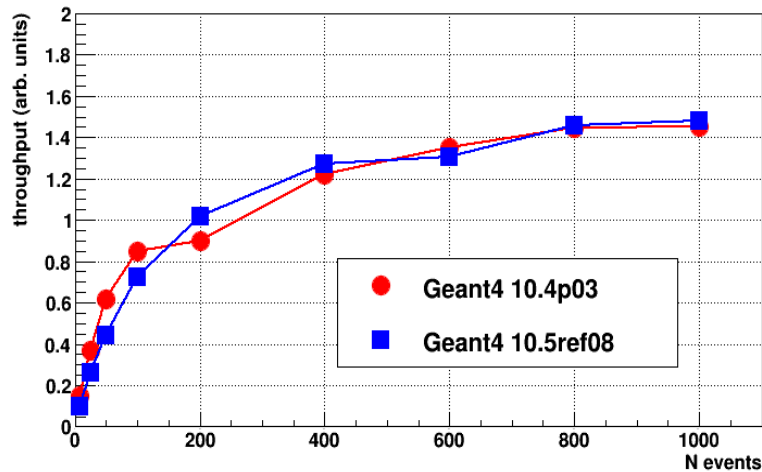


- FTFP_BERT_EMM is the default physics list for CMS
- A new physics list needed for the new HGCal is also tried:
 - FTFP_BERT_EMN – based on Opt4 variant but features not needed for LHC are removed
 - Klein-Nishina model for Compton scattering for all energies
 - Uses Goudsmit-Saunderson model for multiple scattering below 100 MeV
 - No atomic de-excitation


		2018	2023D17	2023D28
Minimum Bias	FTFP_BERT_EMM	1.00	1.29	1.45
(10.4.p03)	FTFP_BERT_EMN	1.09	2.12	1.76
(10.5.ref08)	FTFP_BERT_EMM	1.07	1.26	1.33
	FTFP_BERT_EMN	1.13	2.15	2.30
t-tbar	FTFP_BERT_EMM	1.00	1.69	1.76
(10.4.p03)	FTFP_BERT_EMN	1.09	2.92	3.21
(10.5.ref08)	FTFP_BERT_EMM	1.01	1.66	1.77
	FTFP_BERT_EMN	1.15	3.00	3.28

CPU and RSS for QCD and ttbar


CPU for QCD at 13 TeV



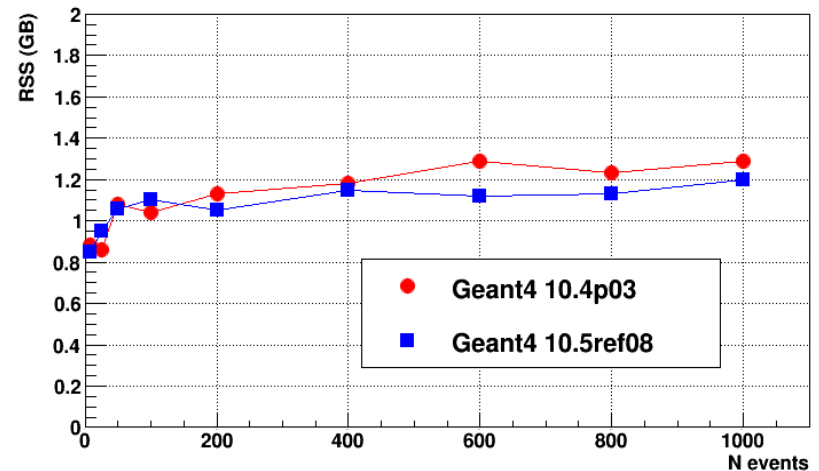
CPU for ttbar at 13 TeV



Memory for QCD at 13 TeV



Memory for ttbar at 13 TeV



Profiling for 10.4p03 (initialization is not excluded)

Rank	Total %	Self	Symbol name
<u>29</u>	5.00	70.73	<u>G4DormandPrince745::Stepper(double const*, double const*</u>
<u>40</u>	4.78	67.72	<u>G4PhysicsVector::Value(double, unsigned long&) const</u>
<u>48</u>	3.96	56.08	<u>G4Mag UsualEqRhs::EvaluateRhsGivenB(double const*, double const*</u>
<u>52</u>	3.17	44.90	<u>dl update slotinfo</u>
<u>33</u>	3.03	42.88	<u>G4Navigator::LocateGlobalPointAndSetup(CLHEP::Hep3Vec</u>
<u>30</u>	1.82	25.84	<u>G4VoxelNavigation::ComputeStep(CLHEP::Hep3Vector const&, CL</u>
<u>71</u>	1.70	24.02	<u>G4Navigator::LocateGlobalPointWithinVolume(CLHEP::Hep</u>
<u>45</u>	1.60	22.70	<u>sim::Field::GetFieldValue(double const*, double*) const</u>
<u>34</u>	1.38	19.59	<u>__tls_get_addr</u>
<u>43</u>	1.36	19.28	<u>update get addr</u>
<u>27</u>	1.33	18.85	<u>G4Navigator::ComputeStep(CLHEP::Hep3Vector const&, CL</u>
<u>57</u>	1.29	18.25	<u>G4UniversalFluctuation::SampleFluctuations(G4Material</u>

Profiling for 10.5ref08 (Initialization is not excluded)

Rank	Total %	Self	Symbol name
<u>48</u>	3.42	45.29	<u>_dl_update_slotinfo</u>
<u>33</u>	3.27	43.31	<u>G4Navigator::LocateGlobalPointAndSetup(CLHEP::Hep</u>
<u>47</u>	3.05	40.34	<u>G4VEmProcess::PostStepGetPhysicalInteractionLengt</u>
<u>32</u>	2.76	36.59	<u>G4DormandPrince745::Stepper(double const*, double</u>
<u>64</u>	2.62	34.66	<u>G4Mag_UsualEqRhs::EvaluateRhsGivenB(double const*</u>
<u>27</u>	1.83	24.20	<u>G4VoxelNavigation::ComputeStep(CLHEP::Hep3Vector</u>
<u>91</u>	1.56	20.73	<u>G4ElasticHadrNucleusHE::HadrNucDifferCrSec(int, d</u>
<u>40</u>	1.52	20.16	<u>update get addr</u>
<u>61</u>	1.45	19.27	<u>G4UniversalFluctuation::SampleFluctuations(G4Mate</u>
<u>35</u>	1.45	19.15	<u>_tls get addr</u>
<u>109</u>	1.36	18.07	<u>_sin_avx</u>
<u>42</u>	1.26	16.72	<u>sim::Field::GetFieldValue(double const*, double*)</u>
<u>20</u>	1.21	16.05	<u>G4SteppingManager::DefinePhysicalStepLength()</u>

Comments on PRELIMINARY results of CPU and memory in 10.5ref08

- These preliminary results confirm our previous observation, that CMS simulation run fully explore concrete CPU node if statistics exceed ~500 events
 - *We see a little speed-up for QCD events*
 - *There is definite speed-up for ~5% for ttbar*
 - *Results strongly depend on hardware state and are not final*
- Preliminary observation for 10.5ref08 using igprof
 - *Profiles are substantially different from 10.4p03*
 - *Initialization time increased*
 - **G4ElastHardNucleusHE initialization dominates – should be reduced**
 - *G4DormandPrince457 methods take less time*
 - *No G4PhysicsVector::Value on the top*

Summary

- Our preliminary analysis demonstrates, that can expect Geant4 10.6 to have several advantages and it can be considered for CMS production in 2021
 - *This release will be definitely faster*
- There are also problems and concerns
 - *Too many warnings*
 - *Rare crashes*
 - *Increased initialization time*
- We propose for 10.6:
 - *To try to reduce initialization time (if possible)*
 - *Provide an option to upload all data before 1st event*
 - *Remove frequent warnings from the release completely*