CMS Simulation

Computing Performance

V. Daniel Elvira (FNAL), V. Ivantchenko (CERN)

Outline

- CPU and memory utilization of G4 module of CMS simulation application
- CPU gains when using approximations, individually and cumulative
- CPU and memory dependence on CMS detector geometry – Run 1, 2 versus HL-LHC configurations
- Integration and stress testing of Geant4 with VecGeom within CMSSW

Hardware and Configurations

(For all tests unless indicated otherwise)

Part of an ATLAS/CMS joint exercise in the context of the Community White paper

- Machine (CERN's OpenLab): olhswep16.cern.ch, Intel(R) Xeon(R) CPU E5-2683 v3 @ 2.00GHz, one thread runs
- Compiler: gcc 6.3
- Geant4: version 10.2, FTFP_BERT physics list – Benchmark (Default means what CMS actually runs FTFP_BERT_EMM)
- Pythia events: 13 TeV Pythia minimum bias (300 events) and ttbar (300 events), pseudo rapidity cut |h|<5.5
- Particle gun: 50 GeV e's, muons, pions with a flat η,φ distribution in $\eta=[-0.8,0.8]$, $\eta=[2,2.7]$, and $\varphi=[0,2\pi]$
- Geometry: 2015 or 2016 and HL-LHC scenarios

Will only show a subset of all these configurations

CPU and Memory Utilization

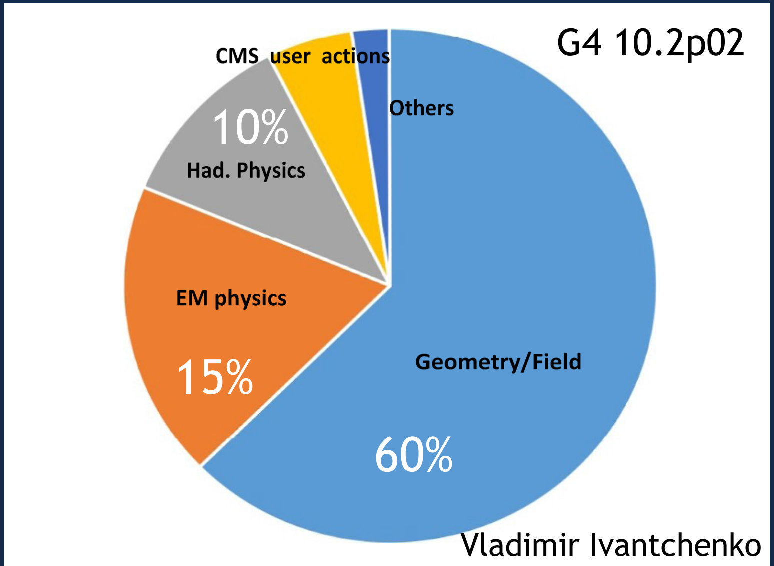

Pie chart of CMS simulation G4 module

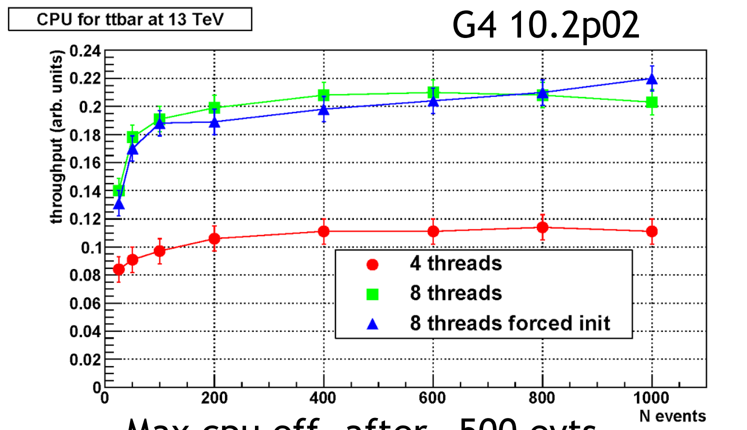

Max cpu eff. after ~500 evts.

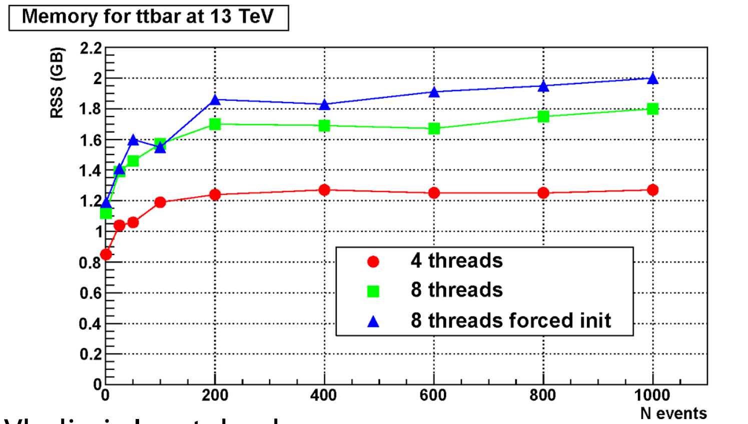

Vladimir Ivantchenko

Throughput, memory per event (ttbar) vs. event number. (Forced init means all data for photon evaporation is read at initialization time.)

Effect of Approximations in CPU

Effect of individual approximations on CPU time performance with respect to the benchmark configuration (FTFP_BERT) – see slide 3

Configuration	MinBias	ttbar
CMS default (FTFP_BERT_EMM)	0.87	0.83
FTFP_BERT Benchmark	1.00	1.00
no Russian roulette	1.33	1.41
no HF shower library	1.68	1.36
no tracking cuts	1.45	1.13
no time cuts	1.05	1.03
no cuts per region	1.07	1.03
no static build	1.05	1.07

HF shower library, Russian roulette, tracking cuts largest contributors
68%-36% 33%-41% 45%-13%

Vladimir Ivantchenko

The cut in range value for the production of secondary particles is 1 mm in the electromagnetic and hadronic calorimeters, 0.01 mm in the pixel detector, 0.1 mm in the strip tracker, 0.002 mm in the muon sensitive volumes, and 1 cm in the support structures. The propagation time cut is 500 ns. A tracking cut is applied for charged particles below 2 MeV inside the vacuum chamber to avoid looping electrons or positrons and parameters associated with particle tracking in EM fields are optimized for performance

Effect of Approximations in CPU

Cumulative effect of approximations on CPU time performance with respect to CMS default – see slide 3

Configuration	MinBias	ttbar
CMS default — What CMS actually runs	1.0	1.0
no Russian roulette	1.28	1.37
no Russian roulette + no tracking and time cuts	1.84	1.63
no Russian roulette + no tracking and time cuts + FTFP_BERT	2.25	2.04
no Russian roulette + no tracking and time cuts + FTFP_BERT + no HF shower library	4.43	3.22
no Russian roulette + no tracking and time cuts + FTFP_BERT + no HF shower library + no cuts per region	4.73	3.41

Vladimir Ivantchenko

The cumulative effect of approximations to the CMS G4 module reduces CPU time utilization by a factor of 4.73 (3.41) for minbias (ttbar) events

CPU Time Dependence on Geometry

	Minimum CPU Time	Bias Memory (RSS)	Top CPU Time	Pair Memory (RSS)
2015	1.00	0.75 GB	1.00	0.79 GB
2016	1.01	0.75 GB	1.01	0.75 GB
2017	1.21	0.70 GB	1.13	0.75 GB
2023D4	1.28	0.87 GB	1.65	0.83 GB
2023D17	1.36	0.78 GB	1.57	0.78 GB

Sunanda Banerjee

- 2015/2016/2017 are Run 2 geometry configurations
- 2023D4/2023D17 are HL-LHC configurations including upgraded pixel, tracker, and muon detectors as well as the novel high granularity HE calorimeter

CPU time/event of Geant4 module increase of 30-65% for HL-LHC wrt Run 2

Number of Volumes

Why is there a 30-65% cpu time difference between 2015 and HL-LHC CMS simulations ?

	Standard	Reflected
Box	1.2M	436k
Tube	92k	2.1k
Trapezoid	238k	150k
Cone	1.9k	0
Polycone	802	32
Polyhedra	4.2k	0
Torus	160	32
PseudoTrapezoid	96	0
TruncatedTube	92	0
UnionSolid	2.8k	976
SubtractionSolid	180k	918
IntersectionSolid	1.1k	0

CMS 2015 Geometry:
3858 solids

4183 logical volumes

2.3 million touchables

	Standard	Reflected
Box	743k	429k
Tube	49k	1k
Trapezoid	160k	141k
Cone	1.9k	0
Polycone	176	32
Polyhedra	16.9M	0
PseudoTrapezoid	96	0
TruncatedTube	92	0
UnionSolid	131k	0
SubtractionSolid	3.4M	594
IntersectionSolid	70	0

Sunanda Banerjee

CMS HL-LHC D4 Geometry:
5177 solids

5372 logical volumes

21.9 million touchables

The CMS HGCal

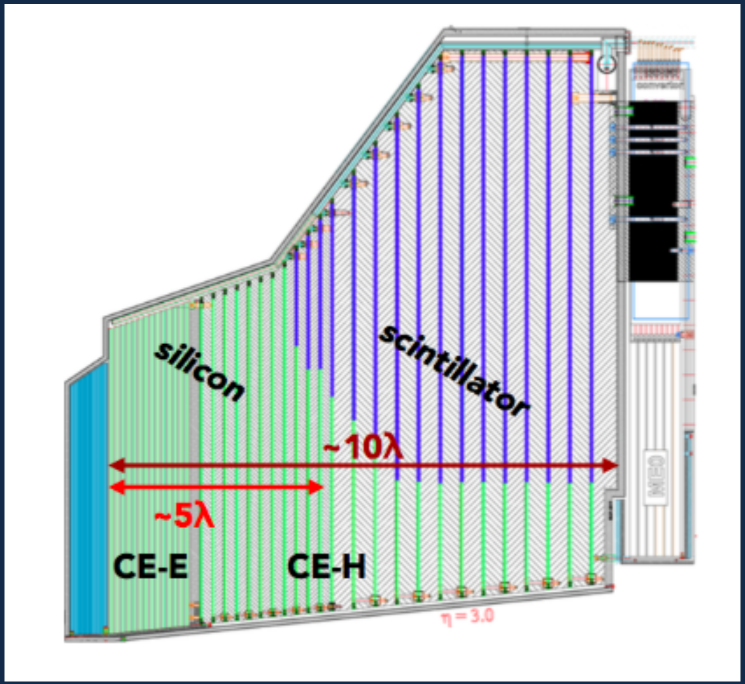

- Covers the [1.5,3] pseudorapidity range
- Consists of 40 layers of silicon and copper/tungsten, brass, or steel
- Readout of 6 million channels with cell sizes of 0.5 cm^2 and 1 cm^2

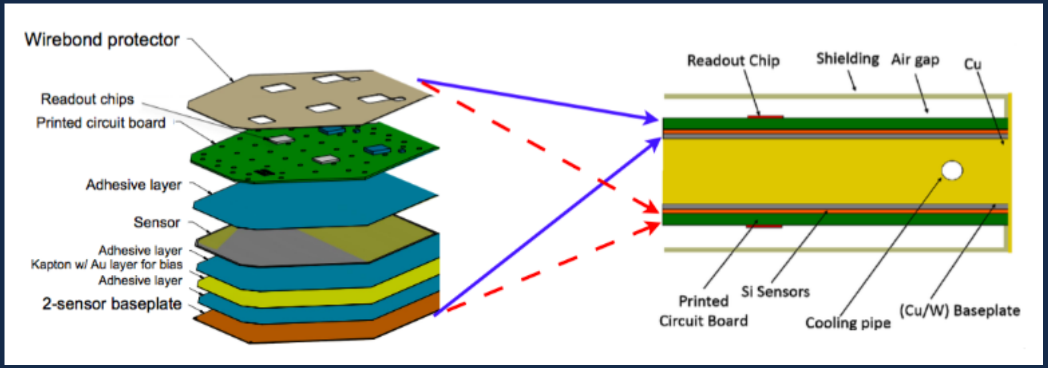

CMS Integration Tests of VecGeom

Plan – migrate to G4 10.4 with VecGeom and "smart tracks" for production of MC samples in the 2018 Collider Run

- Control tracking in field parameters per track
- Tight values for E>Thresh, loos for E<Thresh
- Example: Thresh = 2 MeV may save 10% CPU time

Integrated G4 10.3.ref06 (beta release) with VecGeom to CMSSW (Tests: 3000 $\mu/\pi/e$ at 50 GeV, 500 minbias and 300 ttbar at 14 TeV

	CPU (Native) G4	RSS (Native) G4	CPU (VecGeom)	RSS (VecGeom)
Muon (Barrel)	1	0.54 GB	1.01	0.58 GB
Muon (Endcap)	1	0.58 GB	0.99	0.56 GB
Pion (Barrel)	1	0.60 GB	0.91	0.67 GB
Pion (Endcap)	1	0.60 GB	0.95	0.64 GB
Elec (Barrel)	1	0.59 GB	0.92	0.66 GB
Elec (Endcap)	1	0.60 GB	0.91	0.66 GB
Minimum Bias	1	0.67 GB	0.89	0.76 GB
t-tbar	1	0.74 GB	0.96	0.75 GB

Sunanda Banerjee

CPU gain of 4% (11%) for ttbar (minbias)

when using VecGeom with G4 (scalar mode) with respect to native G4 geometry

CMS Integration Tests of VecGeom

Integrated G4 10.3.ref06 (beta release) with VecGeom to CMSSW (Tests: 3000 $\mu/\pi/e$ at 50 GeV, 500 minbias and 300 ttbar at 14 TeV

VecGeom	CPU (Scalar)	RSS (Scalar)	CPU (Vectorized)	RSS (Vectorized)
Muon (Barrel)	1	0.58 GB	1.02	0.53 GB
Muon (Endcap)	1	0.56 GB	1.04	0.52 GB
Pion (Barrel)	1	0.67 GB	1.05	0.52 GB
Pion (Endcap)	1	0.64 GB	1.01	0.52 GB
Elec (Barrel)	1	0.66 GB	1.11	0.50 GB
Elec (Endcap)	1	0.66 GB	1.12	0.50 GB
Minimum Bias	1	0.76 GB	1.08	0.59 GB
t-tbar	1	0.75 GB	1.05	0.59 GB

Sunanda Banerjee

CPU loss of 5% (8%) for ttbar (minbias)

when using VecGeom with G4 (vector mode) with respect to using G4 (scalar mode)

- Stuck tracks – most in pixel region (polycone, polyhedra) - remains
- Crashes occur in different detectors for tracks // to polycone surface – Fixed

Sunanda/Vladimir (CMS) working with Raman Seghal, Soon Jun, Gabriele Cosmo, etc until robustness issues are resolved

Summary

- CMS G4 module spends 60% of the CPU time in geometry and tracking in field, 15% in EM physics, 10% in HAD physics
- HF shower library saves 68% (36%) time for minbias (ttbar), Russian roulet 33% (41%), tracking cuts 45% (13%)
- Cumulative savings of all approximations are factors of 4.73 and 3.41 for minbias and ttbar respectively
- CPU/event expected to increase by 30-60% for HL-LHC geometry configurations (extended tracker and HGCal)
- Use of VecGeom with Geant4 (scalar mode) yields a 4-11% savings in CPU time
- Use of VecGeom with Geant4 (vector mode) yields a 5-8% time performance penalty – improvement of large factor would come with the GeantV engine