Comparison of Profiling Results for Run 3 and High Pileup LHC Simulation and Reconstruction

Mentee: Heather Thompson (Navajo Tech Univ.)
Mentor: Patrick Gartung (Fermilab)

Abstract

Abstract: The performance of CMS simulation and reconstruction software will be critical given the resource constraints on CPU and memory for the high luminosity LHC. Profiling the CPU and memory usage of the simulation and reconstruction software with every release is essential to ensure that performance remains stable or improves. Several profilers are available for profiling CMS software including Igprof and Intel Vtune. This project involves profiling with both profilers for each new release of CMSSW.

Overview

- Introduction
- Work to do
- Methods

Introduction

- A profiler is software that
  - Records snapshots of code performance on CPU.
  - Reports the sum of time spent in functions and their children
  - Reports the sum of memory allocated and used in functions and their children

- Profilers to be used
  - Igprof - fast and lightweight; handls loaded shared libraries, threads and sub-processes
  - Intel Vtune - Analysis and tuning tool that provides various examinations of performances.

Rank	% total	Counts: to / from this	Counts: Total	Paths: Including child / parent	Paths: Total	Symbol name
	85.22	1,072.75	1,074.28	5	5	edm::WorkerT<edm::stream::EDProducerAdaptorBase>::implDo(edm::Even
[16]	85.22	0.01	1,072.74	5	5	edm::stream::EDProducerAdaptorBase::doEvent(edm::EventTransitionIn
	11.64	146.47	146.47	1	1	cms::CkfTrackCandidateMakerBase::produceBase(edm::Event&, edm::Eve
	7.65	96.30	96.30	1	1	CAHitNtupletEDProducerT<CAHitQuadrupletGenerator>::produce(edm::Ev
	6.33	79.66	79.66	1	1	PrimaryVertexProducer::produce(edm::Event&, edm::EventSetup const&
	5.34	67.21	67.21	1	1	TrackProducer::produce(edm::Event&, edm::EventSetup const&)
	4.13	51.98	51.98	1	1	MuonIdProducer::produce(edm::Event&, edm::EventSetup const&)
	3.50	44.12	44.12	1	1	TrackListMerger::produce(edm::Event&, edm::EventSetup const&)

```
;;;;
;;;;  edm::WorkerT<edm::stream::EDProducerAdaptorBase>::implDo <cycle
[8];42.47;0.050986;3189.507993;edm::stream::EDProducerAdaptorBase::do
;;199.385318;33.207570;  TrackListMerger::produce;[46]
;;;;  PFBlockProducer::produce <cycle 15>;[958]
;;19.921059;50.005972;  produce;[198]
```

Work to do

- Connect to the CMSLPC cluster
- Learn how to run CMS software
- Learn how to run the profiler on CMS software
- Compare the text output of the profiler for each release
  - What are the top 5 functions for CPU usage
  - What are the top 5 functions for memory usage

Method

- Run scripts provided by mentor to gather profiles for each release and each profiler
- Make comparisons of top 5 functions by CPU usage and memory usage for each release and profiler
- Run script provided by mentor to gather profiles for Run 3 and high pileup LHC simulation and reconstruction software

What I did

Provided a script by my mentor, gathered and compared results from a Run 3 and HL-LHC workflow.

For the Run 3 workflow (11834.21)

- Copy the config file needed to run a reconstruction job for Run 3

```
cd ~/nobackup/CMSSW_12_5_X_2022-08-09-1100
cmsenv
cp /uscms_data/d1/gartung/CMSSW_12_5_0_pre2/step3-11834.21.py step3-11834.21.py
```

- Run vtune command to collect profile of reconstruction job

```
source /uscms/home/gartung/nobackup/intel/oneapi/setvars.sh
vtune -collect hotspots -r r11834.21 -resume-after=120 -data-limit=0 -knob
enable-stack-collection=true -knob  stack-size=4096 -knob sampling-mode=sw -- cmsRun
step3-11834.21.py 2>&1 | tee step3-11834.21.log
```

- Generate a Vtune hotspots report to get the top functions by CPU usage

```
vtune -report hotspots -r r11834.21 -format=csv -csv-delimiter=semicolon
>step3-11834.21.hotspots.csv
```

- Generate a Vtune gprof_cc report to get the callgraph of reconstruction

```
vtune -report gprof-cc -r r11834.21 -format=csv -csv-delimiter=semicolon
>step3-11834.21.gprof_cc.csv
```

For the HL-LHC workflow (35234.21)

- Copy the config file needed to run a reconstruction job for HL-LHC

```
cd ~/nobackup/CMSSW_12_5_X_2022-08-09-1100
cmsenv
cp /uscms_data/d1/gartung/CMSSW_12_5_0_pre2/step3-35234.21.py step3-35234.21.py
```

- Run vtune command to collect profile of reconstruction job

```
source /uscms/home/gartung/nobackup/intel/oneapi/setvars.sh
vtune -collect hotspots -r r35234.21 -resume-after=120 -data-limit=0 -knob
enable-stack-collection=true -knob  stack-size=4096 -knob sampling-mode=sw -- cmsRun
step3-35234.21.py 2>&1 | tee step3-35234.21.log
```

- Generate a Vtune hotspots report to get the top functions by CPU usage

```
vtune -report hotspots -r r35234.21 -format=csv -csv-delimiter=semicolon
>step3-35234.21.hotspots.csv
```

Integration build Run 3 workflow cpu hotspots

Summary Report Provides Overall Performance

- Lists instruction set(s) used
- Top time consuming functions

```
Top Hotspots
Function                                  Module                                  CPU Time  % of CPU Time(%)
----------------------------------------  --------------------------------------  --------  ----------------
[Outside any known module]                [Unknown]                                 9.026s              3.5%
CellularAutomaton::createAndConnectCells  libRecoPixelVertexingPixelTriplets.so     5.390s              2.1%
CellularAutomaton::findTriplets           libRecoPixelVertexingPixelTriplets.so     4.310s              1.7%
__memmove_avx_unaligned_erms              libc.so.6                                 3.399s              1.3%
magfieldparam::BCycl<float>::compute      libMagneticFieldParametrizedEngine.so     3.279s              1.3%
[Others]                                  N/A                                     230.122s             90.1%

Top Tasks
Task Type             Task Time  Task Count  Average Task Time
--------------------  ---------  ----------  -----------------
tbb_custom             288.120s          15            19.208s
tbb_parallel_for_each   24.466s         720             0.034s
tbb_parallel_for        12.188s       1,801             0.007s
Collection and Platform Info
    Application Command Line: cmsRun "step3-11834.21.py"
    Operating System: 3.10.0-1160.66.1.el7.x86_64 \S Kernel \r on an \m
    Computer Name: cmslpc139.fnal.gov
    Result Size: 72.3 MB
    Collection start time: 16:53:02 10/08/2022 UTC
    Collection stop time: 17:02:32 10/08/2022 UTC
    Collector Type: User-mode sampling and tracing
```

Sorted by self cost

Sort by cumulative cost

Rank	Total %	Self	Symbol name
89	2.31	112.75	CellularAutomaton::createAndConnectCells(std::vector<HitDoublets const*, std::allocator<HitDoublets const*> > c
102	2.03	99.01	CellularAutomaton::findTriplets(std::vector<HitDoublets const*, std::allocator<HitDoublets const*> > const&, st
99	1.74	84.78	VVIObjF::VVIObjF(float)
154	1.55	75.70	magfieldparam::BCycl<float>::compute(float, float, float&, float&) const
109	1.27	62.12	KFUpdator::update(TrajectoryStateOnSurface const&, TrackingRecHit const&) const

Integration build HL-LHC workflow cpu hotspots

```
Top Hotspots
Function                                        Module                                          CPU Time  % of CPU Time(%)
----------------------------------------------  ----------------------------------------------  --------  ----------------
CellularAutomaton::createAndConnectCells        libRecoPixelVertexingPixelTriplets.so            73.807s             7.1%
TrackListMerger::produce                        pluginRecoTrackerFinalTrackSelectorsPlugins.so   42.657s             4.1%
func@0x16c70                                    liblzma.so.5                                     32.866s             3.2%
PrimaryVertexAssignment::chargedHadronVertex    libCommonToolsRecoAlgos.so                       19.098s             1.8%
DAClusterizerInZT_vect::update                  libRecoVertexPrimaryVertexProducer.so            18.308s             1.8%
[Others]                                        N/A                                             853.907s            82.1%

Top Tasks
Task Type         Task Time  Task Count  Average Task Time
---------------  ---------  ----------  -----------------
tbb_custom        1093.257s          12            91.105s
tbb_parallel_for     5.488s          20             0.274s
Collection and Platform Info
    Application Command Line: cmsRun "step3-35234.21.py"
    Operating System: 3.10.0-1160.66.1.el7.x86_64 \S Kernel \r on an \m
    Computer Name: cmslpc139.fnal.gov
    Result Size: 84.5 MB
    Collection start time: 17:09:10 10/08/2022 UTC
    Collection stop time: 17:30:57 10/08/2022 UTC
```

Output depends on the collection type:
-User-Mode Sampling and Tracing Collection

-Hardware Event-based Sampling Collection.

Sort by cumulative cost

Rank	Total %	Self	Symbol name
45	6.24	356.74	CellularAutomaton::createAndConnectCells(std::vector<HitDoublets const*, std::a
72	3.58	204.66	TrackListMerger::produce(edm::Event&, edm::EventSetup const&)
127	1.89	107.90	DAClusterizerInZT_vect::update(double, DAClusterizerInZT_vect::track_t&, DAClu
120	1.83	104.38	PrimaryVertexAssignment::chargedHadronVertex(std::vector<reco::Vertex, std::al
109	1.59	90.61	VVIObjF::VVIObjF(float, float, int)

Pre-release build Run 3 workflow cpu hotspots

```
Top Hotspots
Function                                  Module                                  CPU Time  % of CPU Time(%)
----------------------------------------  --------------------------------------  --------  ----------------
[Outside any known module]                [Unknown]                                20.252s             4.8%
CellularAutomaton::createAndConnectCells  libRecoPixelVertexingPixelTriplets.so     7.990s             1.9%
CellularAutomaton::findTriplets           libRecoPixelVertexingPixelTriplets.so     6.400s             1.5%
magfieldparam::BCycl<float>::compute      libMagneticFieldParametrizedEngine.so     5.620s             1.3%
VVIObjF::VVIObjF                          libRecoLocalTrackerSiPixelRecHits.so      5.421s             1.3%
[Others]                                  N/A                                     374.413s            89.1%

Top Tasks
Task Type              Task Time  Task Count  Average Task Time
---------------------  ---------  ----------  -----------------
tbb_custom              386.611s          14            27.615s
tbb_parallel_for_each    39.623s         720             0.055s
tbb_parallel_for         19.745s       1,801             0.011s
Collection and Platform Info
    Application Command Line: cmsRun "step3-11834.21.py"
```

113	1.99	6.73	CellularAutomaton::createAndConnectCells(std::vector<HitDoublets const*, std::allo
141	1.72	5.84	CellularAutomaton::findTriplets(std::vector<HitDoublets const*, std::allocator<Hit
180	1.60	5.41	do_lookup_x
139	1.49	5.05	VVIObjF::VVIObjF(float)
208	1.38	4.68	magfieldparam::BCycl<float>::compute(float, float, float&, float&) const

Rank	% total	Counts to / from this	Counts Total	Paths Including child / parent	Paths Total	Symbol name
	2.46	8.34	9.80	1	1	CAHitQuadrupletGenerator::hitNtuplets(IntermediateHitDoublets const&, std::vec
[113]	**2.46**	**6.73**	**1.60**	**1**	**1**	**CellularAutomaton::createAndConnectCells(std::vector<HitDoublets const*, std::**
	0.35	1.20	2.19	1	2	CACell::checkAlignmentAndAct(std::vector<CACell, std::allocator<CACell> >&, st

Pre-release build HL-LHC cpu hotspots

```
Top Hotspots
Function                                        Module                                                CPU Time  % of CPU Time(%)
----------------------------------------------  ----------------------------------------------------  --------  ----------------
CellularAutomaton::createAndConnectCells        libRecoPixelVertexingPixelTriplets.so                  119.409s             7.1%
TrackListMerger::produce                        pluginRecoTrackerFinalTrackSelectorsPlugins.so          58.607s             3.5%
func@0x16c70                                    liblzma.so.5                                            49.649s             2.9%
DAClusterizerInZT_vect::update                  libRecoVertexPrimaryVertexProducer.so                   35.141s             2.1%
PrimaryVertexAssignment::chargedHadronVertex    libCommonToolsRecoAlgos.so                              32.971s             2.0%
[Others]                                        N/A                                                   1390.652s            82.5%

Top Tasks
Task Type         Task Time  Task Count  Average Task Time
----------------  ---------  ----------  -----------------
tbb_custom        1687.470s          14           120.534s
tbb_parallel_for     8.246s          20             0.412s
Collection and Platform Info
    Application Command Line: cmsRun "step3-35234.21.py"
    Operating System: 3.10.0-1160.66.1.el7.x86_64 \S Kernel \r on an \m
    Computer Name: cmslpc178.fnal.gov
    Result Size: 123.0 MB
    Collection start time: 20:08:19 08/08/2022 UTC
    Collection stop time: 20:41:07 08/08/2022 UTC
```

May take several minutes to run all necessary checks

```
    8 AsciiOut              -s AsciiOutputModul                                20        20
    9 fileAction            -s file_close                                       1         1
   10 fileAction            -s file_open                                        2         2

 type    category    Examples: run/evt        run/evt          run/evt
 ---- ------------------- ---------------- ---------------- ----------------
    1 DAClusterizerInZT_vect 1/3
    2 DAClusterizerInZT_vect 1/3              1/3              1/8
    3 MakeMuonWithTEV       1/5
    4 PFRecoTauChargedHadronFromGenericTrackPlugin::operator() 1/1                1/18
    5 TSCPBuilderNoMaterial 1/3              1/3              1/15
    6 TooManyPFDVCandidates 1/3              1/12             1/20
    7 TooManyClusters       1/11             1/6              1/19
    8 AsciiOut              1/11             1/6              1/19
    9 fileAction            End Run: 1
   10 fileAction            pre-events       pre-events

Severity    # Occurrences   Total Occurrences
--------    -------------   -----------------
Warning                45                  45
Error                  20                  20
System                 23                  23
```

What I learned

First off, I've learned that there are many tools that are available for profiling that target specific areas of performance. Working with Intel VTune I seen that it points out the problem for you, it takes out the guessing.