CMS benchmark candidates

David Lange
Andrea Sciabà

Pre-GDB, 8 October 2019

Reference workloads (1/2)

- Defined by CMS in February 2018, last updated in December to generate 2018-like events
  - A more recent version of CMSSW greatly increased the DIGI CPU efficiency
  - No subsequent updates were needed
- GENSIM
  - ttbar events, no input data, Pythia-only (no filters)
  - generation and simulation steps are concatenated event by event
  - The generation step uses a negligible fraction of CPU cycles

Reference workloads (2/2)

- DIGItisation, trigger and pileup simulation
  - Input data is the output of the GENSIM job
  - Premixed minimum bias sample needed
- RECOnstruction and analysis data creation
  - Input data is the output of the DIGI job
- All workloads require access to Frontier for the conditions data

Performance plots

GEN-SIM

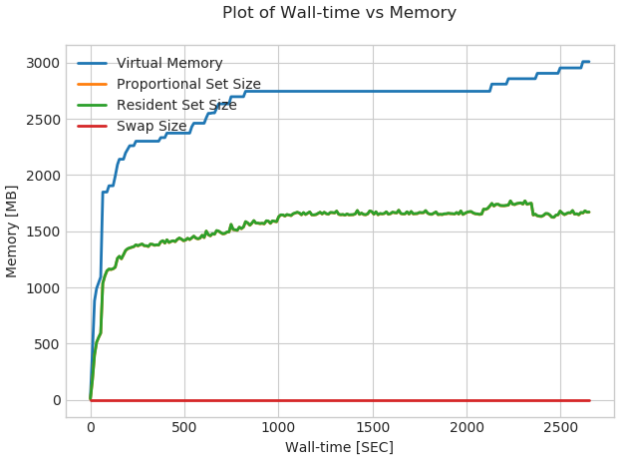

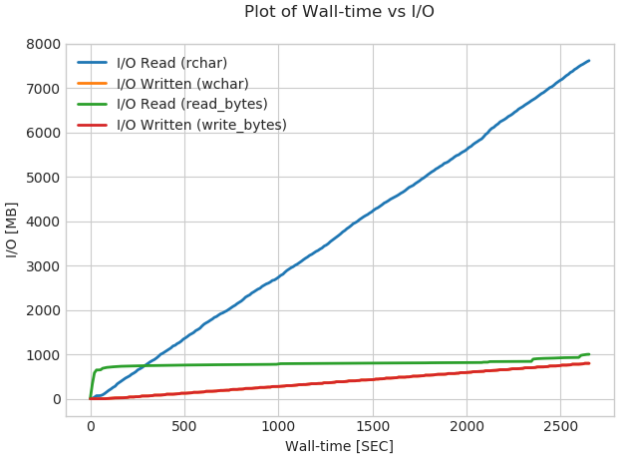

DIGI

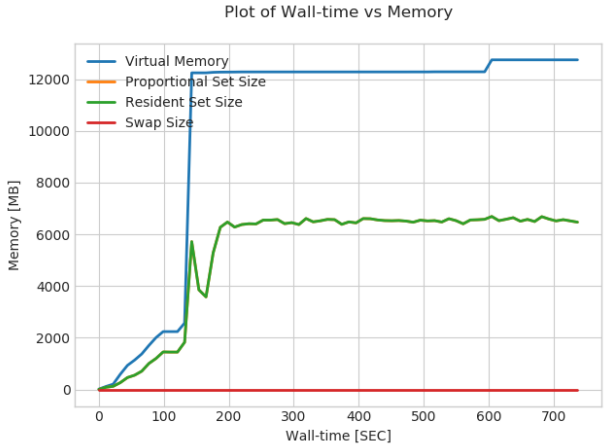

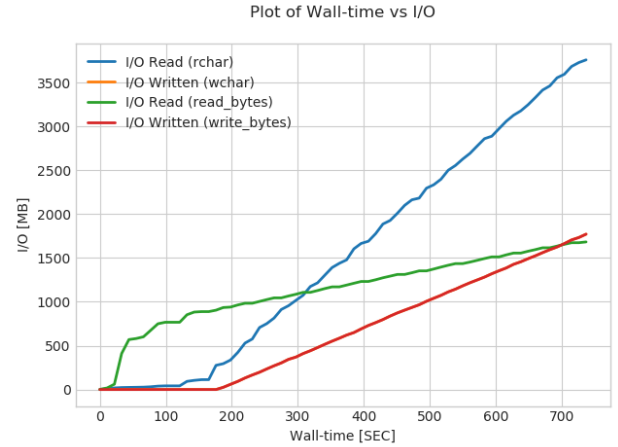

RECO

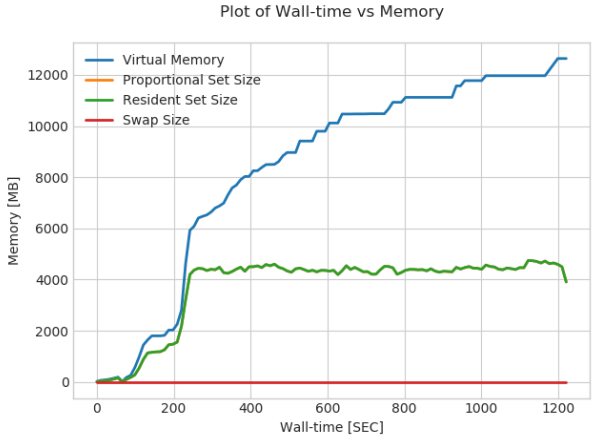

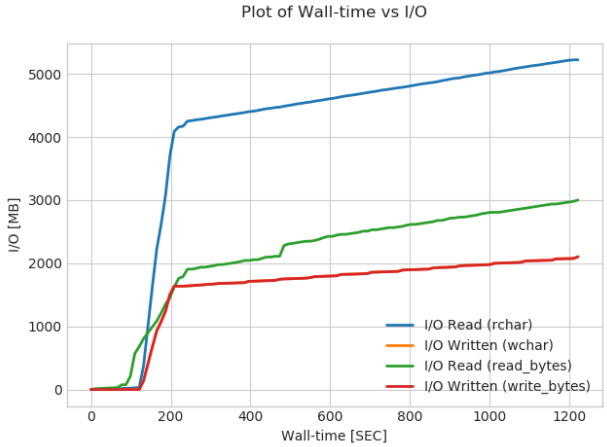

Typical performance

- Results obtained by running the job by hand with a no. of threads < no. of cores
  - Relatively high CPU efficiency
  - Modest memory requirements

Job	Events	Processes/threads	Wallclock time (sec)	CPU efficiency	Throughput per thread (events/sec)	Memory per core (GB)	Read rate per core (MB/sec)	Write rate per core (MB/sec)
GENSIM	1000	8	2651	99%	0.0472	0.19	0.0473	0.0377
DIGI	1000	8	737	78%	0.1696	0.65	0.2854	0.3004
RECO	1000	8	1221	83%	0.1024	0.45	0.3073	0.2153

Challenges (solved)

- First requirement is no need for network connectivity
  - Bundle the input files
    - For minbias one file is enough
  - Dump the needed conditions data into a file to be bundled
    - In theory a straightforward operation
    - In practice it involved an iterative procedure based on parsing error messages to find out what information had to be copied besides the Global Tag!

Score calculation

- Profiting from the TimeReport module in CMSSW to extract timing information
  - Same format for all CMS workloads, unified parser

```
TimeReport> Time report complete in 271.598 seconds
 Time Summary:
 - Min event:   3.34073
 - Max event:   26.2626
 - Avg event:   6.14064
 - Total loop:  156.391
 - Total init:  101.141
 - Total job:   271.598
 - EventSetup Lock:   90.3468
 - EventSetup Get:   35.7631
Event Throughput: 0.511539 ev/s
CPU Summary:
- Total loop:  490.798
- Total init:  76.2363
- Total extra: 0
- Total job:   582.288
Processing Summary:
- Number of Events:  80
```

```
{"copies":1 , "threads per copy":4 , "events per thread" : 20 , "wl-scores": {"digi": 0.5115}, "wl-stats": {"throughput_score": {"score": 0.5115, "avg": 0.5115, "median": 0.5115, "min": 0.5115, "max": 0.5115} , "CPU score": {"score": 0.1630, "avg": 0.1630, "median": 0.1630, "min": 0.1630, "max": 0.1630} }, "log": "ok", "app": {"version":"v0.15","description":"CMS DIGI of ttbar events based on CMSSW_10_2_9","cvmfs_checksum":"7bb1a8eb5227436a7625edeff6c1c89c","bmkdata_checksum":"b7dfab57bfbd132ebdd80335fe368e13","bmk_checksum":"b126af4c6938ccd96f583ea4492859df"} }
```

Feedback

1. Stable, relatively low spread
2. A few crashes observed in GEN-SIM for large numbers of events (> 100)
3. Score spread already small (3%) with just 25 events (now the default)

Conclusions

- CMS workloads found to be stable and not memory-hungry
- Needing a few GB of input data (between 1 and 4 GB depending on the workload)
- All known issues resolved