



Preliminary Profiling Results of LArSoft/ProtoDUNE

Soon Yung Jun (SCS/PDS)
LArSoft Coordination Meeting
Dec 5, 2017

Outline

- Introduction
 - Scope
 - Workflow
 - Tools
- Preliminary results
 - Bench tests
 - Memory profiling (IgProf)
 - CPU profiling (HPCToolkit)
- Summary
 - Remarks
 - Future work

Introduction: Scope

- Profiling LArSoft code: Profiling and optimization work to identify problem areas and solutions in existing workflows.
 - Setup an application to run LArSoft
 - Identify a major production workflow
 - Do preliminary profiling
 - Ease of use of profiling tools
- Resources: the Wilson cluster (tev.fnal.gov)
 - phi[1-4] nodes (4 x Intel Xeon CPU E5-2620@ 2.00GHz)
 - [/cvmfs/fermilab.opensciencegrid.org](http://cvmfs/fermilab.opensciencegrid.org)
 - [/cvmfs/dune.opensciencegrid.org](http://cvmfs/dune.opensciencegrid.org)
 - pnfs-stken:/larsoft
 - stkensrv1n:/dune
 - filesrv01.fnal.gov:/web/sites/g4cpt.fnal.gov

Introduction: Workflow

- LArSoft/protoDUNE: dunetpc v06_57_00
 - 6 GeV proton
 - gen-g4-detsim-reco-ana

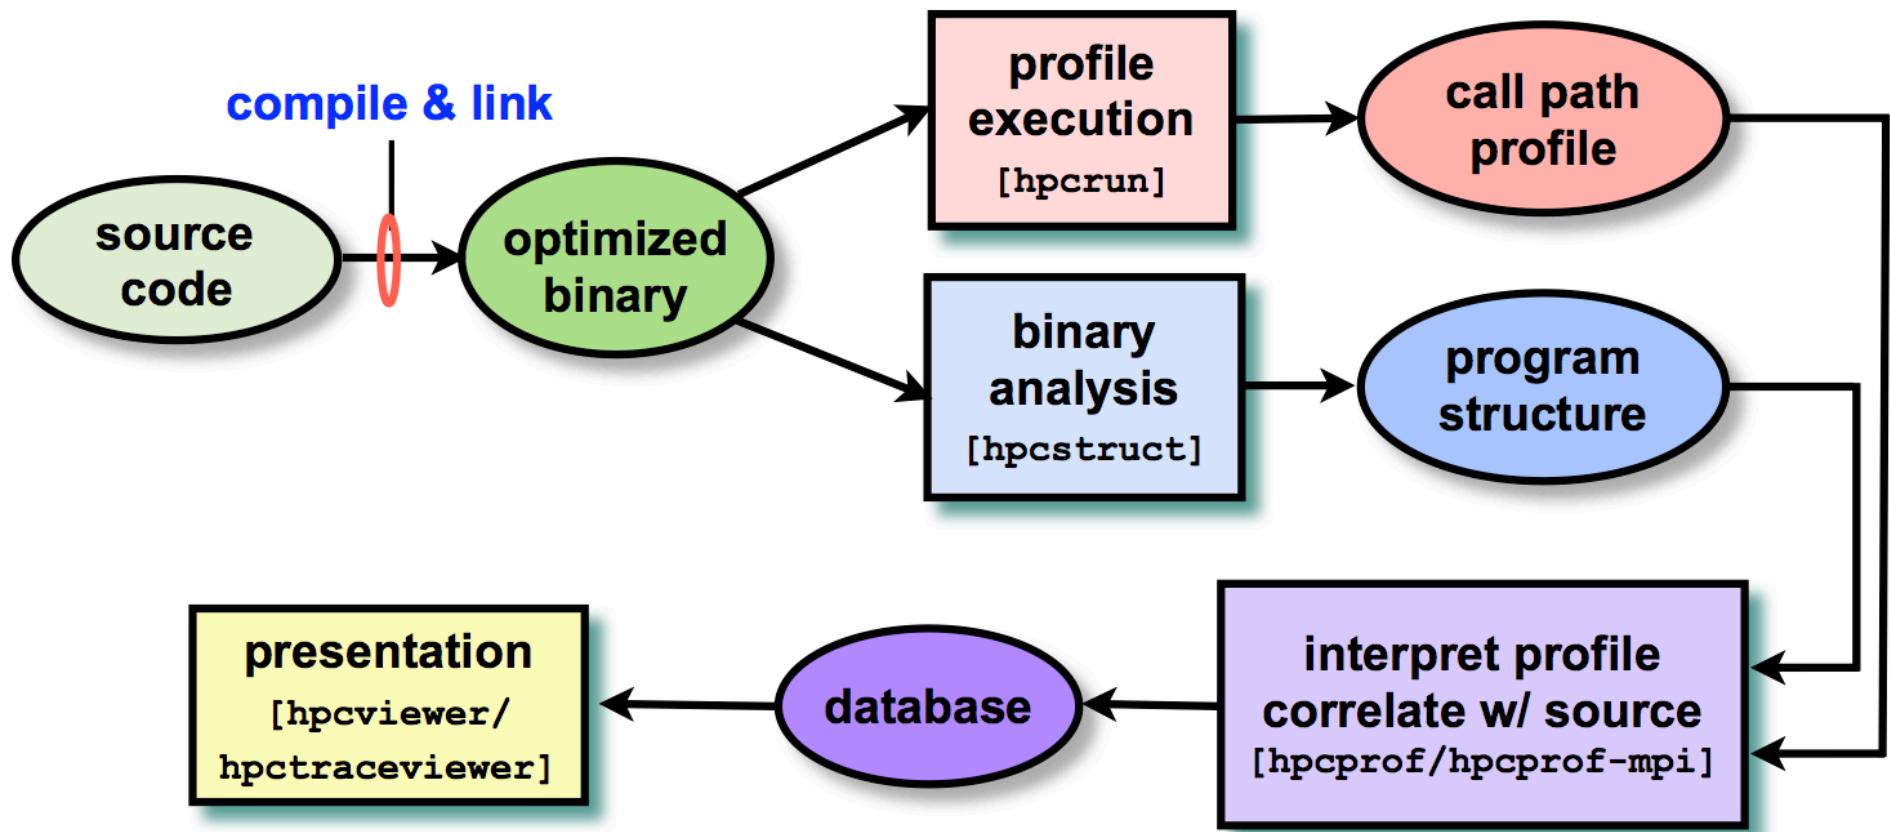
```
source /cvmfs/dune.opensciencegrid.org/products/dune/setup_dune.sh
setup dunetpc v06_57_00 -f Linux64bit+2.6-2.12 -q e14:prof

lar -n 100 -c gen_protoDune_proton_6p0GeV_mono.fcl
lar -c protoDune_g4_3ms.fcl gen_protoDune_proton_6p0GeV_mono.root
lar -c protoDUNE_detsim.fcl gen_protoDune_proton_6p0GeV_mono_g4.root
lar -c protoDUNE_reco.fcl gen_protoDune_proton_6p0GeV_mono_g4_detsim.root
lar -c protoDUNE_ana.fcl gen_protoDune_proton_6p0GeV_mono_g4_detsim_reco.root
```

- 1 GeV cosmic (protoDUNE mmc9)
 - lar -n 100 -c mcc9_gen_protoDune_beam_cosmics_p1GeV.fcl
 - Input files from /pnfs/dune and /pnfs/larsoft
- Thanks to Thomas Junk, Tingjun Yang, Robert Sulej, Dorota Stefan, Heidi Schellman

Introduction: Tools

- Memory: IgProf (<http://igprof.org>) v5.9.6
- CPU: HPC Toolkit (<http://hpctoolkit.org>)



Results: LArSoft

- Link: <http://g4cpt.fnal.gov/LArSoft.html>

LArSoft Profiling and Benchmarking (TBC)

1) Profiling Results

Profiled on [the Wilson cluster](#) using Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00GHz

Application	Version	Performance		
dunetpc	v06_57_00	BenchTest	IgProf	HPCToolkit
dunetpc	v06_54_00	BenchTest	IgProf	HPCToolkit

2) Links for Performance Tools

1. Performance Tools:

[IgProf](#)

[HPCToolkit](#)

[OpenSpeedshop](#)

Results: Bench Test

- CPU and Memory Summary from Art

Art TimeTracker/MemoryTracker

LArSoft/protoDune dunetpc v06_57_00

CPU/Event [sec]

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	2.7e-04	3.7 ± 1.3	5.7 ± 0.4	20.3 ± 5.2	2.1 ± 0.4
cosmic_1p0GeV	6.2 ± 0.6	154.2 ± 24.1	16.2 ± 0.9	676.3 ± 210.5	176.8 ± 33.2

Processor: Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00GHz, cache size : 15360 KB, MemTotal: 32845168 kB

RSS/VSIZE [MB]

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	822 / 256	1433 / 668	1083 / 479	3639 / 1205	1436 / 656
cosmic_1p0GeV	941 / 346	4201 / 3430	2314 / 1709	6899 / 4476	N/A

Results: Profiling Overhead

- Profiling overhead: Profiling/Bench-test (OH)

Profiling OverHead (proton_6p0GeV 100 Events)

1) Overhead CPU

	Default	HPCToolkit	(OH)	IGPROF	(OH)
gen	0.97	1.14	1.18	1.68	1.73
g4	747.76	768.84	1.03	1213.39	1.62
detsim	1132.56	1142.28	1.01	1186.36	1.05
reco	7033.85	7019.12	1.00	21614.96	3.07
ana	21.87	21.55	0.99	816.79	37.34

2) Overhead RealTime

	Default	HPCToolkit	(OH)	IGPROF	(OH)
gen	1.39	2.63	1.89	2.26	1.62
g4	751.84	778.35	1.04	1217.95	1.62
detsim	1134.81	1150.38	1.01	1188.92	1.05
reco	4080.21	4058.92	0.99	9320.36	2.28
ana	22.53	27.70	1.23	822.66	36.51

LArSoft/protoDUNE: IgProf

- https://g4cpt.fnal.gov/larsoft/dunetpc_v06_57_00/igprof.html

IgProf Profiling Information

dunetpc v06_57_00

Memory Profiling Reports

- MEM_TOTAL: the total amount of memory allocated by any function - a snapshot of poor memory locality
- MEM_MAX: the largest single allocation by any function
- MEM_LIVE: memory that has not been freed - snapshot of the heap, i.e. a heap profile

LArSoft/protoDune

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	TOTAL MAX LIVE				
cosmic_1p0GeV	TOTAL MAX LIVE				

- Summary of top memory contenders
 - g4: larg4::LArVoxelReadout::DriftIonizationElectrons(...)
 - detsim: opdet::OpDetDigitizerDUNE::produce(art::Event&)
 - reco geo::GeometryCore::ChannelToWire(unsigned int)
 - ana: sim::SimChannel::TrackIDsAndEnergies

LArSoft/protoDune: HPCToolkit

- https://g4cpt.fnal.gov/larsoft/dunetpc_v06_57_00/hpctoolkit.html

HPCToolkit Profiling Information

LArSoft/protoDune dunetpc v06_57_00

Performance Profiling Reports

- CPU: Inclusive and Exclusive Time (sampling with 200Hz frequency)
- CPI: Cycle per Instruction (PAPI_TOT_CYC/PAPI_TOT_INS)
- FMO: FLOPS/Memory Operations (Computational Intensity)

CPU

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	FUN LIB				
cosmic_1p0GeV	FUN LIB				

CPI/FMO

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	CPI FMO				

LArSoft/protoDUNE: HPCToolkit

- https://g4cpt.fnal.gov/larsoft/dunetpc_v06_57_00/hpctoolkit.html
- Top CPU functions
 - g4: G4Exp (from larg4::OpFastScintillation::PostStepDolt)
 - detsim: __normal_iterator (from sim::SimChannel::Charge)
 - detsim: __GI_memcpy (from art::EDProducer::doEvent)
 - reco: mostly from Eigen_tf namespace (tensorflow-core)
- CPI (Cycle per Instruction)
 - Good balance with minimal stalls
 - detsim/reco = 0.48/0.55
- FMO (Flops per Memory operation)
 - Computational intensity
 - detsim/reco = 0.35/0.24
- Other (PAPI) hardware counters can be easily added

HPC Toolkit: hpcviewer (Demo)

- https://g4cpt.fnal.gov/larsoft/dunetpc_v06_57_00/hpctoolkit.html

Database (for a full GUI analysis)

Sample	gen	g4	detsim	recon	anal
proton_6p0GeV	CPU CPI FMO				
cosmic_1p0GeV	CPU	CPU	CPU	CPU	CPU

Example steps:

- 1) click a link in the table and download the database file
- 2) tar xzf lar-xyz.dat.tgz
- 3) [hpcviewer](#) lar-xyz.dat

The screenshot shows the hpcviewer application window. At the top, there's a menu bar with File, Filter, View, Window, Help. Below the menu is a toolbar with icons for opening files, saving, filtering, and other functions. The main area has two panes. The left pane is a code editor displaying C++ code from SimWireDUNE_module.cc. The right pane is a performance analysis table.

Scope	CPUTIME (usec).[0,0] (I)	CPUTIME (usec).[0,0] (E)
Experiment Aggregate Metrics	1.78e+09 100 %	1.78e+09 100 %
__gnu_cxx::__normal_iterator<std::pair<unsigned short, std::vector<sim::IDE, std::allocator<	1.28e+08 7.2%	1.28e+08 7.2%
inflate_fast	1.13e+08 6.4%	1.13e+08 6.4%
Legacy35tZeroSuppressService::filter(std::vector<short, std::allocator<short> > const&, uns	7.67e+07 4.3%	6.68e+07 3.8%
int TStreamerInfo::ReadBuffer<TVirtualCollectionProxy>(TBuffer&, TVirtualCollectionProxy co	2.13e+08 12.0%	6.29e+07 3.5%
__read_nocancel	6.25e+07 3.5%	6.25e+07 3.5%
IdealAdcSimulator::count(double, unsigned int, unsigned int) const	5.79e+07 3.3%	5.79e+07 3.3%
hc2cf_32	5.14e+07 2.9%	5.14e+07 2.9%
__GI_memcpy	5.07e+07 2.8%	5.07e+07 2.8%
void util::LArFFT::DoInvFFT<float>(std::vector<TComplex, std::allocator<TComplex> >&, std:	1.95e+08 10.9%	4.38e+07 2.5%
r2cb_32	4.23e+07 2.4%	4.23e+07 2.4%
TComplex::operator=(TComplex const&)	4.19e+07 2.4%	4.19e+07 2.4%
ExponentialChannelNoiseService::addNoise(unsigned int, std::vector<float, std::allocator<float>	4.14e+07 2.3%	4.02e+07 2.3%
void util::LArFFT::Convolute<float>(std::vector<float, std::allocator<float> >&, std::vector<TCo	3.73e+07 2.1%	3.73e+07 2.1%
__gnu_cxx::__normal_iterator<std::pair<unsigned short, std::vector<sim::IDE, std::allocator<	3.73e+07 2.1%	3.73e+07 2.1%
SimWireDUNE::produce(art::Event&)	1.27e+09 71.1%	3.61e+07 2.0%
__ieee754_log	3.43e+07 1.9%	3.43e+07 1.9%

Summary

- Identified a major workflow of the LArSoft/protoDUNE simulation and reconstruction chain and provided preliminary profiling results
 - proton 6 GeV and cosmic 1 GeV (mmc9)
 - Total memory churns with IgProf
 - Top CPU functions and libraries with HPCToolkit
- Future work
 - Extend measurements for DUNE-FD simulation (?)
 - Try to resolve an issue converting raw data of OpenSpeedshop to database (conflict with forked processes other than lar - 9 sh/bash, 1 awk, 1 cat)
 - Understand instability of measurements (occasional deadlocks from reco)