The ALICE Simulation Report

Ivana Hřivnáčová (IPNO IN2P3/CNRS), Sandro Wenzel (CERN)
For the ALICE Collaboration

23rd Geant4 Collaboration Meeting, 30 August 2018, Lund

The ALICE simulation environment: From Run2 to Run3

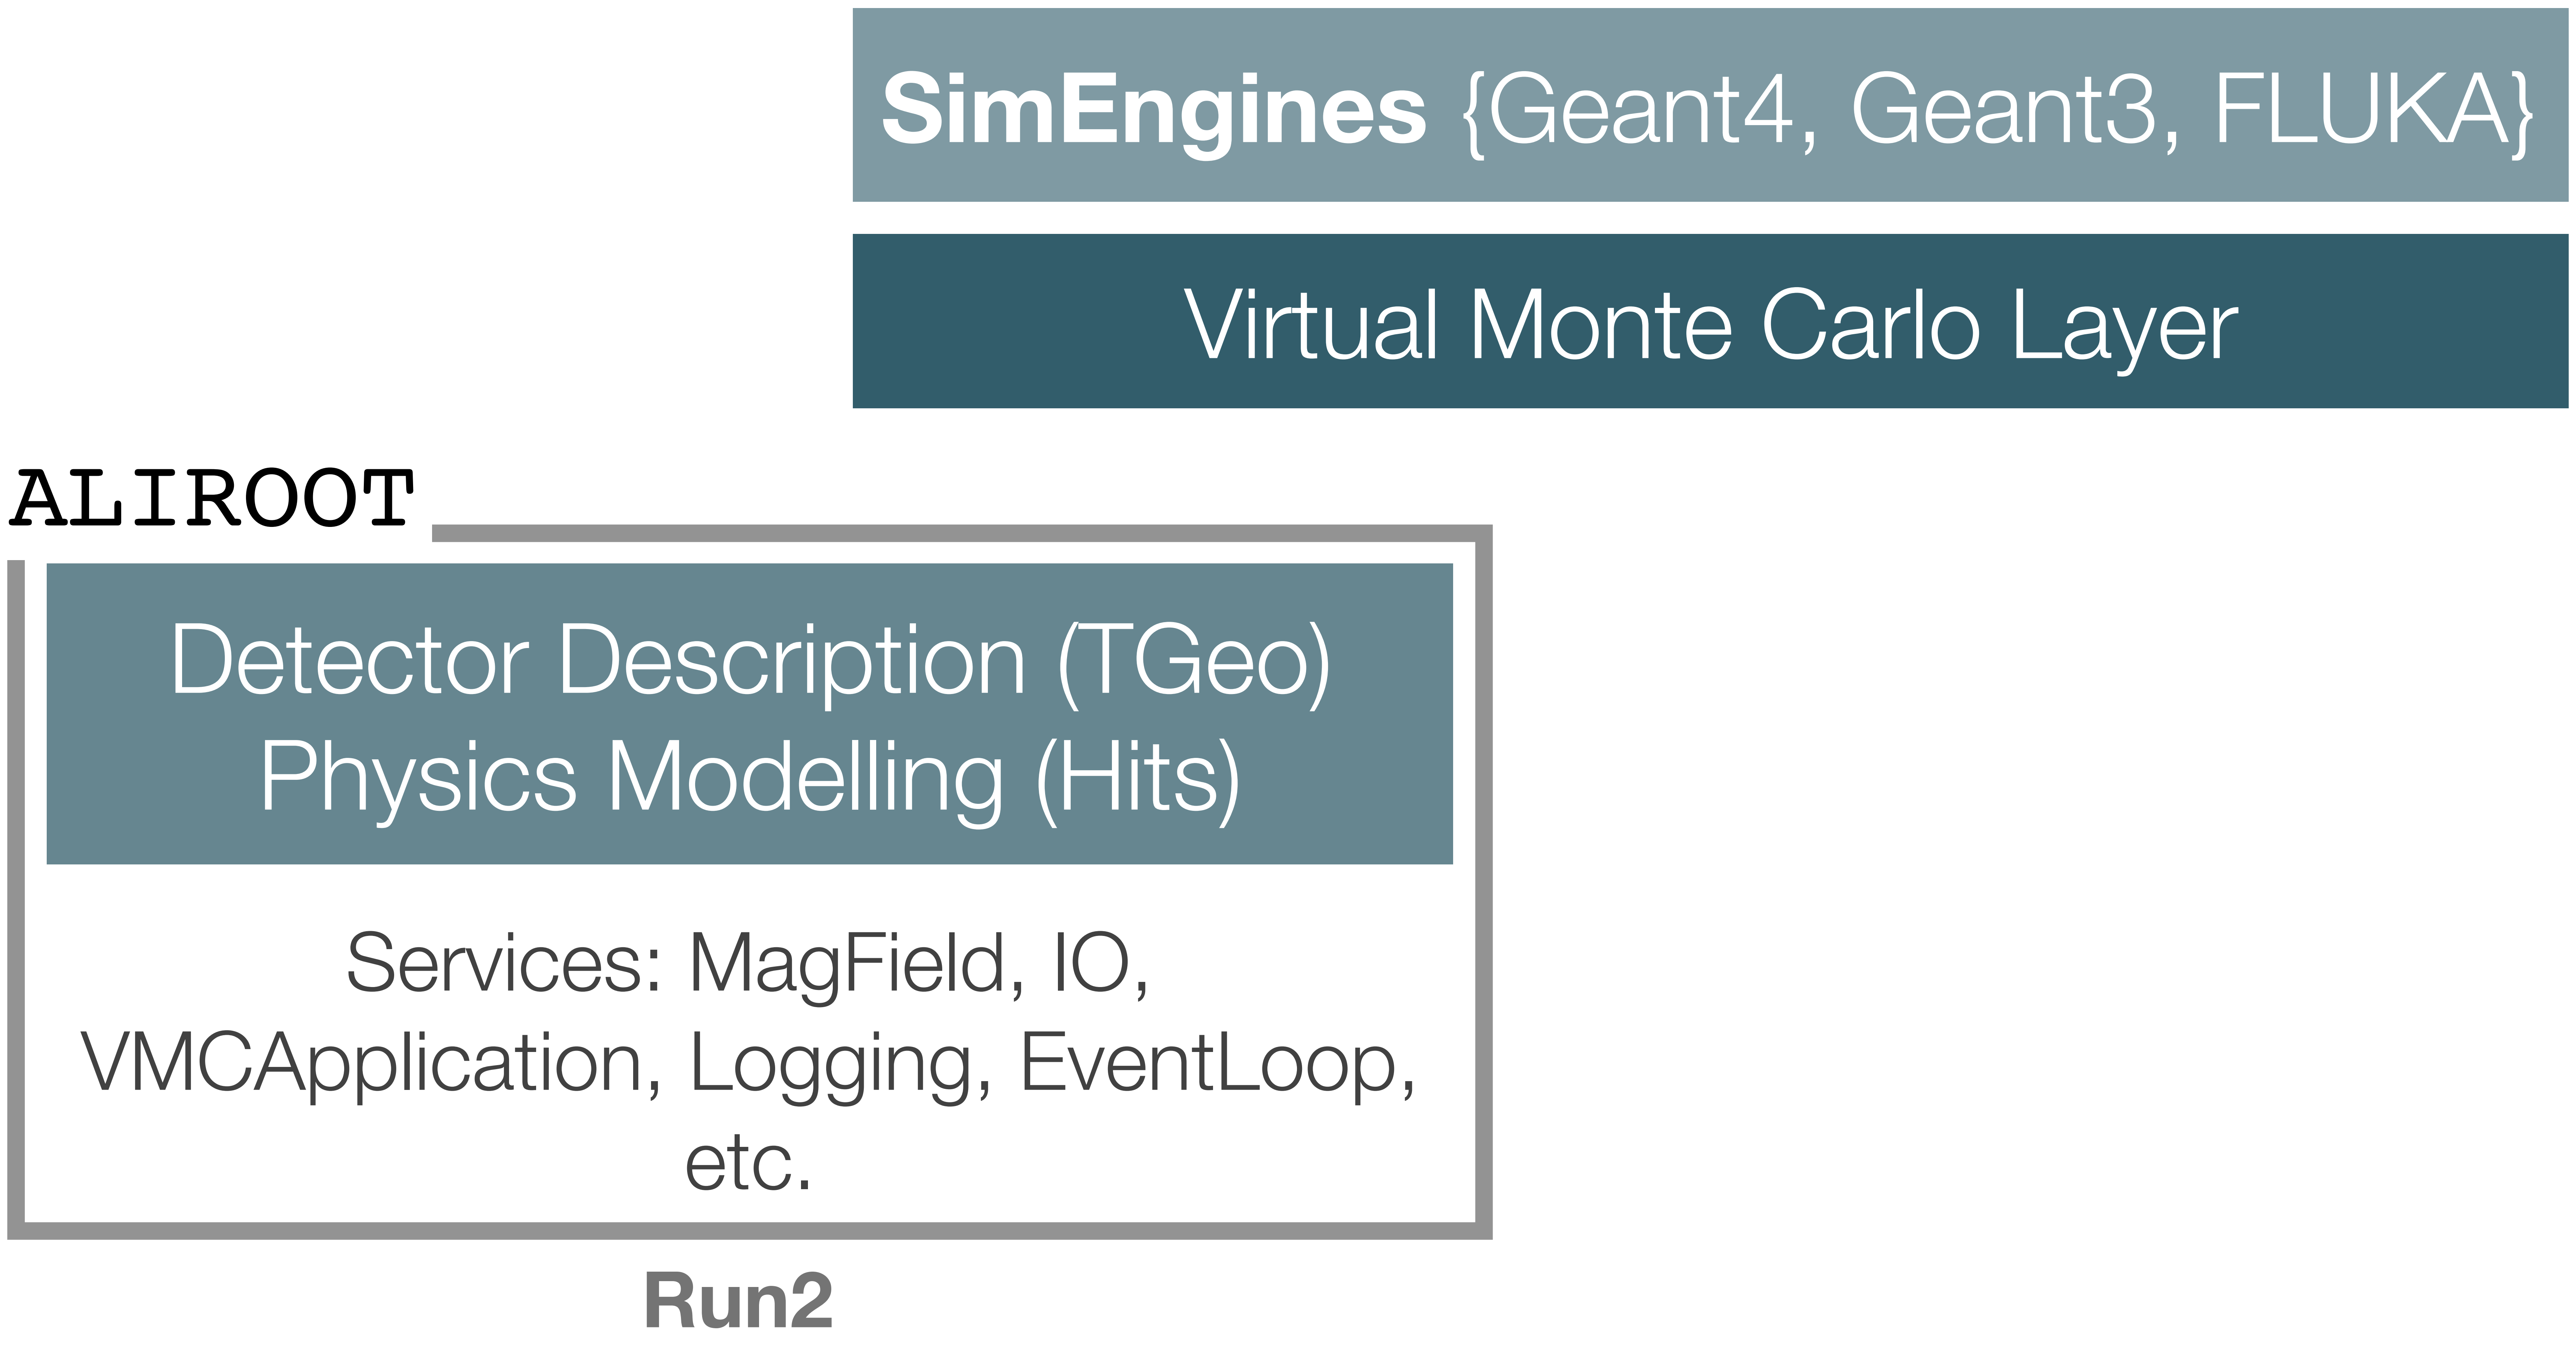

The ALICE simulation environment: From Run2 to Run3

SimEngines {Geant4, Geant3, FLUKA}

Virtual Monte Carlo Layer

ALIROOT

Detector Description (TGeo)
Physics Modelling (Hits)

Services: MagField, IO, VMCApplication, Logging, EventLoop, etc.

Run2

ALICE-O2

Detector Description (TGeo)
Physics Modelling (Hits)

FairRoot

FairMQ

Run3

Keep the big picture; Make in-house code smaller; Potentially benefit from new developments (e.g., FairMQ)

Run 2 - AliRoot
GRID Productions

The list of ALICE Geant4 GRID productions in 2018

- The large effort in the validation of Geant4 simulation
- General purpose productions
- Specific checks, some done together with G3 equivalent ones for comparison

The physics analysis has not yet moved to Geant4 due to remaining issues in the General Purpose productions

Production (2018)	Name	Status	Events generated	Running time	Saving time	Disk size	AliPhysics version
Common productions							
LHC18g8b	Pb-Pb, 5.02 TeV, Run3 Hijing + Nuclei, 0-10% centrality, low field, G4, ALIROOT-7745	Software update	0	-	-	0 B	v5-09-33-01-1
LHC18g8a	Pb-Pb, 5.02 TeV, Run3 Hijing + Nuclei, 0-10% centrality, high field, G4, ALIROOT-7745	Software update	0	-	-	0 B	v5-09-33-01-1
LHC18f2c	p-p, 5.02 TeV - G4 MC for pt resolution studies, w/o delta ray fix, 17p/q anchor, CENT wSDD, ALIROOT-7846	Completed	109600	149d 8:58	14:40	46.41 GB	v5-09-20b-01-1
LHC18f2b	p-p, 5.02 TeV - G4 MC for pt resolution studies, 17p/q anchor, CENT wSDD, ALIROOT-7846	Completed	115200	149d 22:20	15:49	48.86 GB	v5-09-20e-01-1
LHC18d7_extra	p-p, 5.02 TeV - G4 MC production with FTFP_INCLXX_EMV physics list, 17p/q anchor, CENT wSDD, extra statistics, ALIROOT-7783	Completed	5558400	18y 168d	18d 5:36	1.831 TB	v5-09-20e-01-1
LHC18d7	p-p, 5.02 TeV - G4 MC production with FTFP_INCLXX_EMV physics list, 17p/q anchor, CENT wSDD, ALIROOT-7783	Completed	555200	1y 350d	2d 15:26	187.2 GB	v5-09-20e-01-1
LHC18d5	p-p, 5.02 TeV - MC production with fix for the DRAY option in G4 VMC, LHC17p, with SDD, ALIROOT-7776	Completed	557600	1y 154d	3d 0:18	190.9 GB	v5-09-20e-01-1
LHC18c9b	p-p, 5.02 TeV - G4 production with default stepper and PurifyKine off, anchored to 17p/q, ALIROOT-7704	Completed	15340800	32y 74d	62d 19:37	6.73 TB	v5-09-24-01-1
LHC18c3	p-p, 5.02 TeV - G4 production with Nystrom stepper anchored to 17p/q, ALIROOT-7707	Completed	15251600	38y 15d	63d 19:35	5.08 TB	v5-09-20c-01-1
Analysis productions							
LHC18b13b2	A-p, 8.16 TeV, MC production for muon Upsilon analysis anchored to 16s, Geant4, ALIROOT-7696	Completed	22256000	271d 18:59	12d 7:56	59.8 GB	v5-09-20b-01-1
LHC18b13a2	p-A, 8.16 TeV, MC production for muon Upsilon analysis anchored to 16r, Geant4, ALIROOT-7696	Completed	21930000	277d 23:21	12d 8:26	59.76 GB	v5-09-20b-01-1
LHC18c5b	p-p, 5.02 TeV - Geant4 MC production for J/Psi, Psi(2S) and Upsilon analysis at forward rapidity in pp 17p/q, ALIROOT-7694	Completed	44135000	1y 116d	22d 16:28	840.1 GB	v5-09-20b-01-1
LHC18b12b_extra	p-p, 5.02 TeV - MC to study differences in the DCA resolution, Geant4 anchored to LHC17p, extra statistics, ALIROOT-7700	Completed	550400	1y 232d	2d 22:53	188.3 GB	v5-09-20b-01-1
LHC18b12b	p-p, 5.02 TeV - MC to study differences in the DCA resolution, Geant4 anchored to LHC17p, ALIROOT-7700	Completed	23600	19d 19:20	2:00	8.251 GB	v5-09-20b-01-1
LHC18a4b_fast	p-p, 5.02 TeV - HF production anchored to LHC17p/q (pp reference run) with Geant4, HFE, FAST, ALIROOT-7659	Completed	30471350	101y 145d	1y 43d	10.51 TB	v5-09-20b-01-1
LHC18a4b_cent	p-p, 5.02 TeV - HF production anchored to LHC17p/q (pp reference run) with Geant4, HFE, CENT, ALIROOT-7659	Completed	17374750	53y 106d	211d 6:36	6.058 TB	v5-09-20b-01-1
LHC18a4a_fast	p-p, 5.02 TeV - HF production anchored to LHC17p/q (pp reference run) with Geant4, D2H, FAST, ALIROOT-7659	Completed	46136100	131y 67d	1y 153d	14.68 TB	v5-09-20b-01-1
LHC18a4a_cent	p-p, 5.02 TeV - HF production anchored to LHC17p/q (pp reference run) with Geant4, D2H, CENT, ALIROOT-7659	Completed	25926900	74y 203d	337d 8:48	8.332 TB	v5-09-20b-01-1

Run 2 - AliRoot
GRID Productions (2)

- The last Geant4 big production (LHC18d7_extra):
  - **5.5M events, took 18y 168d**
  - Geant4 10.1.p03, FTFP_BERT_EMV physics list with a special MSC model in EMCAL region, with the default Geant4 stepper
- The latest productions configuration
  - Switch from FTFP_BERT_EMV to FTFP_INCLXX_EMV physics list (~20% slowdown)
    - The FTFP_BERT list shows much less light nuclei (d, t, 3He), which come from secondary particles. This list gave results that are much closer to G3 and data.
  - Switch from the default to Nystrom stepper (~9% speed-up)
- (Further) detailed computing performance studies will be done with the final validated Geant4 configuration.
  - Until then concrete profile numbers are subject to change and not representative

Run 2 - AliRoot

Geant4 Validation - Ongoing

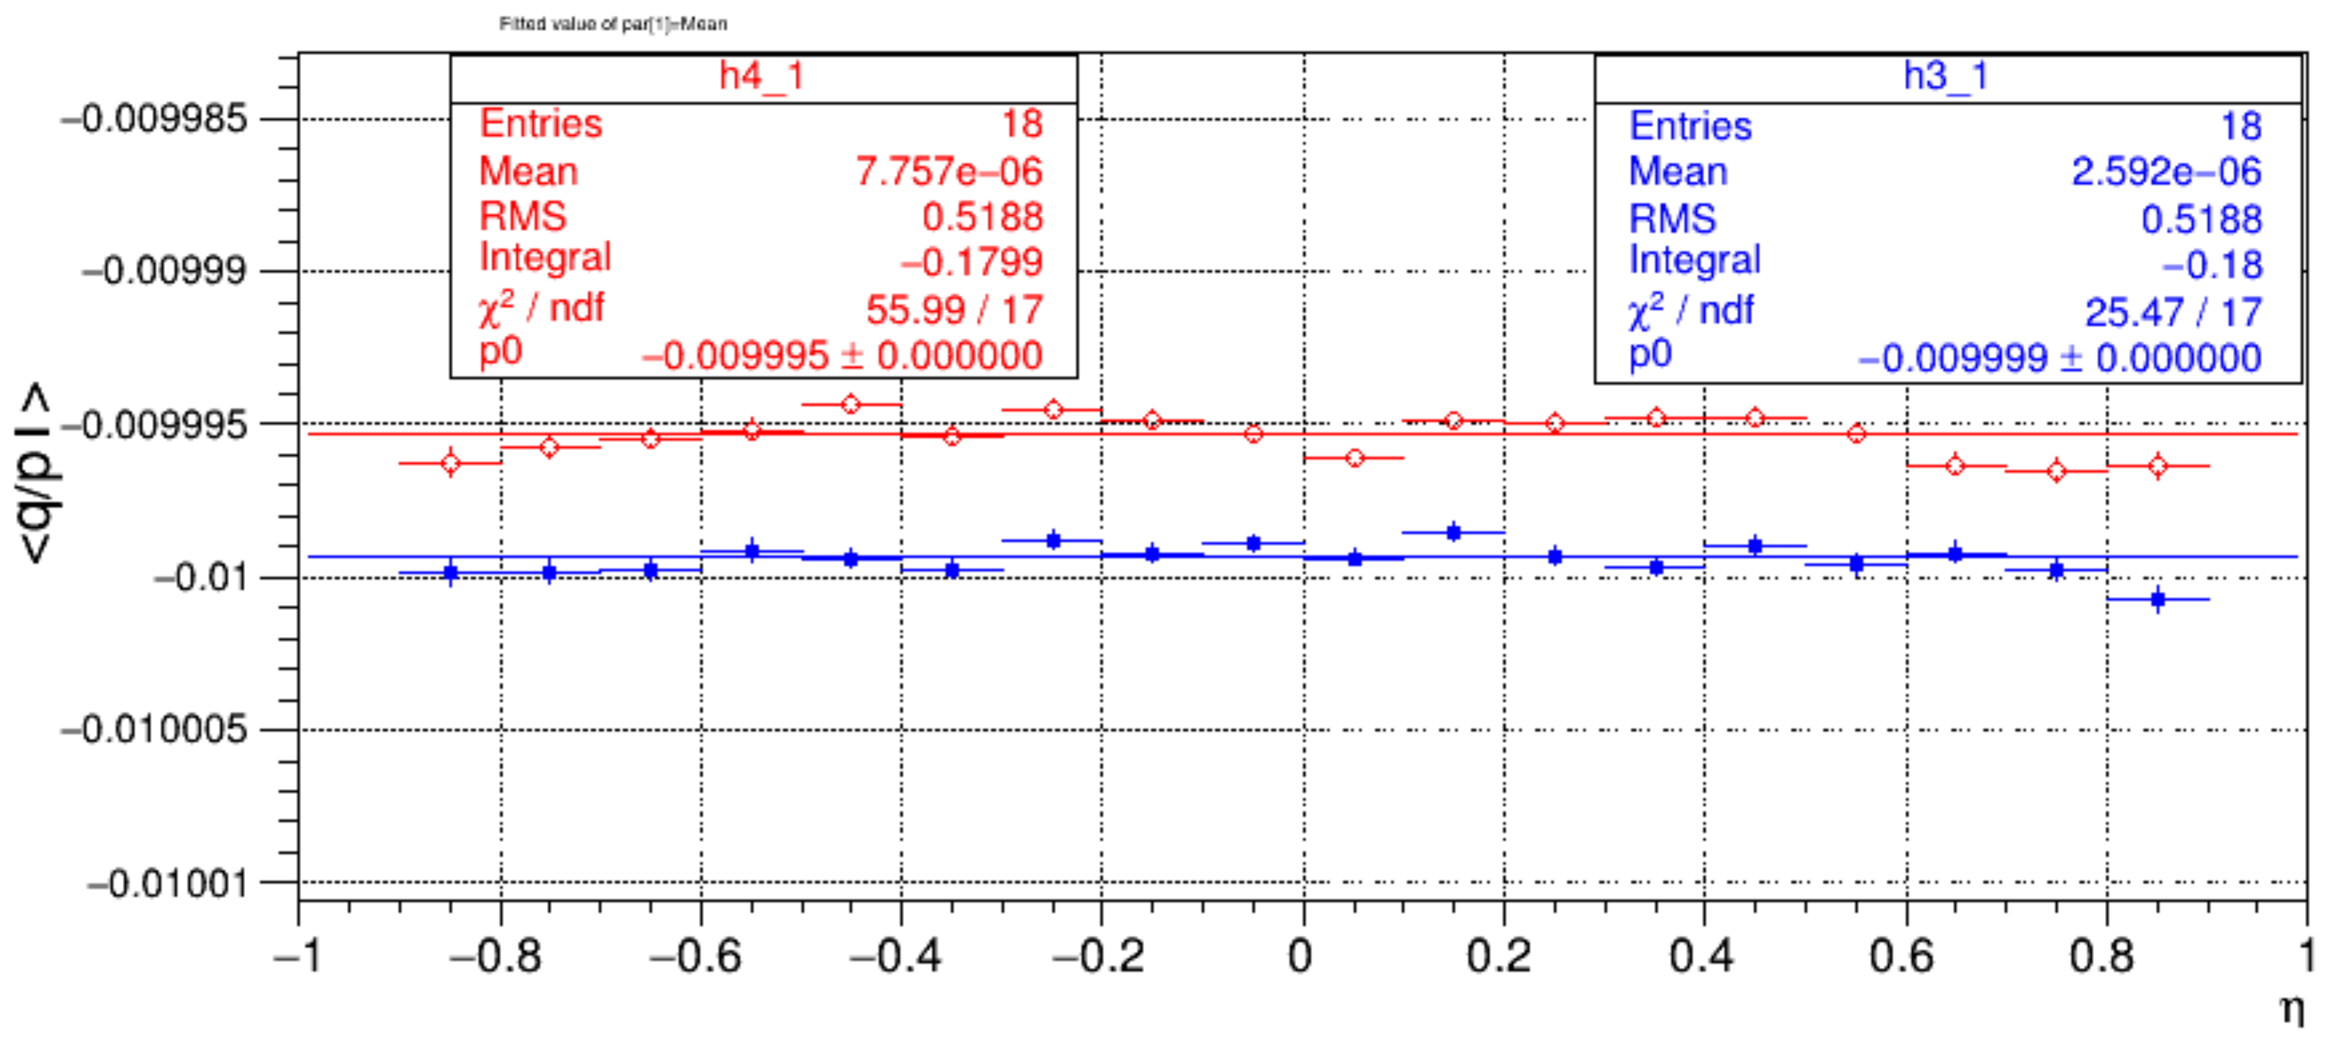

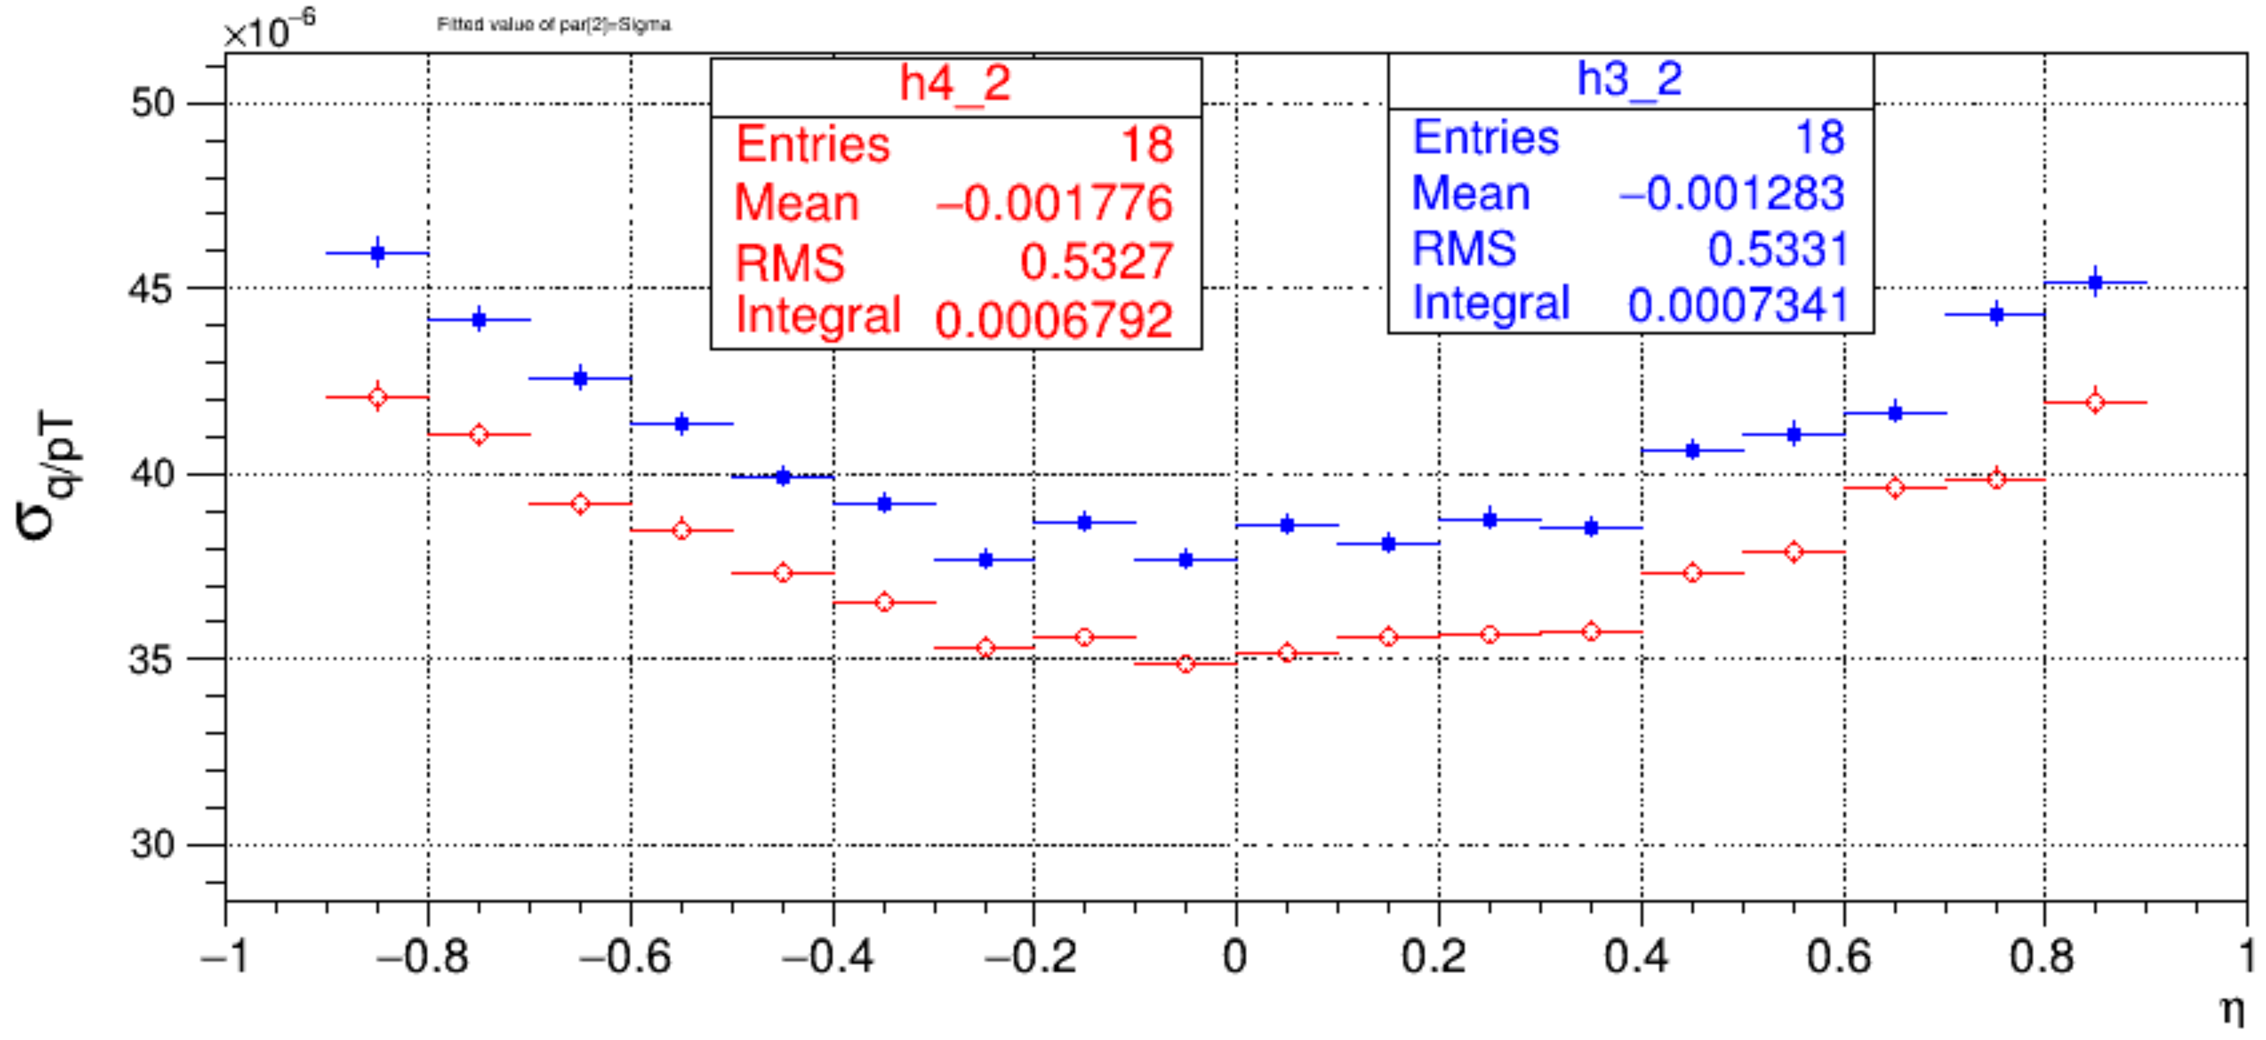

- **q/pT mean and resolution for the pT=100 GeV muons (i.e. q/pT_true = -0.01) vs eta.**
- Obtained by Kalman-fitting the MC track positions along the transported particles, accounting for the energy loss.
- The resolution (width of the q/pT) is somewhat better for G4 compared to G3, presumably due to the smaller multiple scattering, but *there is a clear bias in <q/pT>*
    - We suspect it is transport-related, would be interesting to redo this exercise with constant Bz field)

Run 2 - AliRoot

Geant4 Validation - Ongoing (2)

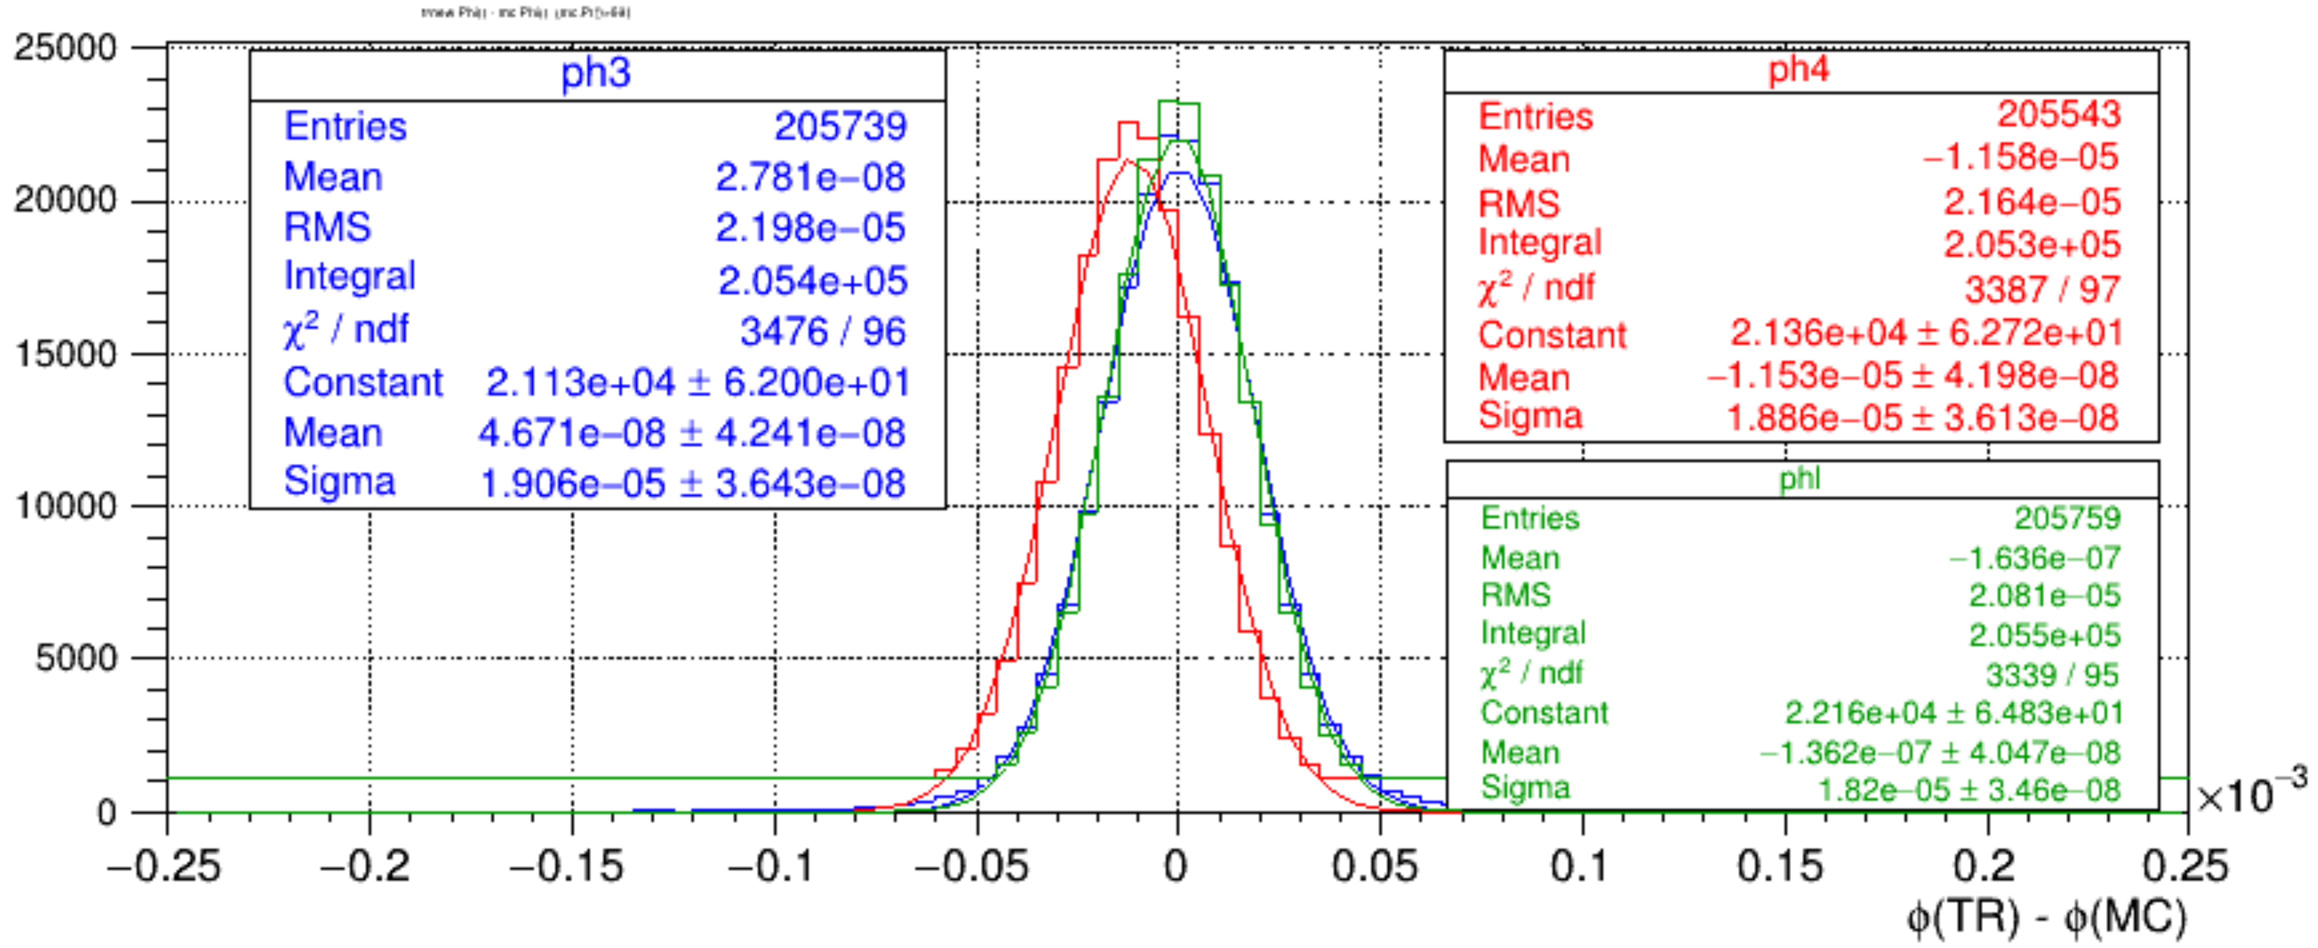

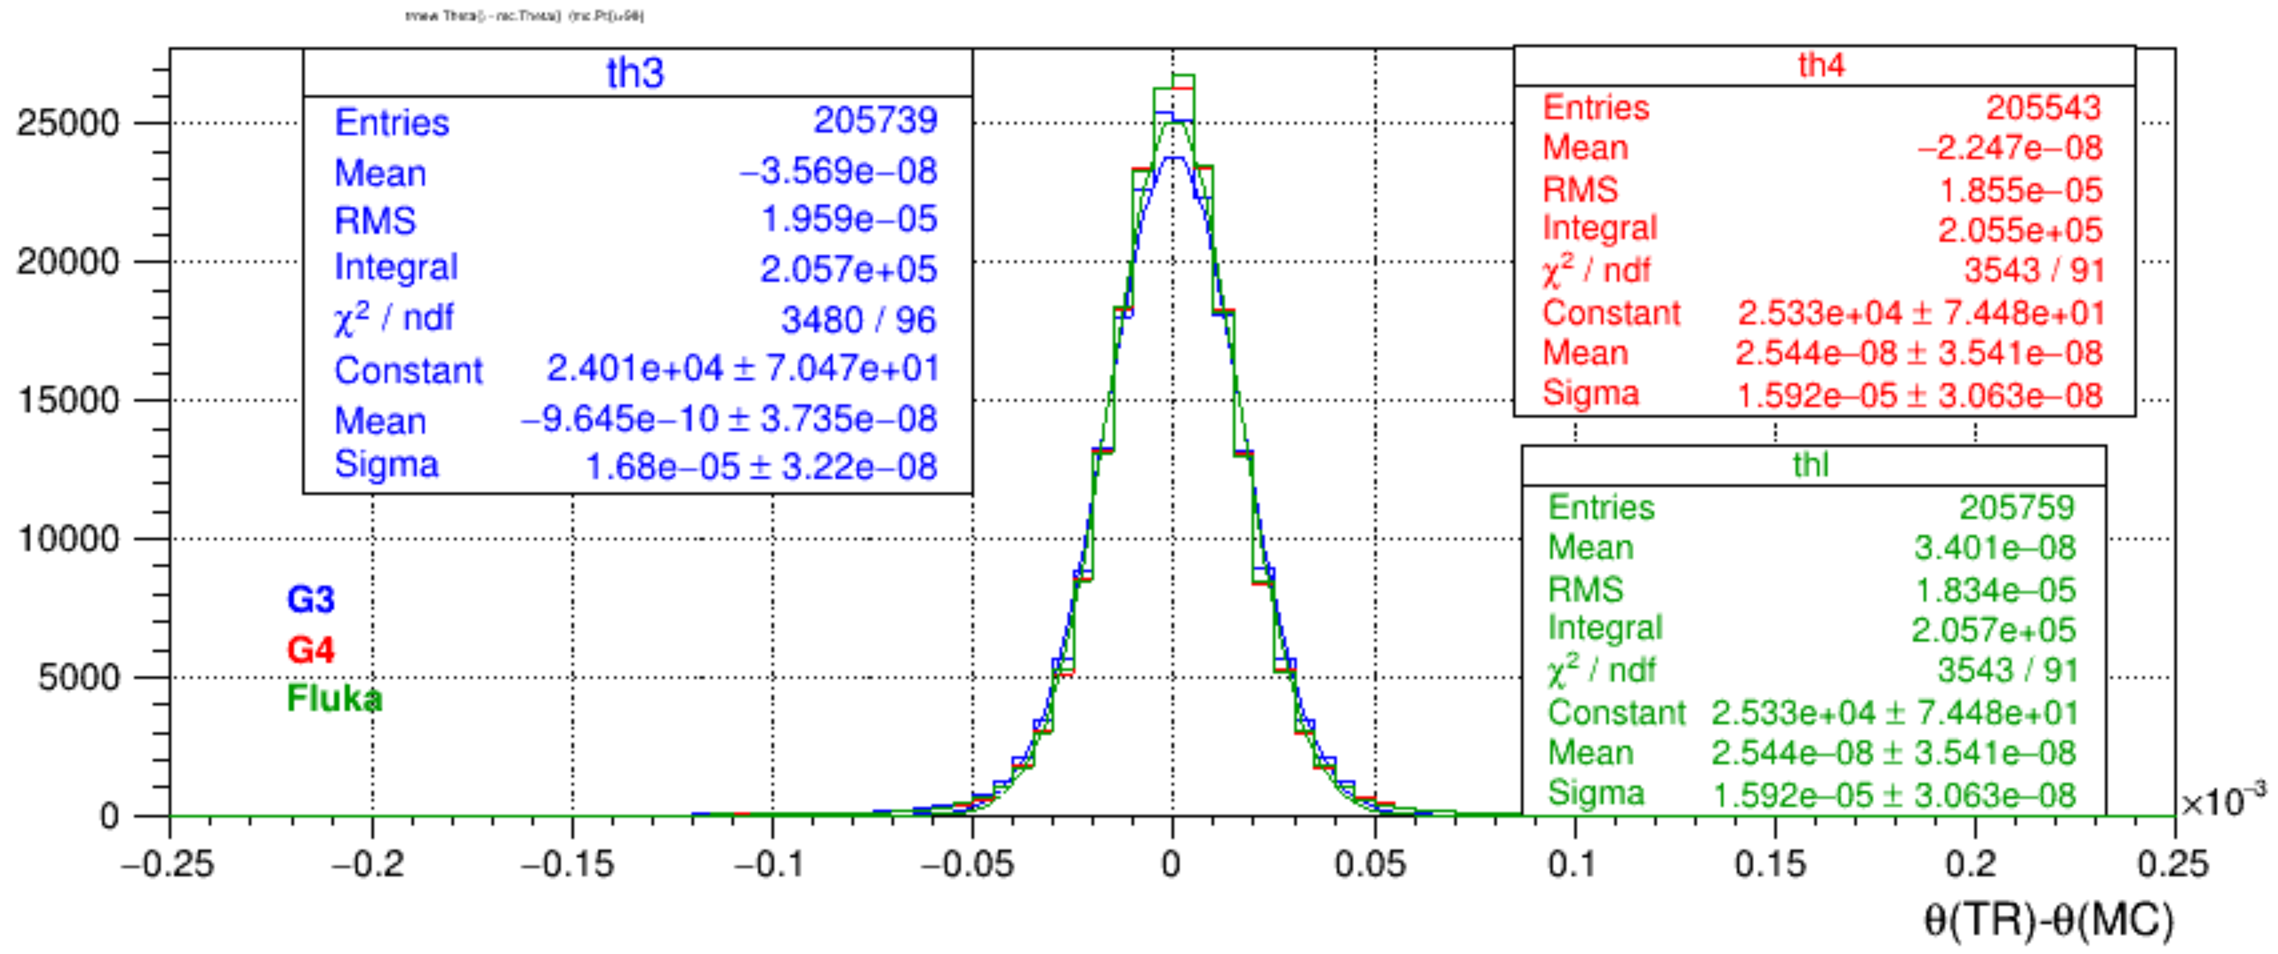

- **The difference in the azimuthal and polar angle between generated and refitted track in the MC track emission point.**
- Again, in the bending direction G4 shows a small but clear bias.

Run 2 - AliRoot

Current

- Pb-Pb collision can be very demanding: may have up to 100k primaries in the collision to transport
  - ~O(GBs) of memory / event
  - ~O(h) of CPU time / event
- Bad for scheduling and efficiently using given resource (packing problem)
- Prevents access to (opportunistic) HPC

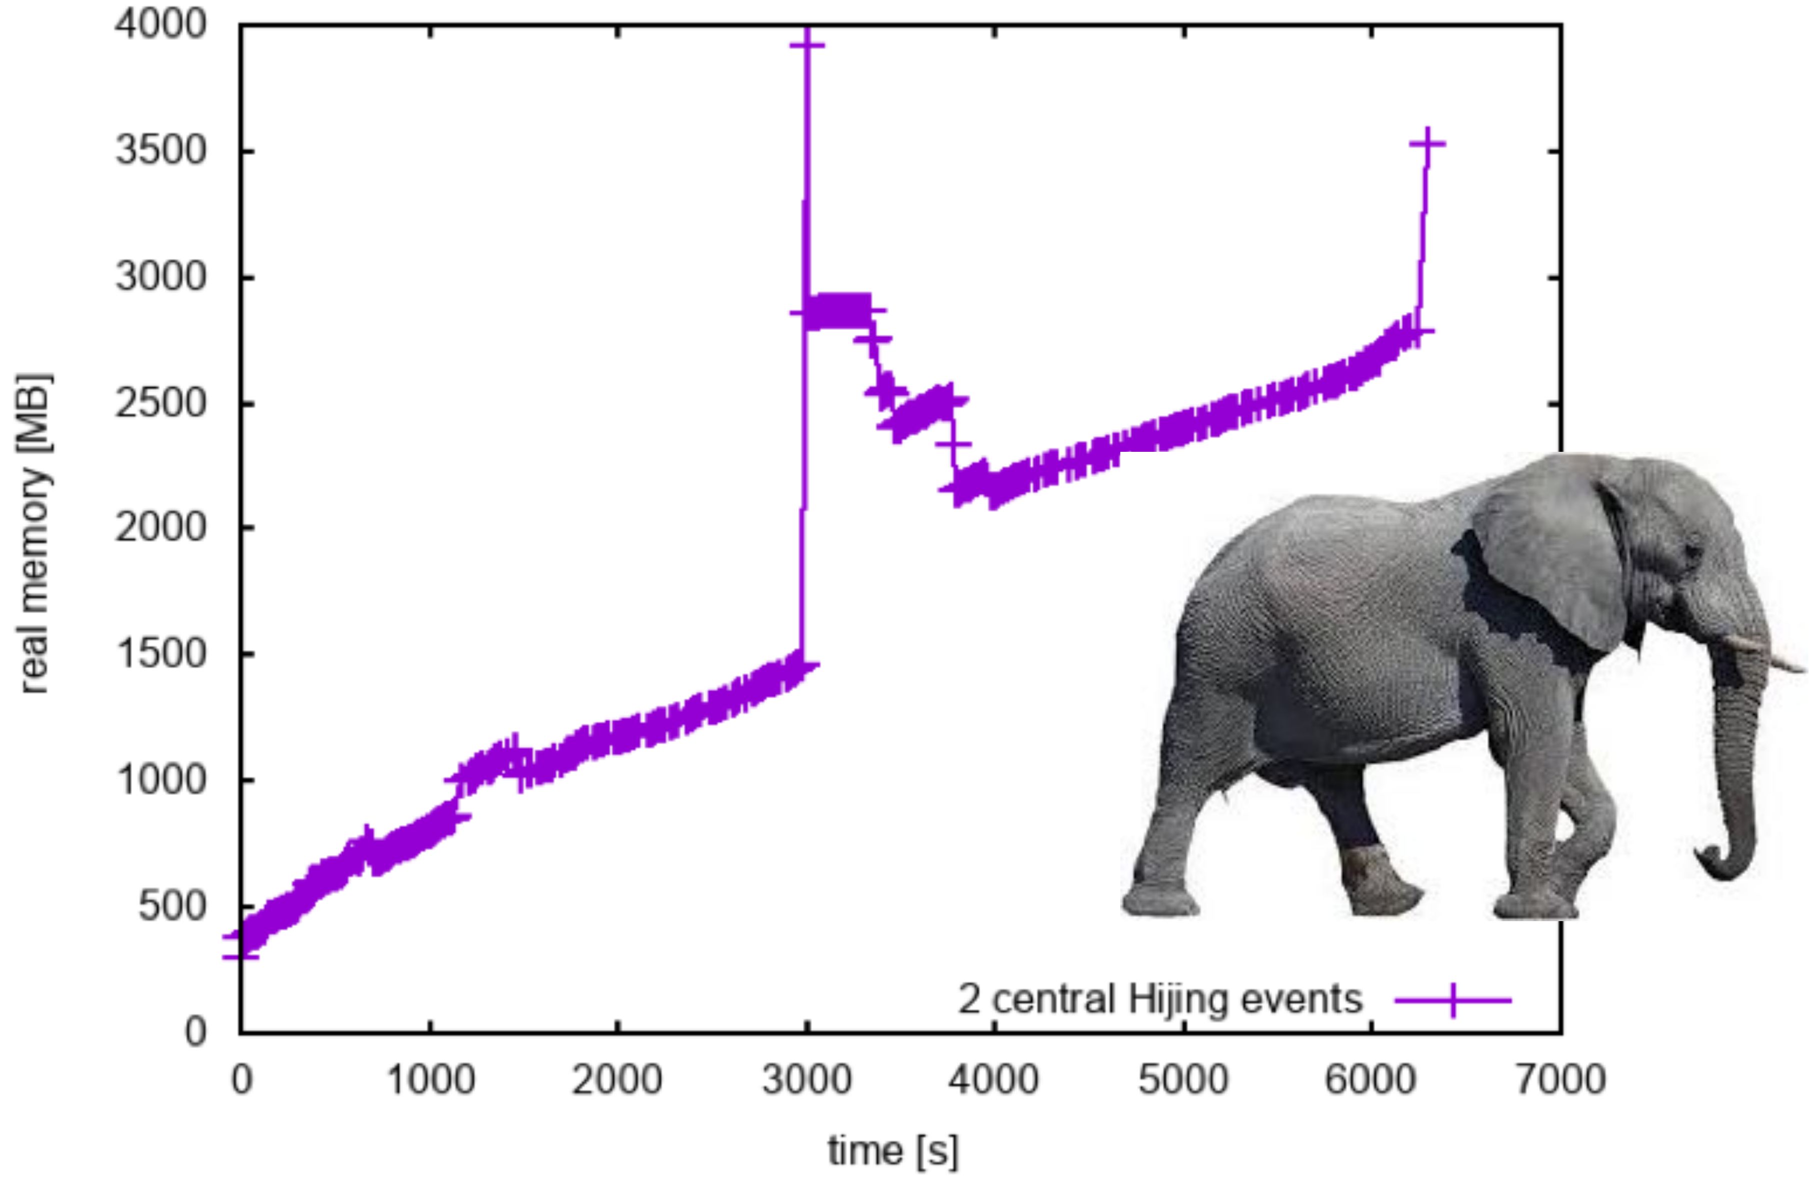

Run 3 - ALICE-O2

Future

- **Goal: A simulation system running on anything from laptop to many-core and HPC facilities**
- Ingredients put forward here:
  - **Independent actors** based on heterogenous multi-processing and message passing
  - **Event splitting** and **collaborative simulation parallelism**

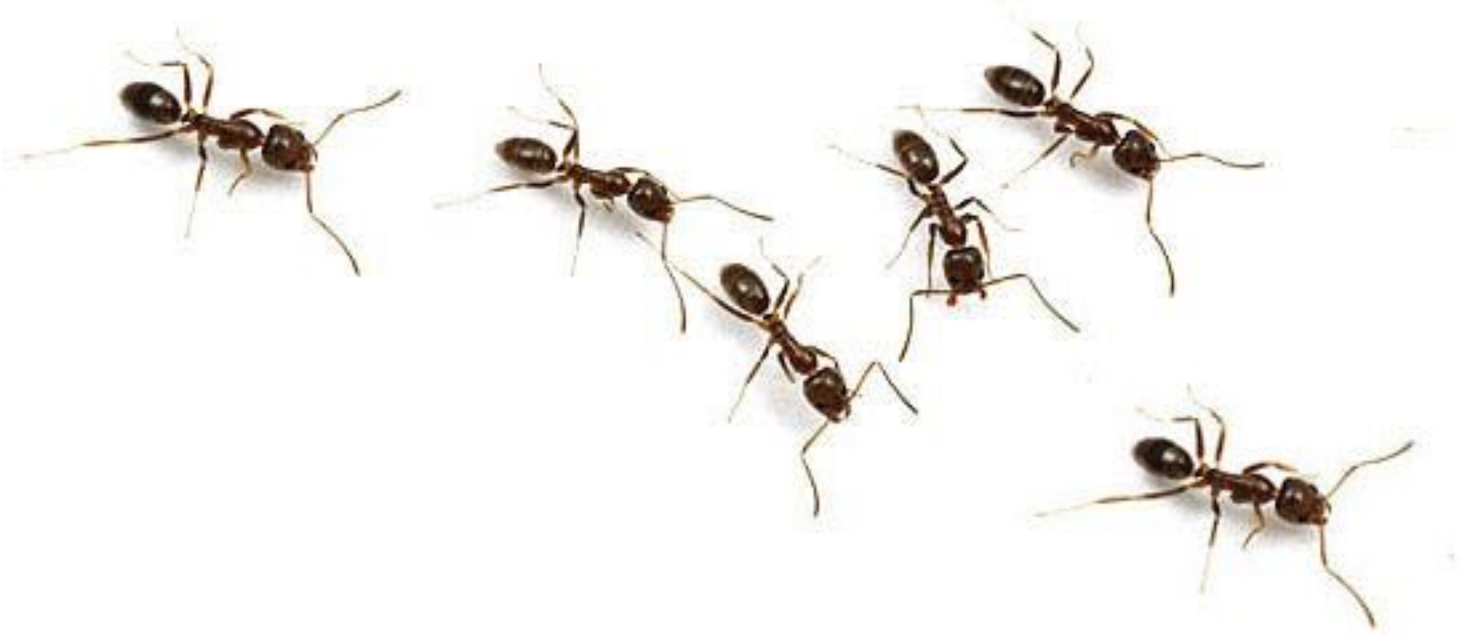

Run 3 - ALICE-O2

Multi-processing Parallelism

Key Ingredients

1. Break existing monolithic simulation into few actors with specialized concern with FairMQ
    - FairMQ = Fair(MessageQueue) is an abstract messaging library for C++
2. Splitting an event into sub-events
3. Collaborative Parallel Simulation
    - Memory “catastrophe” avoided by copy-on-Write late-forking
    - Works exceptionally well

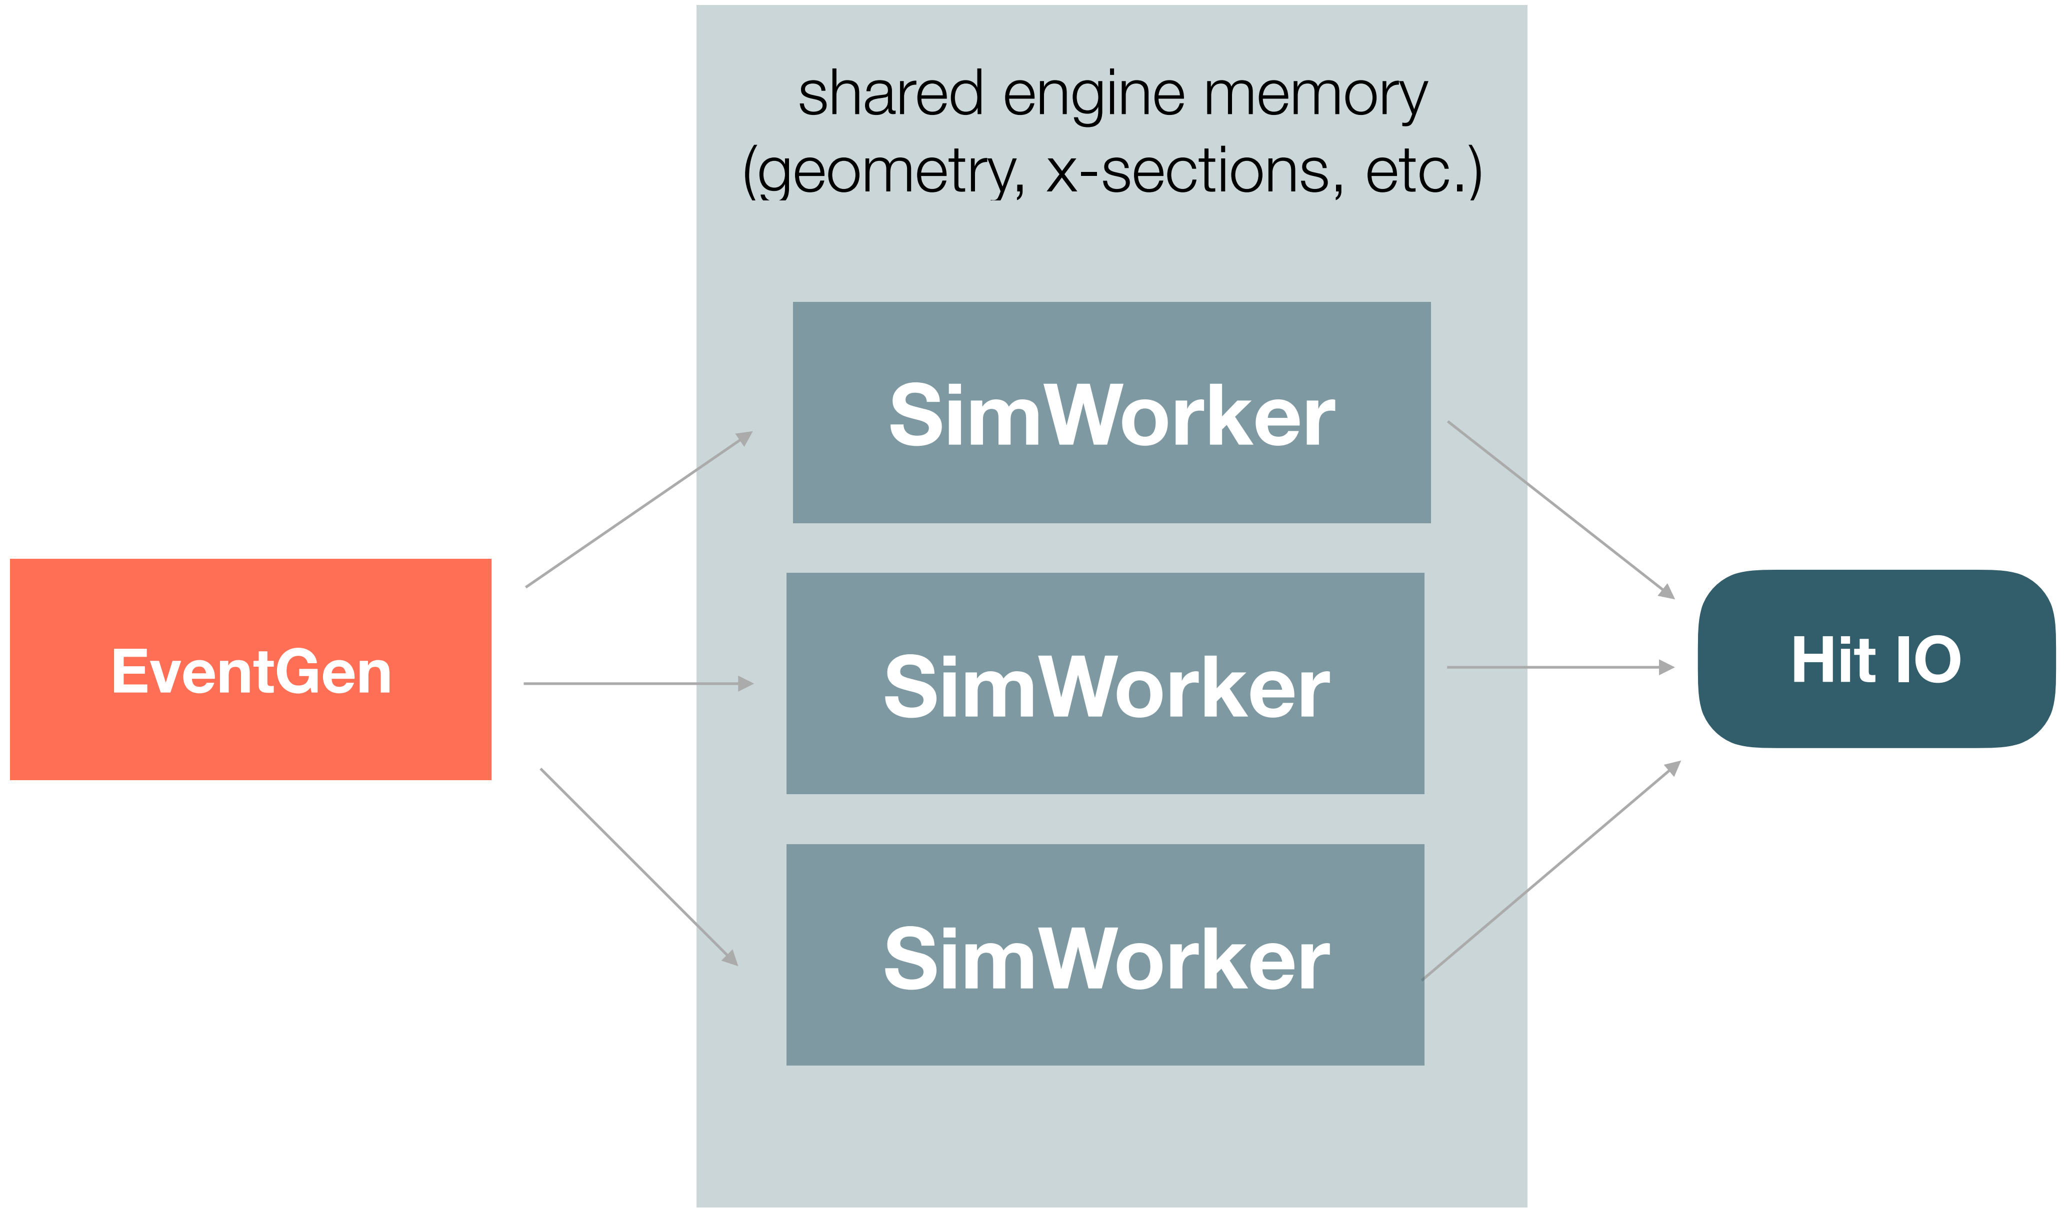

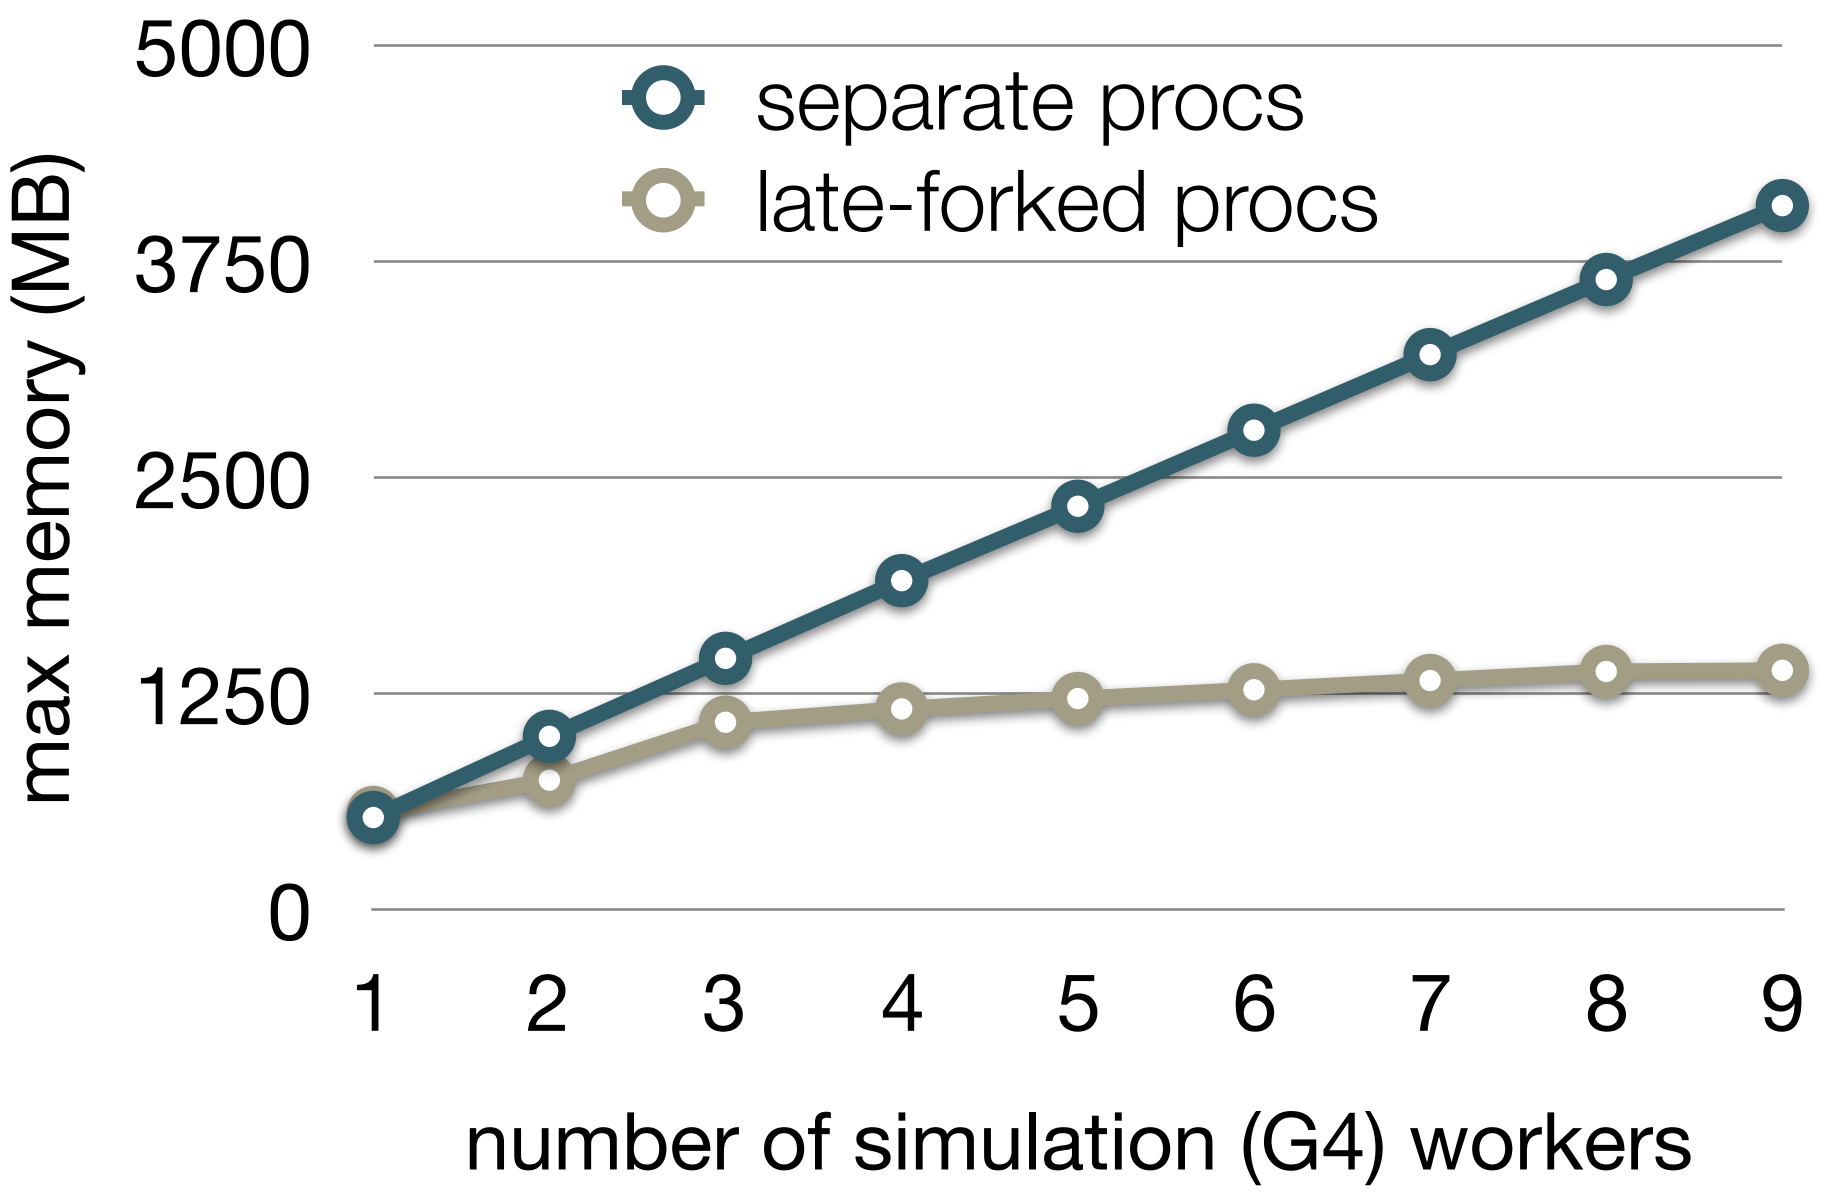

Run 3 - ALICE-O2

Multi-processing Parallelism

Key Ingredients

1. Break existing monolithic simulation into few actors with specialized concern with FairMQ
   - FairMQ = Fair(MessageQueue) is an abstract messaging library for C++
2. Splitting an event into sub-events
3. Collaborative Parallel Simulation

See more at the CHEP 2018 talk:
S. Wenzel: A scalable and asynchronous detector simulation system based on ALFA

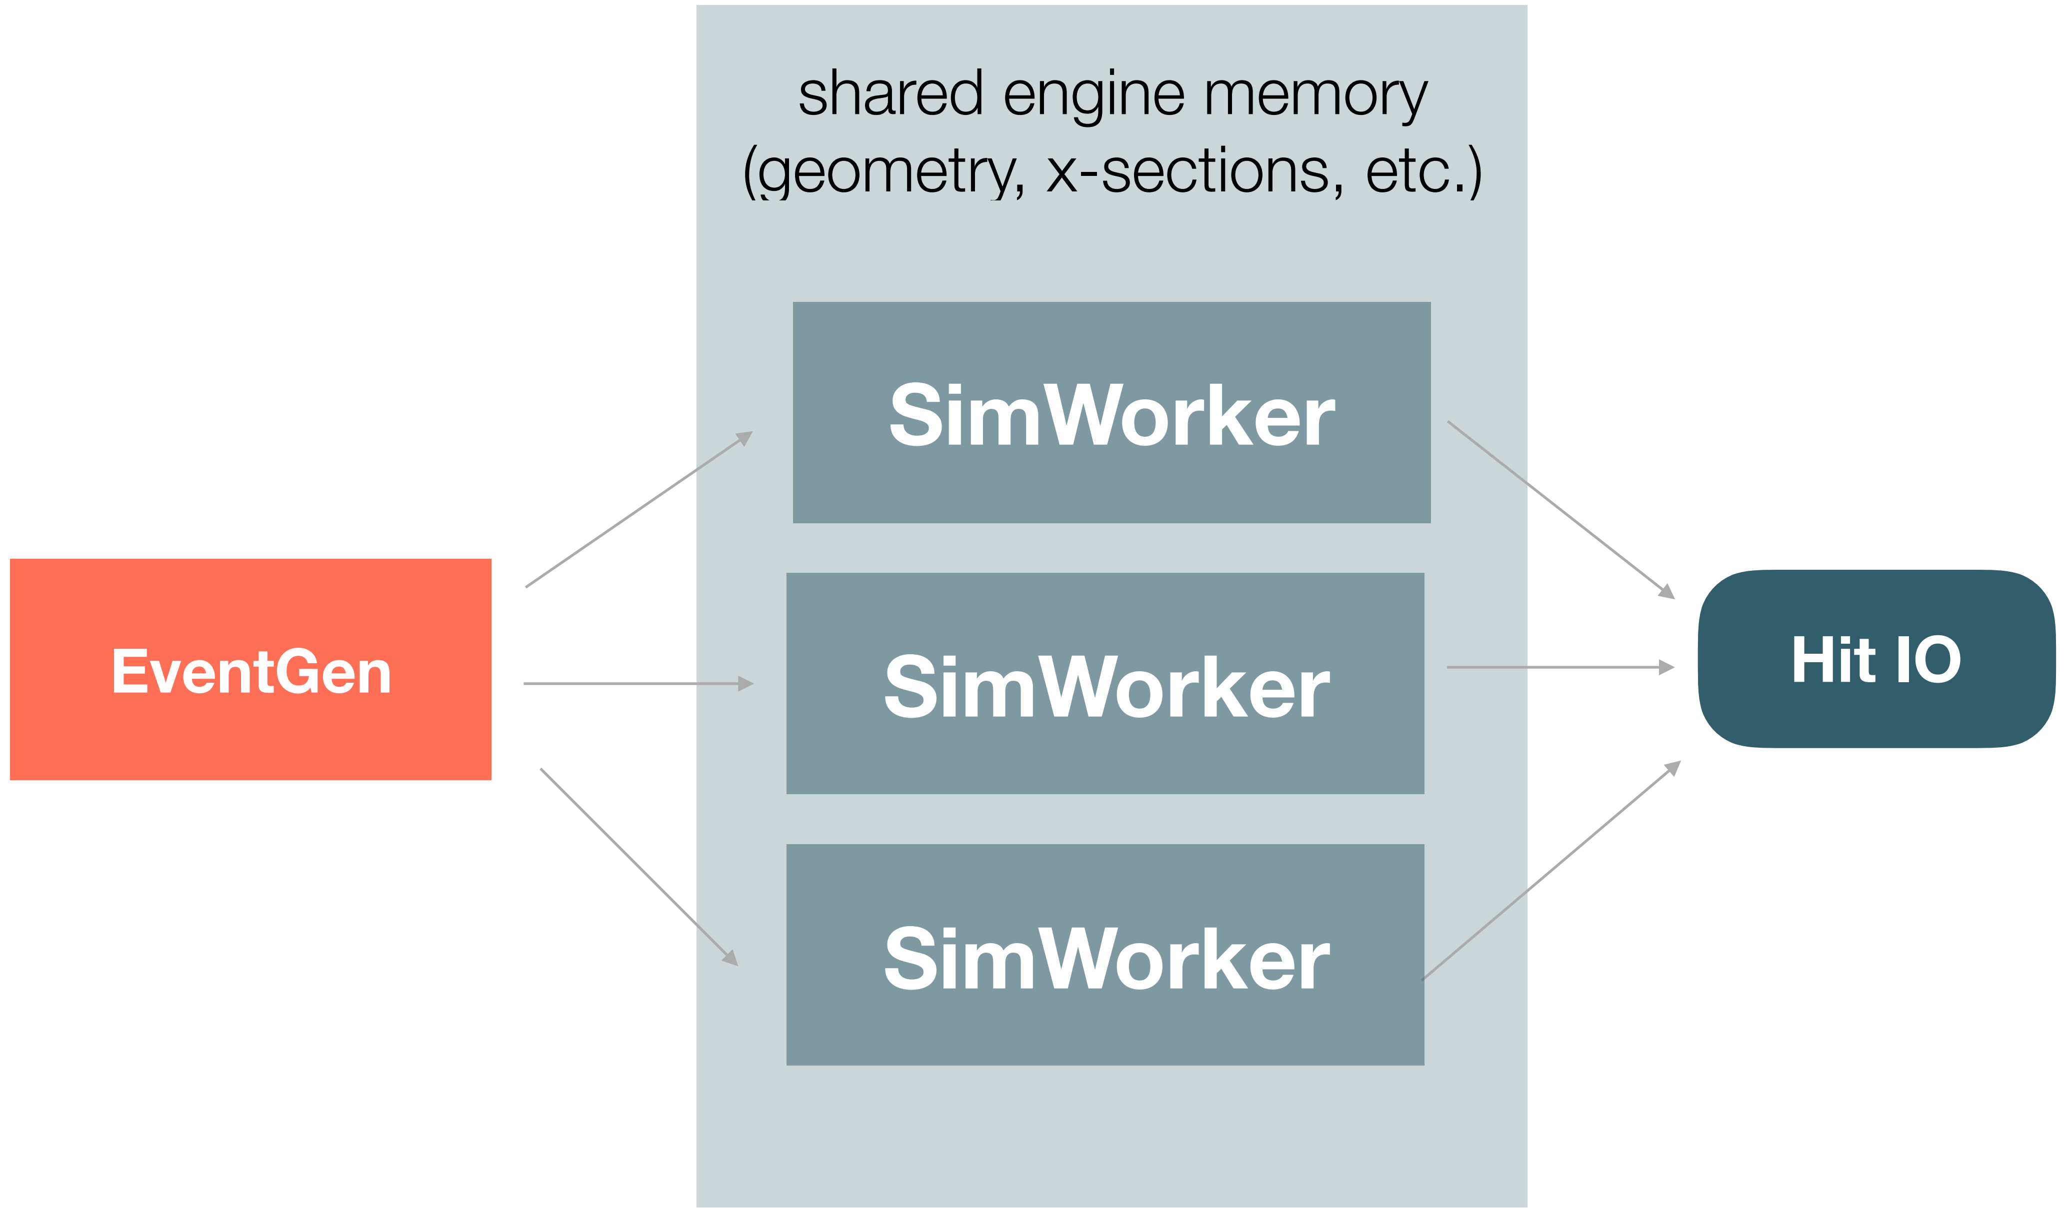

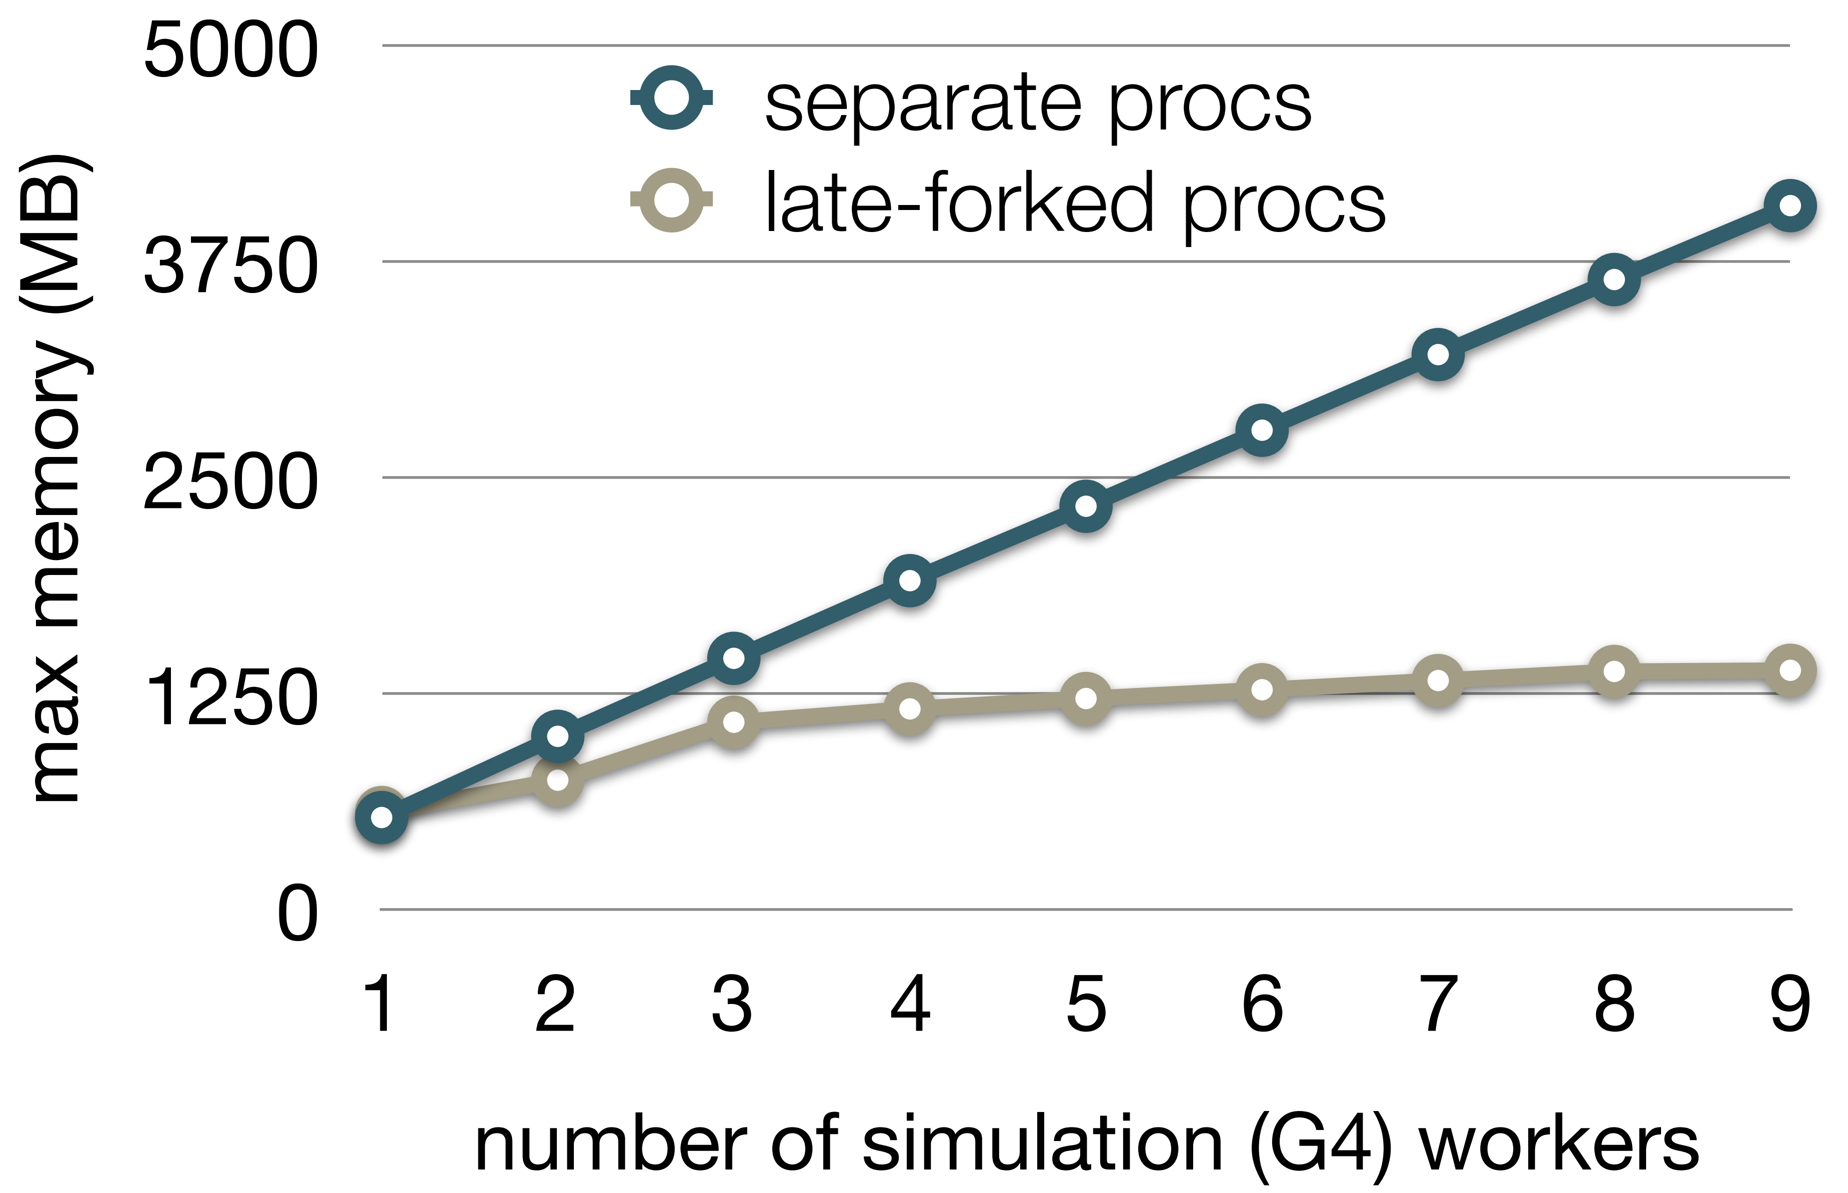

Run 3 - ALICE-O2

Multi-threading Parallelism

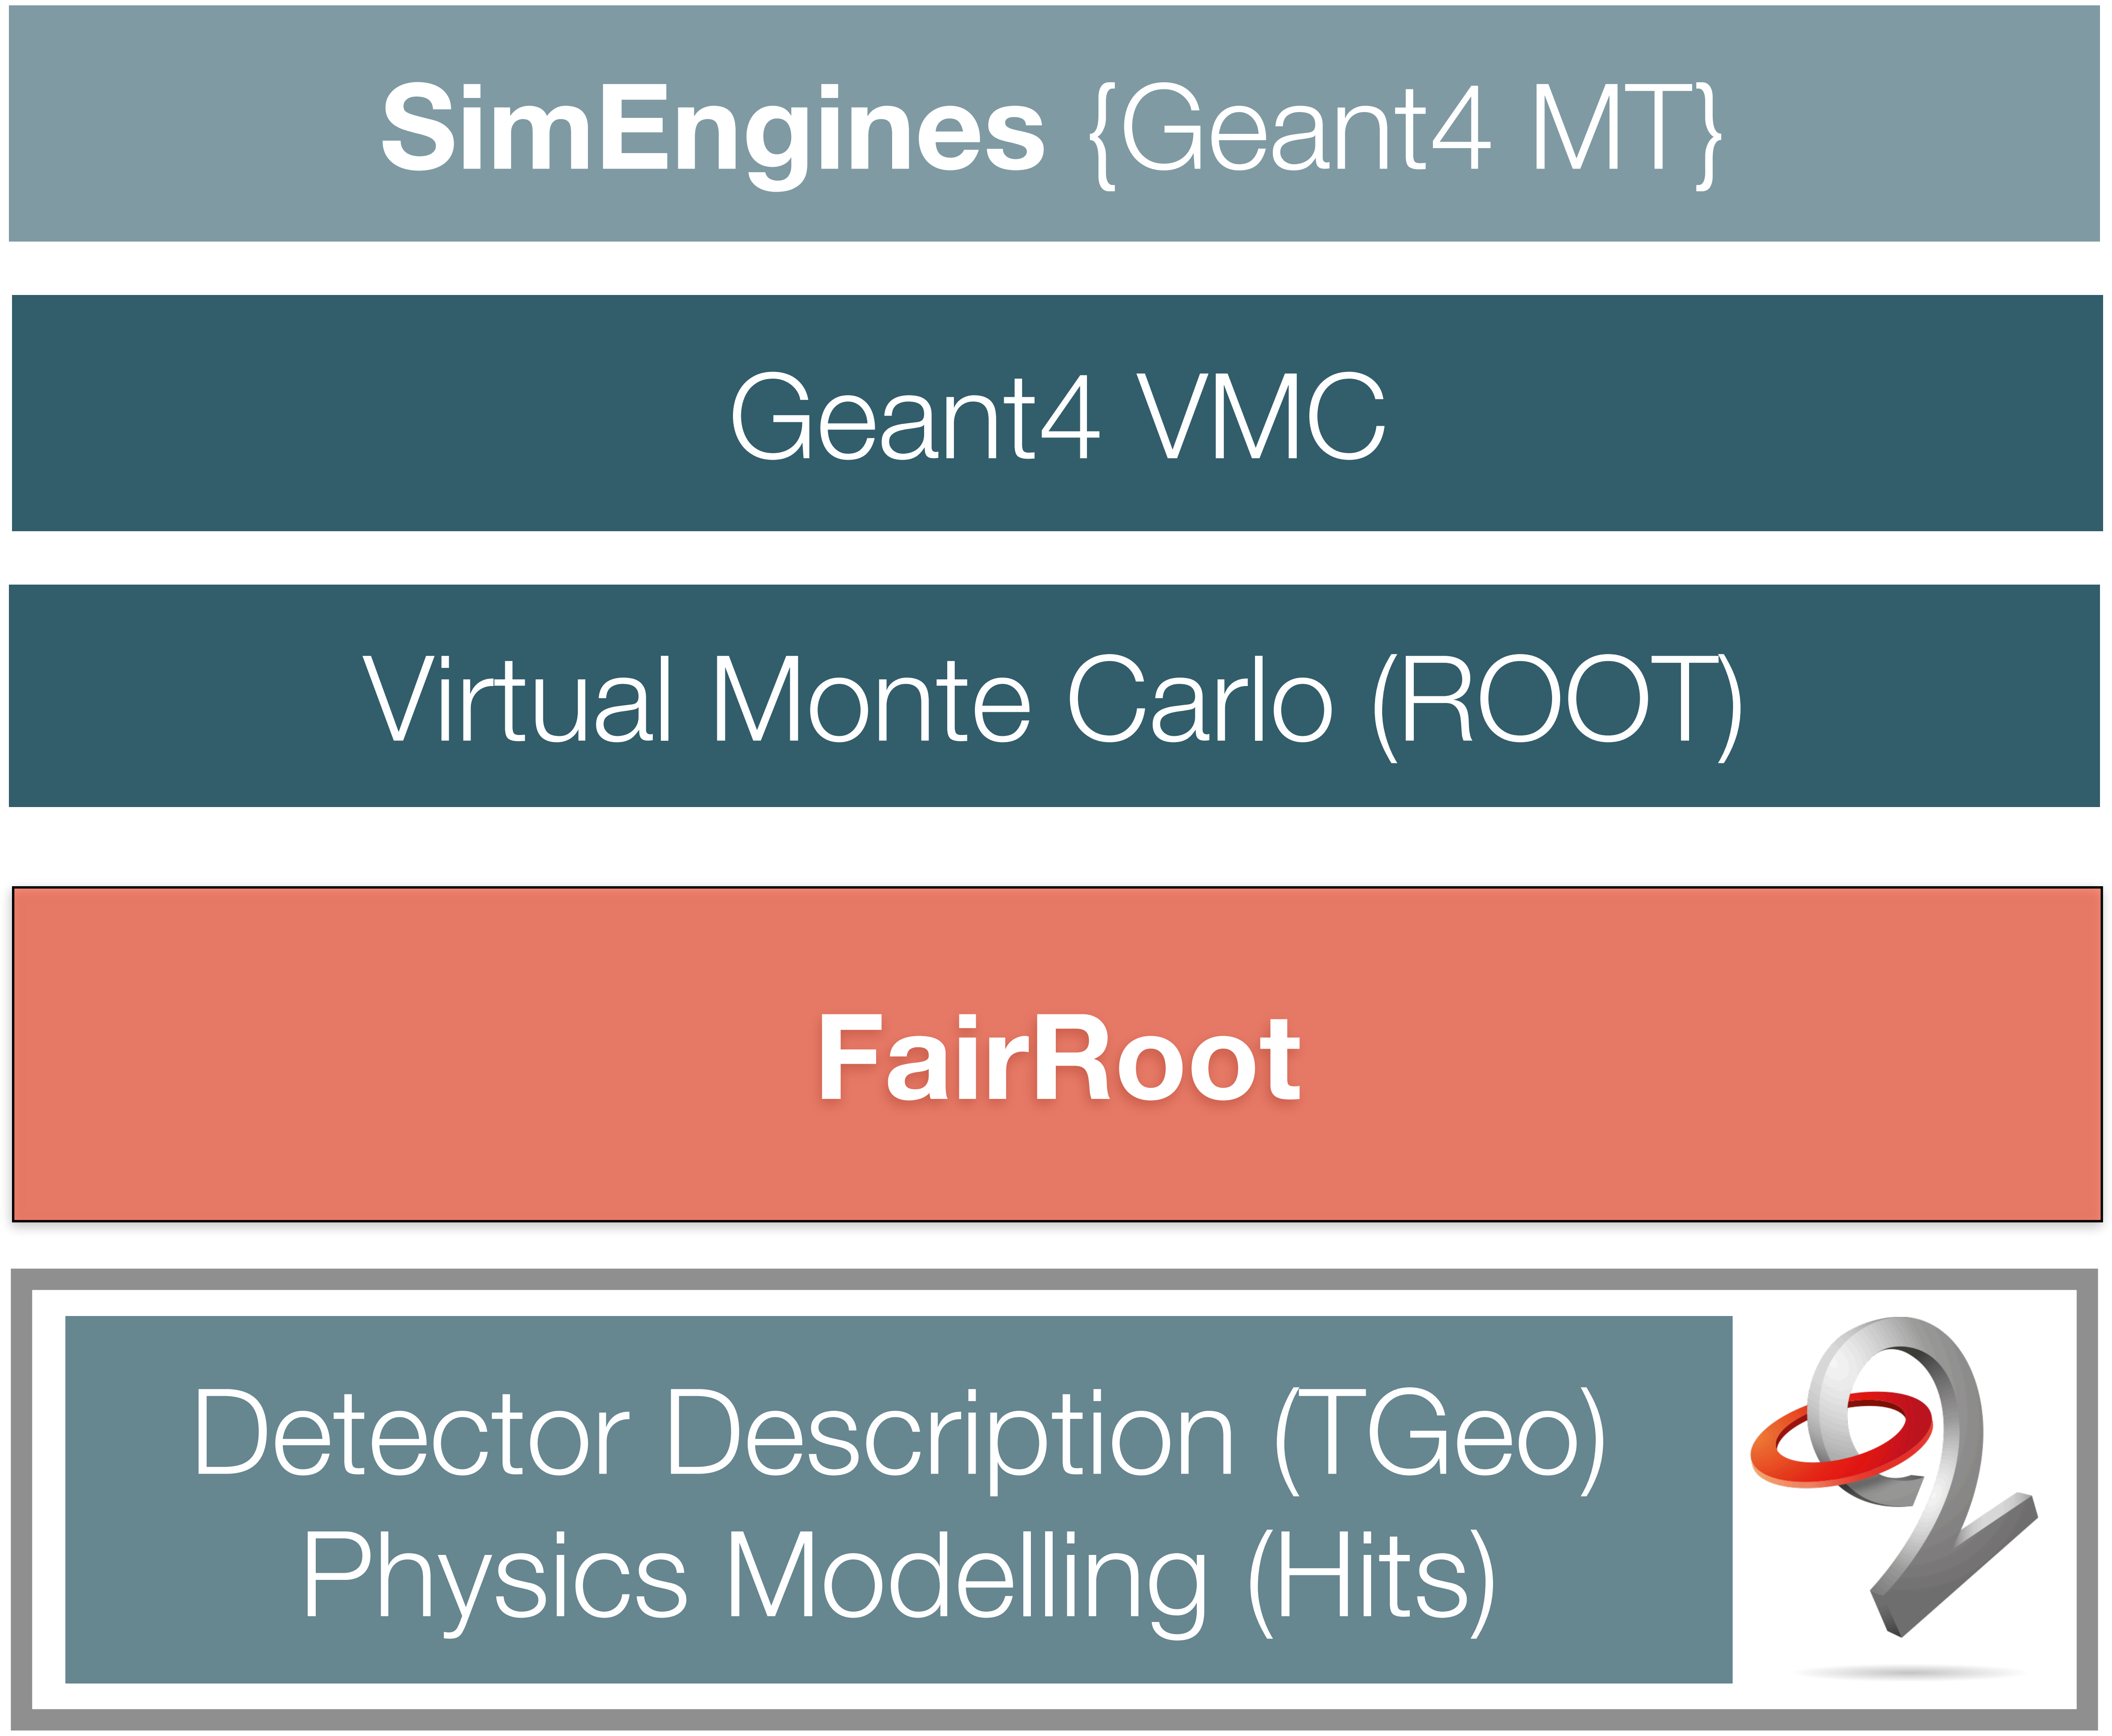

- Complementary to the multi-processing approach
- Event level parallelism
- Migration to MT concerned all packages:
  - VMC (ROOT), Geant4 VMC, FairRoot, O2
- The FairRoot simulation examples and the O2 simulation program were enhanced with the multi-threading run option
- The MT mode was added in the standard O2 tests

Run 3 - ALICE-O2

Multi-threading Parallelism

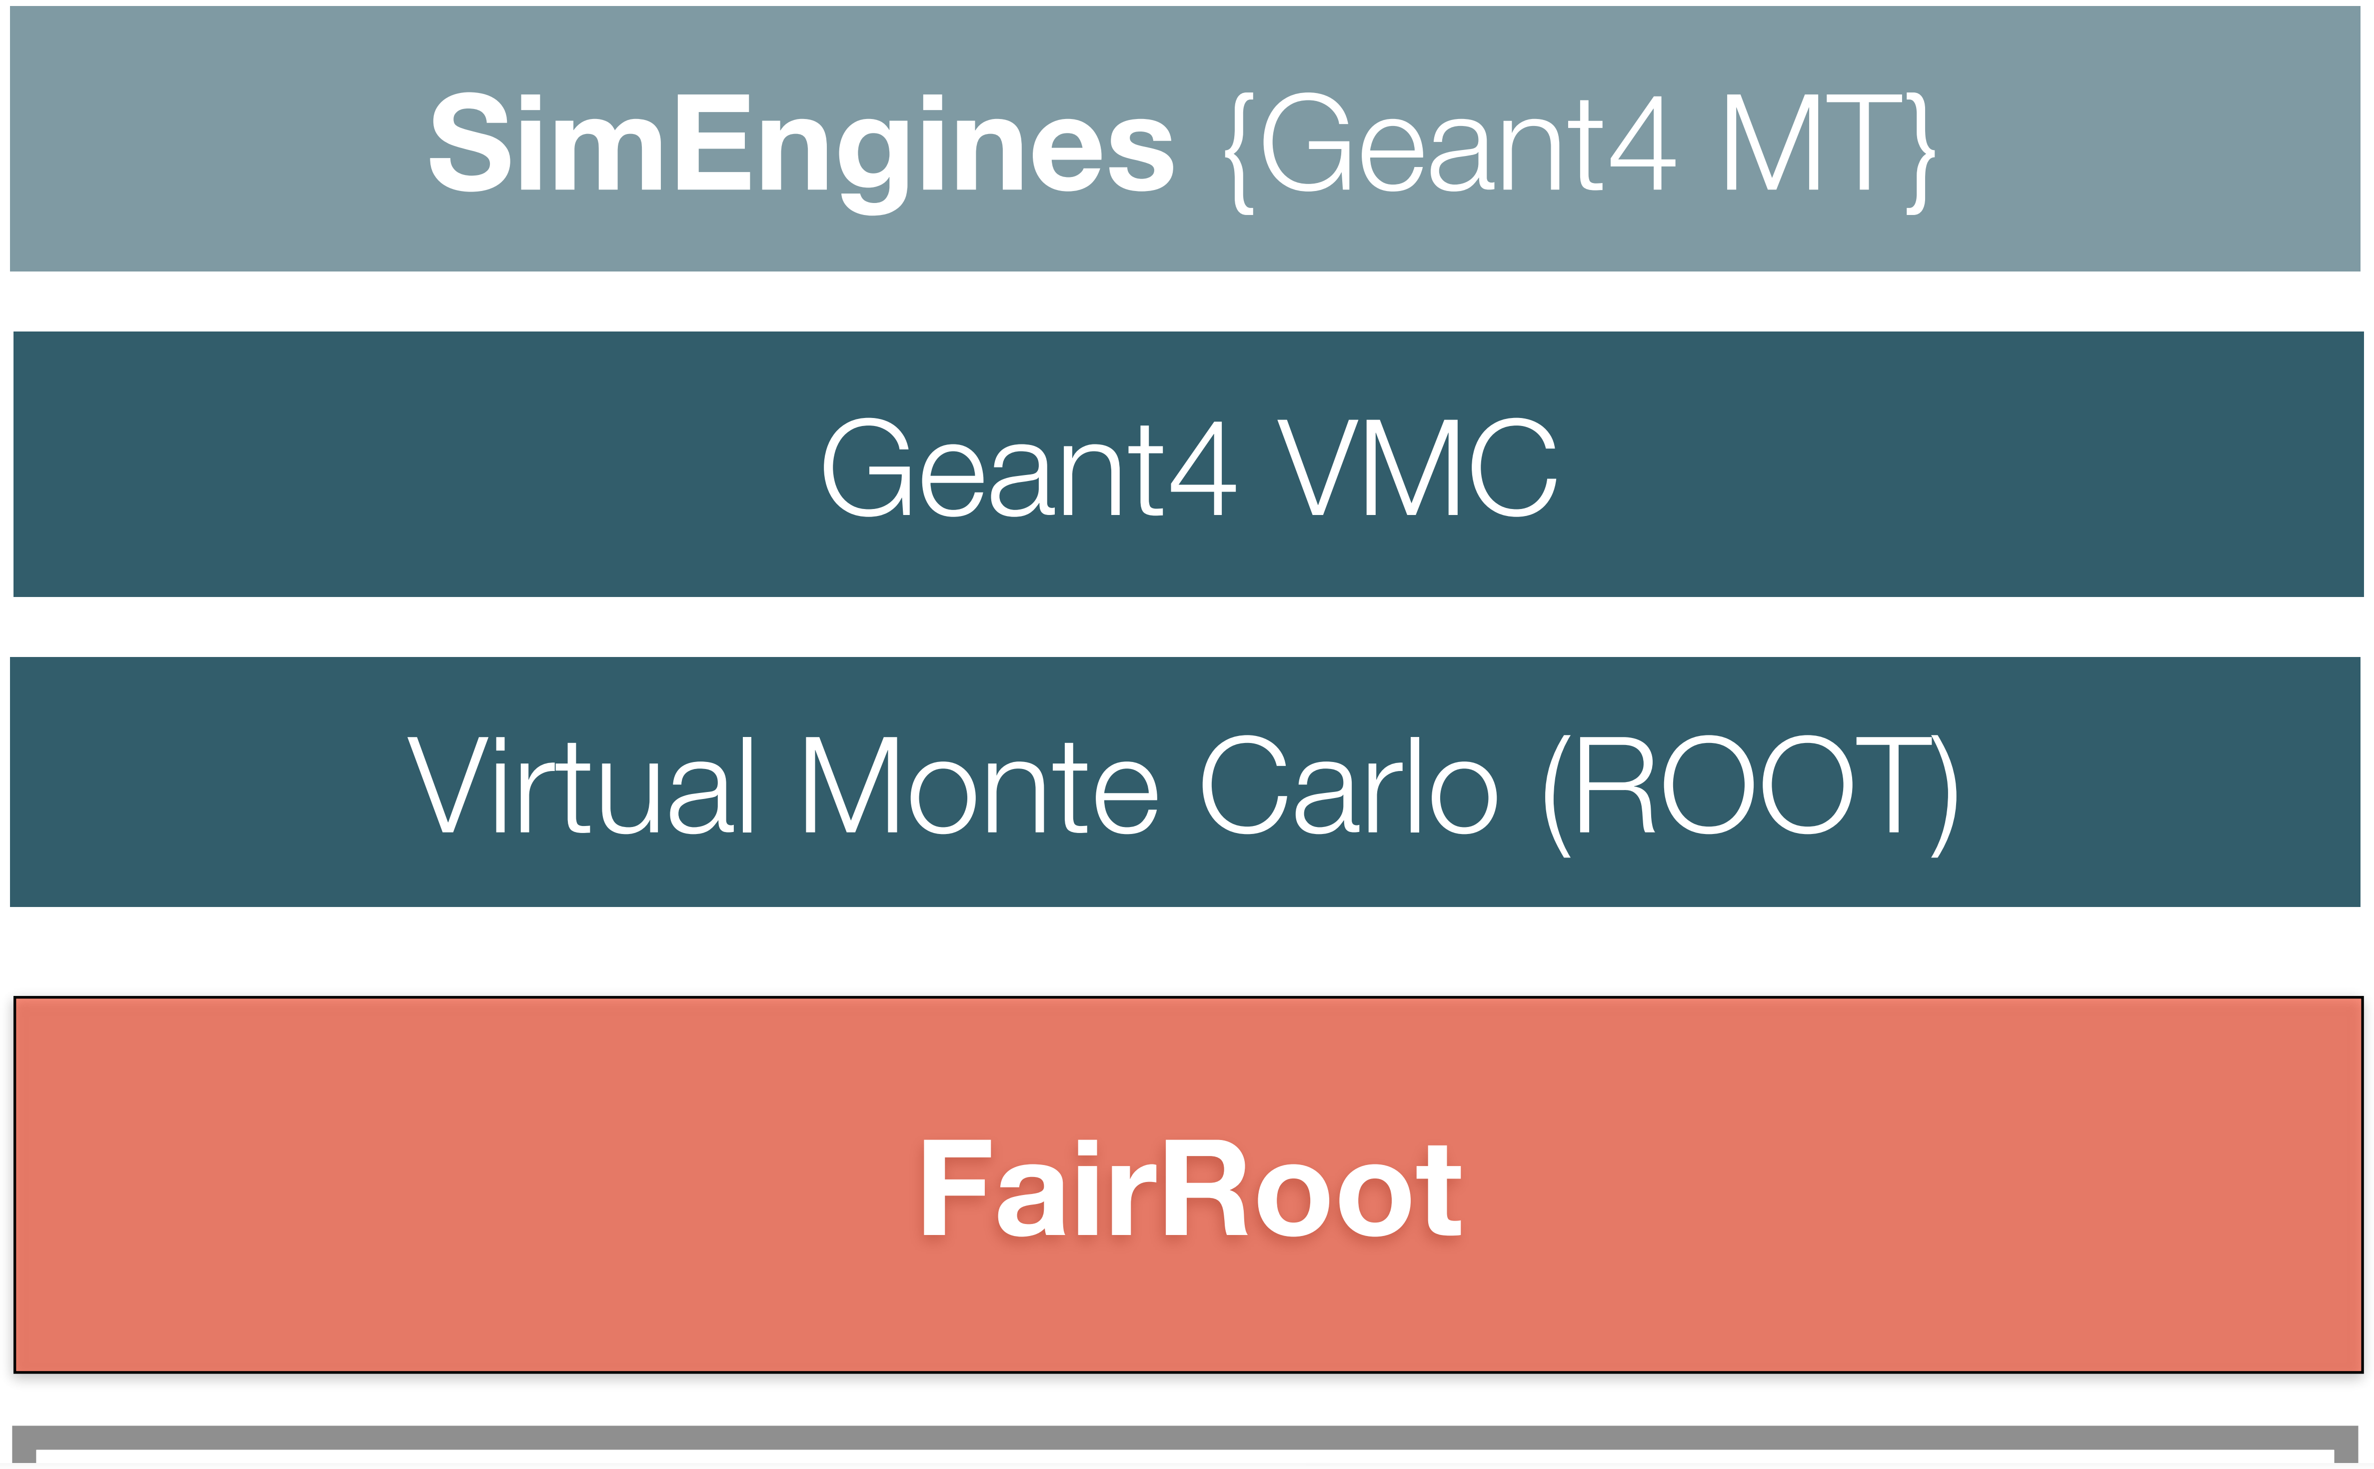

See more at the CHEP 2018 talk:
I.Hrivnacova: FairRoot and O2 Multi-threading Simulation

- Complementary to the multi-processing approach
- Event level parallelism
- Migration to MT concerned all packages:
  - VMC (ROOT), Geant4 VMC, FairRoot, O2
- The FairRoot simulation examples and the O2 simulation program were enhanced with the multi-threading run option
- The MT mode was added in the standard O2 tests

Conclusions

- The Run2 ALICE simulation
  - Geant4 validation is ongoing by ALICE DPG group
  - Tuning the configuration to get the best results without loosing performance
- The Run3 ALICE simulation aims to be HPC ready
  - Multi-processing parallelism based on FairMQ:
    - we sub-event parallelized Geant4, Geant3 and FLUKA at the same time
    - do event generation, transport and IO asynchronously
    - have elasticity and are agile (can change deployment easily and scale across nodes)
  - Multi-threading parallelism - with Geant4 MT