



LHCb workloads in the benchmarking suite

Andrea Valassi (IT-DI-LCG)

With contributions from and many thanks to:

Concezio Bozzi, Ben Couturier, Domenico Giordano, Servesh Muralidharan

Benchmarking Pre-GDB, 8 October 2019 – CERN

<https://indico.cern.ch/event/739897>

For more details, see previous talks:

Benchmarking WG meeting, 30th August 2019, <https://indico.cern.ch/event/782626>

System performance modelling WG meeting, 6th March 2019 <https://indico.cern.ch/event/772026>

System performance modelling WG meeting, 6th February 2019 <https://indico.cern.ch/event/772025>



Current LHCb workload: gen-sim

- **A typical LHCb “MC simulation job”**
 - Including both event generation and detector simulation
 - Single-threaded and single-process
 - This type of job represents ~90% of CPU consumed by LHCb in WLCG
- Extremely easy to run
 - It takes around 30 minutes to process 5 events
- Current metric: score is #events per (wall-time) 1000 seconds
 - In the past:
 - One score including both generation and simulation
 - Including the first event (slower due to I/O)
 - In the latest versions:
 - Two new separate score for generation and simulation, excluding first event
 - Old gen-sim score is still measured and recorded

LHCb gen-sim workload metrics

- In the past: parse the total event loop for generation and simulation

One of 32 job logs:

```
INFO Algorithm (millisec) | <user> | <clock> | min max | entries | total (s) |
INFO -----
INFO EVENT LOOP | 350681.855 | 351971.374 | 226829.970 587473.2 | 5 | 1759.857 |
```

In this job: average 352s (= 1760 / 5) per event – i.e. 2.84 events per 1000s

Benchmarking score: {"score": 91.1255, "avg": 2.8477, "median": 2.8431, "min": 2.8102, "max": 2.9122} {"copies":32, "threads_x_copy":1, "events_x_thread": 5, "throughput_score": {"s

Across 32 jobs: average 2.85 events per 1000s – i.e. 91.1 events per 1000s node throughput (= 2.85 x 32)

- Now: separately parse generation and simulation, skip the first event

Event #3 GEN

```
2019-02-01 10:17:12 UTC GaussGen INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 117214431, 0]
PYTHIA Warning in SpaceShower: pT2nextQCD: small daughter PDF
2019-02-01 10:19:53 UTC GaussGenPrevPrev INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 935330033, 0]
2019-02-01 10:19:53 UTC GaussGenPrev INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 1199519958, 0]
2019-02-01 10:19:53 UTC GaussGenNext INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 233623463, 0]
PYTHIA Error in BeamRemnants: setKinematics: kinematics construction failed
```

One of 32 job logs:

```
2019-02-01 10:19:53 UTC MainEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 1707412442, 0]
2019-02-01 10:20:35 UTC GiGaGetMainEvent INFO Number of extracted MCParticles 'MC/Particles' 1406
2019-02-01 10:20:35 UTC GiGaGetMainEvent INFO Number of extracted MCVertices 'MC/Vertices' 1510
2019-02-01 10:20:35 UTC PrevPrevEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 1184460771, 0]
2019-02-01 10:20:35 UTC PrevEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 549931932, 0]
2019-02-01 10:20:35 UTC NextEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 17447920, 0]
2019-02-01 10:20:59 UTC GiGaGetNextEvent INFO Number of extracted MCParticles 'MC/Particles' 742
2019-02-01 10:20:59 UTC GiGaGetNextEvent INFO Number of extracted MCVertices 'MC/Vertices' 895
2019-02-01 10:20:59 UTC NextNextEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 1401020779, 0]
2019-02-01 10:20:59 UTC GaussGen INFO Evt 8804, Run 7140089, Nr. in job = 4 with seeds [7140089, 8804, 195401829, 0]
```

Event #3 SIM

Pile-up FILE: /tmp/results/lhcb-n32-e5-1549015048_7694/proc_1/out_1.log

iCPU=xx	[time,	evt,	#part,	##	#vert,	##	G/S]	time/k#part
iCPU=00	[32.000,	8801,	0,	0,	0,	0,	GEN]	---
iCPU=00	[14.000,	8802,	0,	0,	0,	0,	GEN]	---
iCPU=00	[161.000,	8803,	0,	0,	0,	0,	GEN]	---
iCPU=00	[82.000,	8804,	0,	0,	0,	0,	GEN]	---
iCPU=00	[11.000,	8805,	0,	0,	0,	0,	GEN]	---
iCPU=00	[365.000,	8801,	9638,	5,	10387,	5,	SIM]	37.87093
iCPU=00	[573.000,	8802,	19742,	5,	21208,	5,	SIM]	29.02441
iCPU=00	[66.000,	8803,	2148,	2,	2405,	2,	SIM]	30.72626
iCPU=00	[151.000,	8804,	5443,	3,	5901,	3,	SIM]	27.74205
iCPU=00	[305.000,	8805,	12255,	5,	12933,	5,	SIM]	24.88780

Comments:

- GEN and SIM time ~uncorrelated
- SIM time ~proportional to #particles (but first event SIM is slower)
- SIM time / particle more stable as a throughput metric than time / event

Metrics for this job:

Event #1 SIM: includes slow initialization (e.g. magnetic field, conditions)

[FIRSTEVENT] Total time in magnetic field SIM initialization 45.0



Trident analysis

- Servesh ran the bmk on a node with 16 virtual processors (8 cores 2xHT)
 - 16 identical jobs executed in parallel, i.e. roughly ~ at the same time
 - Trident plot shows all 16 jobs
 - Correlated with timing info from a single log of one of these 16 jobs
- Summary of results:
 - *GEN slightly more efficient than SIM (visible as higher IPC regions)*
 - Initialization of GEN and SIM clearly visible, before and during 1st event
 - Should skip 1st event for benchmarking steady-state GEN and SIM
 - Some patterns to be better understood (memory peak in a GEN event)

[START OF JOB] Time 0s

ApplicationMgr INFO Application Manager Configured successfully
2019-02-28 16:14:13.471 UTC DetectorPersist... INFO Added successfully Conversion service:XmlCnvSvc
...

[SIGNAL GENERATOR INITIALIZATION] Time 10s

2019-02-28 16:14:23.337 UTC GaussGen SUCCESS Requested to process 5 events
2019-02-28 16:14:23.462 UTC GaussGen INFO Create BeamForInitialization
[photos splash]
[pythia splash]
[taula splash]

2019-02-28 16:14:41.023 UTC Generation.Sign... INFO Generating **Signal events of D*(2010)- D*(2010)+**

2019-02-28 16:14:24.379 UTC Generation.Sign...SUCCESS Using as production engine Pythia8Production
2019-02-28 16:14:41.024 UTC Generation SUCCESS Requested to generate EventType 27163076

[MINBIAS GENERATOR INITIALIZATION - PILEUP #1/4] Time 30s

2019-02-28 16:14:41.089 UTC GaussGenPrevPrev SUCCESS Requested to process 5 events
2019-02-28 16:14:41.089 UTC GenerationPrevPrev... INFO Generating **Minimum Bias events.**
2019-02-28 16:14:41.089 UTC GenerationPrevPrev...SUCCESS Using as production engine Pythia8Production
2019-02-28 16:14:57.589 UTC GenerationPrevPrevSUCCESS Requested to generate EventType 30000000

[MINBIAS GENERATOR INITIALIZATION - PILEUP #2/4] Time 45s

2019-02-28 16:14:57.591 UTC GaussGenPrev SUCCESS Requested to process 5 events
2019-02-28 16:14:57.591 UTC GenerationPrev.... INFO Generating **Minimum Bias events.**
2019-02-28 16:14:57.591 UTC GenerationPrev....SUCCESS Using as production engine Pythia8Production
2019-02-28 16:15:14.066 UTC GenerationPrev SUCCESS Requested to generate EventType 30000000

[MINBIAS GENERATOR INITIALIZATION - PILEUP #3/4] Time 60s

2019-02-28 16:15:14.067 UTC GaussGenNext SUCCESS Requested to process 5 events
2019-02-28 16:15:14.068 UTC GenerationNext.... INFO Generating **Minimum Bias events.**
2019-02-28 16:15:14.068 UTC GenerationNext....SUCCESS Using as production engine Pythia8Production
2019-02-28 16:15:30.561 UTC GenerationNext SUCCESS Requested to generate EventType 30000000

[MINBIAS GENERATOR INITIALIZATION - PILEUP #4/4] Time 75s

2019-02-28 16:15:30.563 UTC GaussGenNextNext SUCCESS Requested to process 5 events
2019-02-28 16:15:30.563 UTC GenerationNextN... INFO Generating **Minimum Bias events.**
2019-02-28 16:15:30.564 UTC GenerationNextN...SUCCESS Using as production engine Pythia8Production
2019-02-28 16:15:47.279 UTC GenerationNextNextSUCCESS Requested to generate EventType 30000000

[SIGNAL SIMULATION INITIALIZATION] Time 90s

2019-02-28 16:15:47.282 UTC MainEventGaussSim SUCCESS Requested to process 5 events

[MINBIAS SIMULATION INITIALIZATION - PILEUP #1/4] Time 115s

2019-02-28 16:16:09.422 UTC PrevPrevEventGa...SUCCESS Requested to process 5 events

[MINBIAS SIMULATION INITIALIZATION - PILEUP #2/4] Time 115s

2019-02-28 16:16:09.434 UTC PrevEventGaussSim SUCCESS Requested to process 5 events

[MINBIAS SIMULATION INITIALIZATION - PILEUP #3/4] Time 115s

2019-02-28 16:16:09.446 UTC NextEventGaussSim SUCCESS Requested to process 5 events

[MINBIAS SIMULATION INITIALIZATION - PILEUP #4/4] Time 115s

2019-02-28 16:16:09.457 UTC NextNextEventGa...SUCCESS Requested to process 5 events

[GENERATE EVENT #1/5 - SIGNAL AND MINBIAS PILEUP] Time 115s

2019-02-28 16:16:09.471 UTC ApplicationMgr INFO Application Manager Started successfully
2019-02-28 16:16:09.490 UTC GaussGen INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 516570798, 0]
PYTHIA Warning in SpaceShower::pT2nextQCD: weight above unity
PYTHIA Error in BeamRemnants::setKinematics: kinematics construction failed
2019-02-28 16:16:28.813 UTC GaussGenPrevPrev INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 1161403360, 0]
2019-02-28 16:16:28.816 UTC GaussGenPrev INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 281376746, 0]
2019-02-28 16:16:28.846 UTC GaussGenNext INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 209151265, 0]
2019-02-28 16:16:28.912 UTC GaussGenNextNext INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 1073365045, 0]

[SIMULATE EVENT #1/5 - INITIALIZATION OF MAGNETIC FIELD] Time 135s

--- Warning from G4Material::G4Material() define a material with density=0 is not allowed.
The material /dd/Materials/Vacuum will be constructed with the default minimal density: 1e-25g/cm3
2019-02-28 16:16:29.342 UTC MagneticFieldSvc INFO **Opened magnetic field file :**
/cvmfs/lhcb.cern.ch/lib/lhcb/DBASE/FieldMap/v5r7/cdf/field.v5r0.c1.down.cdf
...

[SIMULATE EVENT #1/5 - INITIALIZATION OF GEOMETRY AND PHYSICS] Time 150s

2019-02-28 16:16:42.507 UTC GiGaLVcwnv WARNING GiGaVoiConv Misalignment: Multiple indistinguishable misalignable detector elements found inside /dd/Geometry/AlterMagnetRegion/T/IT/ITV_410_tto4_ShortLayer - Please Check Geometry DB
2019-02-28 16:16:52.397 UTC MainEventGaussSim INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 1637558138, 0]
2019-02-28 16:16:52.400 UTC GiGa.GiGaMgr INFO GiGaRunManager:: Geometry will be extracted from GiGaGeo [LHCb detector info printouts]
[HADRONIC PROCESSES SUMMARY splash]
...

[SIMULATE EVENT #1/5 - SIMULATE SIGNAL AND MINBIAS PILEUP] Time 180s

2019-02-28 16:17:17.528 UTC RichG4HistoSet2 INFO Now Booking Rich G4 Histo Set2
...
2019-02-28 16:18:28.753 UTC GiGaGetMainEvent INFO Number of extracted MCParticles 'MC/Particles' 3991
2019-02-28 16:18:28.754 UTC GiGaGetMainEvent INFO Number of extracted MCVertices 'MC/Vertices' 4354
...

2019-02-28 16:19:56.208 UTC NextNextEventGa... INFO Evt 8801, Run 7140089, Nr. in job = 1 with seeds [7140089, 8801, 143414944, 0]
2019-02-28 16:19:56.230 UTC GiGaGetNextNext... INFO Number of extracted MCParticles 'MC/Particles' 2
2019-02-28 16:19:56.230 UTC GiGaGetNextNext... INFO Number of extracted MCVertices 'MC/Vertices' 3
2019-02-28 16:19:56.236 UTC IODataManager INFO Referring to dataset 00071400_000000089_1.sim by its file ID:F626C0AF-743B-E911-AC75-0242AC110002

[GENERATE EVENT #2/5 - SIGNAL AND MINBIAS PILEUP] Time 340s

2019-02-28 16:19:56.705 UTC GaussGen INFO Evt 8802, Run 7140089, Nr. in job = 2 with seeds [7140089, 8802, ...]
2019-02-28 16:20:05.663 UTC GaussGenNextNext INFO Evt 8802, Run 7140089, Nr. in job = 2 with seeds [7140089, 8802, 1456434655, 0]

[SIMULATE EVENT #2/5 - SIGNAL AND MINBIAS PILEUP] Time 350s

2019-02-28 16:20:05.667 UTC MainEventGaussSim INFO Evt 8802, Run 7140089, Nr. in job = 2 with seeds [7140089, 8802, 1360117217 ...]

2019-02-28 16:25:30.401 UTC NextNextEventGa... INFO Evt 8802, Run 7140089, Nr. in job = 2 with seeds [7140089, 8802, 618267676, 0]
2019-02-28 16:25:32.805 UTC GiGaGetNextNext... INFO Number of extracted MCParticles 'MC/Particles' 275
2019-02-28 16:25:32.805 UTC GiGaGetNextNext... INFO Number of extracted MCVertices 'MC/Vertices' 261

[GENERATE EVENT #3/5 - SIGNAL AND MINBIAS PILEUP] Time 680s

2019-02-28 16:25:33.255 UTC GaussGen INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 117214431, 0]
PYTHIA Warning in SpaceShower::pT2nextQCD: **small daughter PDF**

2019-02-28 16:27:14.261 UTC GaussGenPrevPrev INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 935330033, 0]

[SIMULATE EVENT #3/5 - SIGNAL AND MINBIAS PILEUP] Time 780s

2019-02-28 16:27:14.264 UTC MainEventGaussSim INFO Evt 8803, Run 7140089, Nr. in job = 3 with seeds [7140089, 8803, 2073811835, 0]
2019-02-28 16:27:41.809 UTC /dd/Structure/L... ERROR Gap not found!
2019-02-28 16:27:41.813 UTC GiGaGetMainEvent INFO Number of extracted MCParticles 'MC/Particles' 1419
2019-02-28 16:27:41.813 UTC GiGaGetMainEvent INFO Number of extracted MCVertices 'MC/Vertices' 1501

[GENERATE EVENT #4/5 - SIGNAL AND MINBIAS PILEUP] Time 820s

2019-02-28 16:27:55.174 UTC GaussGen INFO Evt 8804, Run 7140089, Nr. in job = 4 with seeds [7140089, 8804, 195401829, 0]
PYTHIA Warning in MultipartonInteractions::pTnext: weight above unity
PYTHIA Error in Pythia::check: charge not conserved
PYTHIA Abort from Pythia::forceHadronLevel: check of event revealed problems
PYTHIA Warning in TauDecays::decay: unknown tau production, assuming unpolarized and uncorrelated
2019-02-28 16:28:45.095 UTC GaussGenPrevPrev INFO Evt 8804, Run 7140089, Nr. in job = 4 with seeds [7140089, 8804, 1305292043, 0]

[SIMULATE EVENT #4/5 - SIGNAL AND MINBIAS PILEUP] Time 870s

2019-02-28 16:28:45.125 UTC MainEventGaussSim INFO Evt 8804, Run 7140089, Nr. in job = 4 with seeds [7140089, 8804, 1348109561, 0]

2019-02-28 16:30:00.626 UTC NextNextEventGa... INFO Evt 8804, Run 7140089, Nr. in job = 4 with seeds [7140089, 8804, 185454740, 0]

[GENERATE EVENT #5/5 - SIGNAL AND MINBIAS PILEUP] Time 945s

2019-02-28 16:30:00.755 UTC GaussGen INFO Evt 8805, Run 7140089, Nr. in job = 5 with seeds [7140089, 8805, 974091012, 0]

[SIMULATE EVENT #5/5 - SIGNAL AND MINBIAS PILEUP] Time 950s

2019-02-28 16:30:06.773 UTC MainEventGaussSim INFO Evt 8805, Run 7140089, Nr. in job = 5 with seeds [7140089, 8805, 379163024, 0]

2019-02-28 16:32:56.143 UTC GiGaGetNextNext... INFO Number of extracted MCParticles 'MC/Particles' 2242
2019-02-28 16:32:56.143 UTC GiGaGetNextNext... INFO Number of extracted MCVertices 'MC/Vertices' 2293

[JOB FINALIZATION AND END] Time 1120s

2019-02-28 16:32:56.411 UTC ApplicationMgr INFO Application Manager Stopped successfully
2019-02-28 16:32:56.411 UTC GaussGen SUCCESS 5 events processed

2019-02-28 16:32:57.569 UTC ApplicationMgr INFO Application Manager Finalized successfully
2019-02-28 16:32:57.570 UTC ApplicationMgr INFO Application Manager Terminated successfully





- Correlation of log file and Trident timeline was long and tedious
 - Suggestion for BMK: extract the timeline info during log parsing (done for benchmark calculation anyway) to make Trident correlation easier
 - Suggestion for Trident: allow user-defined time start (t_0) and time scale

Future LHCb workloads?

- Include multi-process or multi-threaded MC simulation?
 - WIP in LHCb on multi process simulation (GaussMP)
 - Plan is launch first productions within ~weeks on KNL HPC at CINECA
 - WIP in LHCb on multi-threaded simulation (Gaussino)
 - Expected to reach full production quality next year
 - Reminder: MC sim is and will remain ~90% of LHCb CPU in WLCG!
- Include reconstruction workloads from the HLT?
 - Most active area of WIP in LHCb in view of Run3
 - Both CPU and GPU implementations are being worked on
 - These workloads are relevant for procurement for the LHCb trigger farm