Status of GeantV Integration in CMSSW

Kevin Pedro, Sunanda Banerjee (FNAL) October 4, 2019





Introduction

- Integration testing of GeantV w/ CMSSW has several goals:
 - Demonstrate benefits of co-development between R&D team & experiments
 - Exercise capabilities of CMSSW framework to interface with external processing (ExternalWork mechanism) and handle track-level parallelization in detector simulation
 - o Measure any potential CPU penalties when running GeantV in CMSSW
 - Estimate cost of adapting to new interfaces and eventually migrating to new (and potentially backward-incompatible) tools such as GeantV
 - Thinking forward to HPC/GPU solutions
- *Not* planning to migrate CMS simulation to GeantV
 o Geant4 collaboration has not endorsed the project
 - o This is an R&D exercise

GeantV Integration Tests in CMSSW

- Repositories: <u>install-geant</u>, <u>SimGVCore</u>
- ✓ Generate events in CMSSW framework, convert HepMC to GeantV format
- ✓ Build CMSSW geometry natively and pass to GeantV engine (using TGeo)
- Using constant magnetic field, limited EM-only physics list
- ✓ Calorimeter scoring adapted
- ✓ Run GeantV using CMSSW ExternalWork feature:

o Asynchronous, non-blocking, task-based processing

✓ Output in CMS format, immediately suitable for digitization etc.



Geant4 vs. GeantV Scoring

- Sensitive detectors (SD) and scoring trickiest to adapt
 - \circ Necessary to test "full chain" (simulation \rightarrow digitization \rightarrow reconstruction)
 - o Significantly more complicated than Geant4 MT



Geant4 shares memory, but each event processed in separate thread Each event processed in multiple threads, mixed in with other events

Duplicate SD objects per event per thread, then aggregate
 → 4 streams, 4 threads = 16 SD objects

o GeantV TaskData supports this approach

- ➢ Use template wrappers to unify interfaces and operations
 - o Ensure exact same SD code used for Geant4 & GeantV
 - o Minimize overhead (no branching or virtual table)



- Each ScoringClass object has instance of CaloSteppingAction
 o Some additional memory overhead from duplicated class members
 - Attempt to minimize this by storing volume maps in magic static struct
- Merged ScoringClass output copied to cache attached to Event object
 - GeantV may consider event finished before CMSSW has written output
 → copy to cache, then immediately clear ScoringClass objects (avoid possible race conditions)

Physics Validation

- Settings:
 - o Geant4 10.4p2 w/ VecGeom v0.5 (scalar)
 - o GeantV pre-beta-7 w/ VecGeom v1.1
 - o All CMS-specific G4 optimizations disabled
 - o Same production cuts (default 1mm)
 - Single thread (reproducible pRNG sequences)
 - \circ Generate 1000 events w/ single electron, E=100 GeV, $\eta=1.0,\,\phi=1.1$
- Tests: (same generated events used for G4 and GV)
- 1. No field (B = 0)
- 2. Constant field (B = 3.8 T)

(more in backup)

1. Energy Deposits for 100 GeV e- (B=0)

100 GeV Electron B=0 EB (Geant4 vs GeantV)

100 GeV Electron B=0 EE (Geant4 vs GeantV)



- The number of entries differs by 0.3% (7.4%) in EB (EE)
- The means differ by 0.2% for EB and 2.5% for EE

1. Hit Time for 100 GeV e- (B=0)

100 GeV Electron B=0 EB (Geant4 vs GeantV)

100 GeV Electron B=0 EE (Geant4 vs GeantV)



- Means differ by 0.07% for EB and 0.13% for EE
- GeantV and Geant4 applications provide roughly the same distributions

2. Energy Deposits for 100 GeV e- (B=3.8)

100 GeV Electron B=3.8 EB (Geant4 vs GeantV)

100 GeV Electron B=3.8 HB (Geant4 vs GeantV)



- The number of entries differ by 0.4% (23.3%) in EB (HB)
- The means differ by 2.2% for EB and 8.8% for HB

2. Hit Time for 100 GeV e- (B=3.8)

100 GeV Electron B=3.8 EB (Geant4 vs GeantV)

100 GeV Electron B=3.8 HB (Geant4 vs GeantV)



- The means differ by 0.03% for EB and 1.15% for HB
- There is a small difference in the physics results of GeantV and Geant4 applications in the presence of B-field

Performance Tests

- Settings:
 - o GeantV pre-beta-7+ (63468c9b)
 - Enabled: vectorized multiple scattering, field (not physics)
 - \circ Generate 500 events, 2 electrons w/ E = 50 GeV, random directions
 - Keep # events / thread constant (copy & concat 500 generated events)
 - o Use same generated events in G4 and GV
 - o Keep unused threads busy
 - o Disable output
- Machine: FermiCloud VM w/
 Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz, 4096 KB cache
 sse4.2 instructions
- Track wall clock time & memory with CMSSW TimeMemoryInfo tool
 Measures VSIZE, RSS per event
 - Calculate speedup from wall time (divided by # threads used, since # events / thread is constant)

Characterization

- VM CPU has relatively small cache
 - o Known that major component of GeantV speedup arises from smaller library → fewer cache misses
- To characterize CMSSW performance results, first run built-in GeantV FullCMS standalone test
 - Single thread, settings as close to previous slide as possible (see test script: <u>testStandalone.sh</u>)
 - o NB: different physics list used in standalone vs. CMSSW
- Results:
 - o GeantV: RealTime=756.002s CpuTime=753.09s
 - o Geant4: User=1617.36s Real=1618.52s
- \rightarrow 2.14× speedup (standalone)

Time Performance



- GV 2.6× faster than G4 single thread, still \sim 2.2× faster in MT
- GV single track mode similar to basketized
- G4 has better scaling w/ # threads than GV

Memory Performance



- Memory grows ~linearly w/ # threads (expected)
- GV uses more memory than G4 (expected)
- Single track mode uses similar memory to basketized

What Would Be Next?

- To complete the goals of CMS R&D studies for the paper:
 Full magnetic field map
 - o Test on machines w/ different cache sizes
- Stretch goals/notes for future similar projects:
 - o Random number generator
 - o Adapt scoring classes for other detectors (beyond calorimeters)
 - Combine w/ other simulation improvements
 - Notably Russian Roulette & HF shower library, which give largest gains
- If GeantV project were to continue:
 - o Better solution for geometry conversion than TGeo
 - o Sensitive volume/detector functionality
 - Vectorized hadronic physics
 - o Improve threading, memory management, and ownership models
 - o Decouple event loading & task launching in ExternalLoop mode
 - Event-wise scoring rather than current thread-wise scoring w/ TaskData

Conclusions

- CMS studies met ~all goals laid out
 - Co-development led to improvements and bug fixes in GeantV to facilitate experiments' use
 - o One of the first projects to exercise CMSSW ExternalWork feature
 - Physics validation & CPU measurements show very positive results
 - Path to adapt interfaces efficiently is laid out:
 "Rosetta stone" mostly contained in StepWrapper and VolumeWrapper

Geant4	GeantV StepWrapper	
<u>StepWrapper</u>		
<u>VolumeWrapper</u>	<u>VolumeWrapper</u>	

- Demonstrator to test major elements of GeantV-CMSSW integration is ready
 - \circ Up to 2.6× speedup in CMSSW application
 - o Will finalize results for paper

Backup

Template Wrappers

- ➢ Goal: use *exact same* SD code for Geant4 and GeantV
- **Problem:** totally incompatible APIs

o Example: G4Step::GetTotalEnergyDeposit() VS. geant::Track::Edep()

Solution: template wrapper with unified interface
 e.g. StepWrapper<T>::getEnergyDeposit()
 SD and a subscribe the surgement

• SD code *only calls* the wrapper

• Wrapper stores pointer to T (minimize overhead)

- Current wrappers:
 - o BeginRun
 - o BeginEvent
 - o Step
 - o Volume
 - o EndEvent
 - o EndRun

Traits

• Collect Geant4/GeantV-specific types and wrappers into unified **Traits** class: struct G4Traits {

```
typedef G4Step Step;
typedef sim::StepWrapper<Step> StepWrapper;
};
struct GVTraits {
typedef geant::Track Step;
typedef sim::StepWrapper<Step> StepWrapper;
};
```

• Provides standardized typenames to be used by SD class: template <class Traits> class CaloSteppingActionT : ..., public Observer<const typename Traits::Step *> { public: void update(const Step * step) override { update(StepWrapper(step)); } private: // subordinate functions with unified interfaces

```
void update(const StepWrapper& step);
};
```



- SD interface & implementation in **Calo** (.icc file), w/ unimplemented wrapper interfaces
- G4/GV wrapper specializations in CaloG4/GV, w/ specific instances of templated SD class → isolate dependencies

Scoring Approaches

- Two approaches to scoring in CMSSW:
- 1. Inherit from **G4VSensitiveDetector** (Geant4 class) \rightarrow automatically initialized for geometry volumes marked as sensitive
- Inherit from SimWatcher (CMSSW standalone class)
 → need to specify names of watched geometry volumes
- CaloSteppingAction is a demonstrator class w/ approach 2

 Simplified version of ECAL and HCAL scoring
 Less dependent on Geant4 interfaces
- "Real" SD code uses approach 1
- More work to extract Geant4 dependencies will be necessary
 - o Some SD class methods directly from Geant4 (via inheritance)
 - o Need to mock up Geant4-esque interfaces w/ dummy classes for GeantV

More Physics Validation

- 3. Generate 1000 events of single electrons at 2, 10 and 50 GeV at a fixed direction and compare GeantV against Geant4 with magnetic field off and on at 3.8 Tesla
- 4. Generate 100 events of 50 GeV double electrons at 50 GeV with $-3 < \eta < 3$ and $0 < \phi < 2\pi$, run in multi-threaded mode (4 threads), B = 0 Tesla
- 5. Repeat multi-threaded test with B = 3.8 Tesla



- Number of hits is the same for all 3 energies. The differences are at the level of 0.1/0.3/0.2% for 2, 10 and 50 GeV
- The means differ by 0.8/0.6/0.4% at the three energies

3. Energy Deposit with B = 3.8



- Number of hits is the same for all 3 energies. The differences are at the level of 27.7/6.7/1.3% for 2, 10 and 50 GeV
- The means differ by 0.5/1.6/1.7% at the three energies

4. Energy Deposit with B = 0, MT

50 GeV Electrons B = 0 MultiThreads EB (Geant4 vs GeantV)

50 GeV Electrons B = 0 MultiThreads EE (Geant4 vs GeantV)



- Events are generated with 50 GeV electrons having random direction within a limited range of η and ϕ
- The agreement is pretty good in the B=0 option for both # of hits as well as in the shape of the distributions for EB and EE

4. Hit Times with B = 0, MT

50 GeV Electrons B = 0 MultiThreads EB (Geant4 vs GeantV)

50 GeV Electrons B = 0 MultiThreads EE (Geant4 vs GeantV)



• Hit time distributions are also in good agreement for the B=0 option in EB as well as in EE

5. Energy Deposit with B = 3.8, MT

50 GeV Electrons B = 3.8 Tesla MultiThreads EB (Geant4 vs GeantV)

50 GeV Electrons B = 3.8 Tesla MultiThreads EE (Geant4 vs GeantV)



- Same events (50 GeV electrons, random direction within a limited range of η and ϕ) are simulated in a uniform B-field option of 3.8 Tesla
- The agreement is still good for both # of hits as well as in the shape of the distributions for EB and EE

5. Hit Times with B = 3.8, MT

50 GeV Electrons B = 3.8 Tesia MultiThreads EB (Geant4 vs GeantV)

50 GeV Electrons B = 3.8 Tesla MultiThreads EE (Geant4 vs GeantV)



• Hit time distributions are also in reasonable agreement for the B = 3.8 Tesla option in EB as well as in EE

CMS Simulation Optimizations

	Relative CPU usage	
Configuration	MinBias	ttbar
No optimizations	1.00	1.00
Static library	0.95	0.93
Production cuts	0.93	0.97
Tracking cut	0.69	0.88
Time cut	0.95	0.97
Shower library	0.60	0.74
Russian roulette	0.75	0.71
FTFP_BERT_EMM	0.87	0.83
VecGeom (scalar)	0.87	0.93
Mag. field step,track	0.92	0.90
All optimizations	0.16	0.24

- HF shower library, Russian Roulette have largest impacts
- VecGeom, mag. field improvements entered production in past ~year
 - Enabled by validating and using latest Geant4 versions
- Cumulative effects: overall, simulation is 6.2× (4.1×) faster for MinBias (ttbar) vs. default Geant4 settings
- CMS full simulation is at least 8× faster than ATLAS