HEP Event Reconstruction with Cutting Edge Computing Architectures

Giuseppe Cerati, Mike Wang, Allison Hall, Sophie Berkman (FNAL)
Boyana Norris, Brian Gravelle (University of Oregon)

SciDAC4-HEP Presentation - Sep. 27, 2018

Introduction

- Pilot project (3 years) with Fermilab physicists and UOregon computer scientists

- P5 panel recommendation: Higgs as a tool for discovery, physics of ν mass
  - HL-LHC, LArTPC program (SBN, DUNE)
- Reconstruction for current LHC and LArTPC experiments is challenging
- Future experiments even more challenging:
  - larger sizes or more granular: more detector channels
  - higher beam intensities: more data to process
- Reconstruction compute time does not scale well

- Need large speedups in reconstruction to reach design detector sensitivity and enable discoveries!

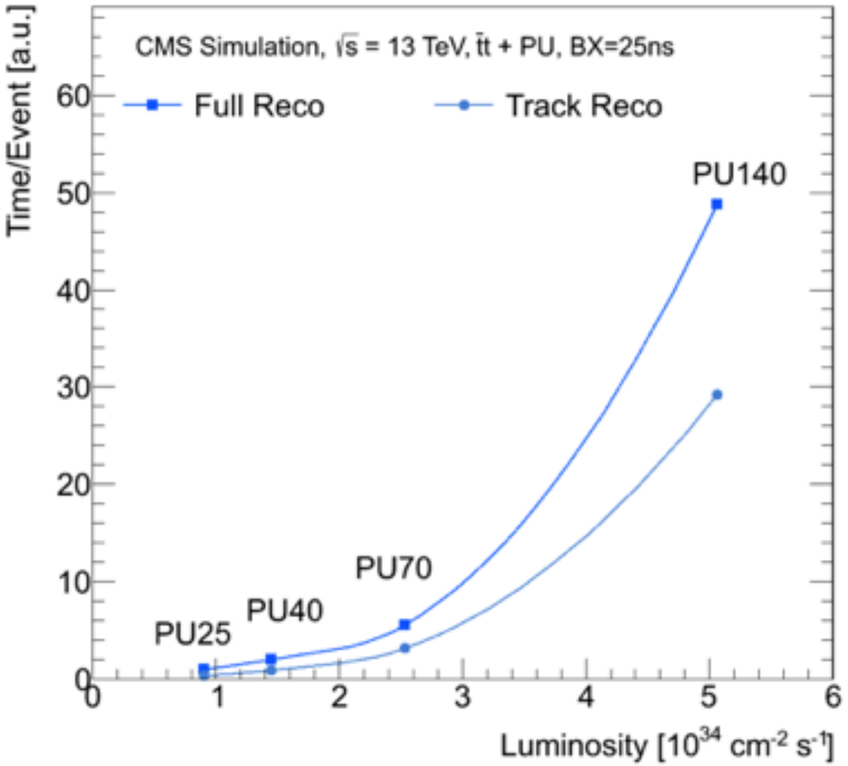

Project goals

- Accelerate HEP event reconstruction using modern parallel architectures
- Objectives of the project are the following:
  - Identify key algorithms which are top time contributions in the reconstruction workflows
  - Re-design the algorithm to make efficient usage of parallelism, both at data- and instruction-level
  - Deploy the new code in the experiments' framework
  - Explore execution on different architectures and platforms

- Focus on two areas:
  1. Novel parallel algorithm for charged particle tracking in CMS
     - expand UCSD/Princeton/Cornell collaboration, work on algorithm optimization (physics+computing) and deployment
  2. Pioneer similar techniques for reconstruction in LArTPC detectors
     - brand new effort, build on top of knowledge acquired through the CMS tracking effort

CMS tracking prototype

- Collaboration with UCSD/Cornell/Princeton
- Tracking is the reconstruction of charged particle trajectories through thin silicon layers
- Faster tracking enables higher reconstruction accuracy and throughput of online and offline data
- Single-thread 10x faster than CMSSW, additional 60x from 100+ threads on KNL
  - first tracking iteration
- First physics performance results are comparable with CMSSW and can be further improved
- High event throughput on Skylake is promising towards usage in High Level Trigger for LHC Run3

CMS data with PU~50, HL-LHC will have PU~200

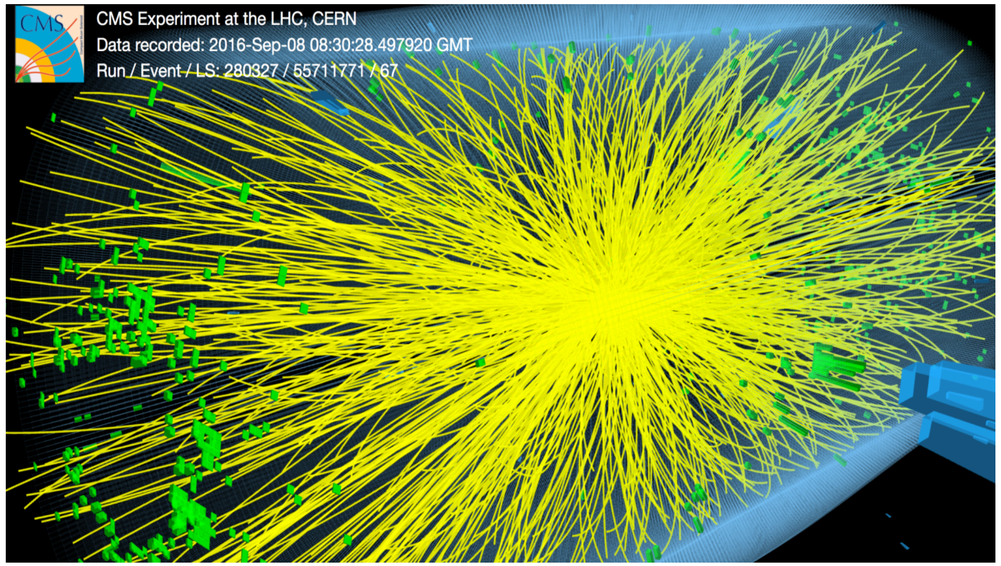

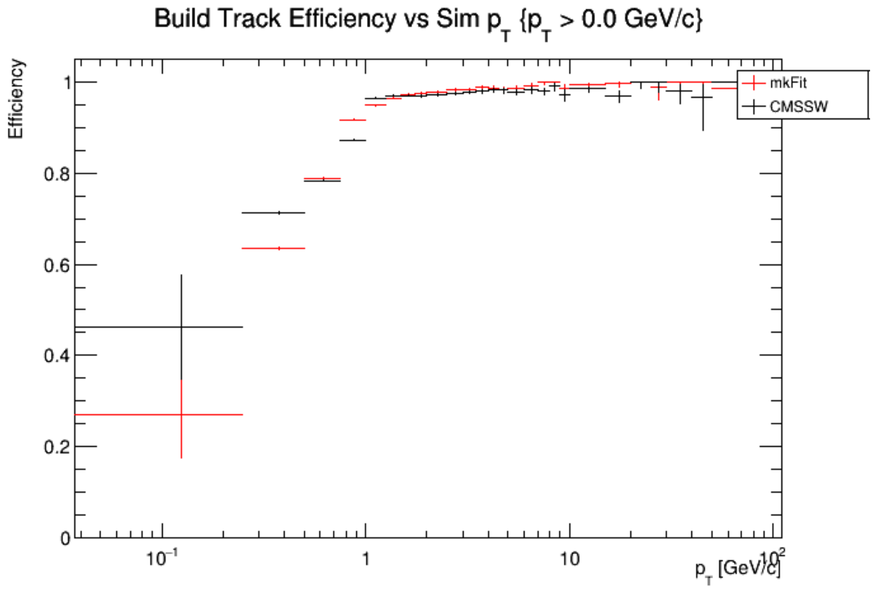

Fermilab

Profiling with TAU Commander

- Profiling the code with advanced tools gives a better understanding of scalability problems and bottlenecks
- Developed tools to analyze TAU Commander (RAPIDS Institute) profiling data with Pandas data frames
  - scaling performance, memory efficiency, correlation between metrics, per-function analysis including breakdown of stalls
- Test various platforms, incl. those used in CMS HLT
- First analysis led to speedups of ~8%, more detailed analysis ongoing will provide indications for further improvements

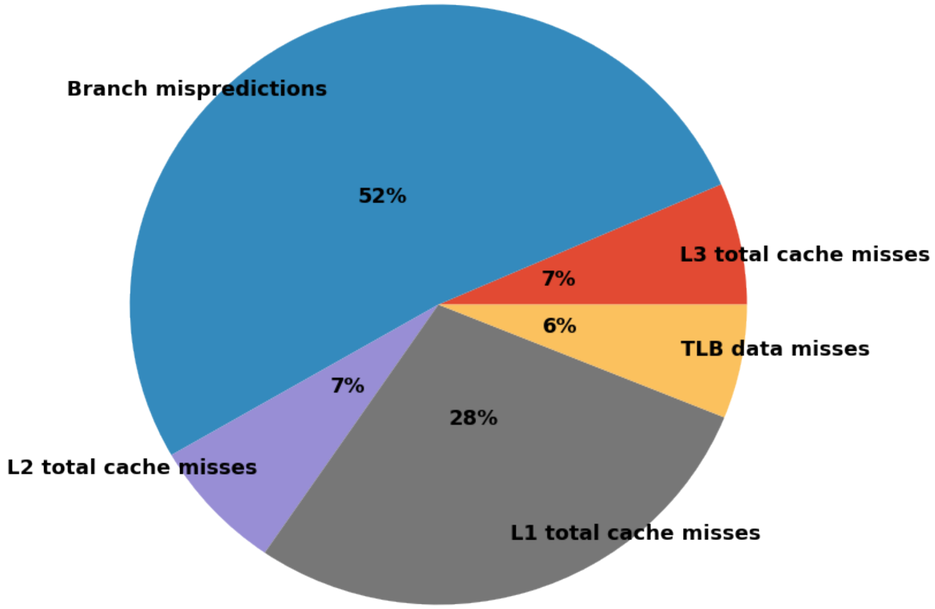

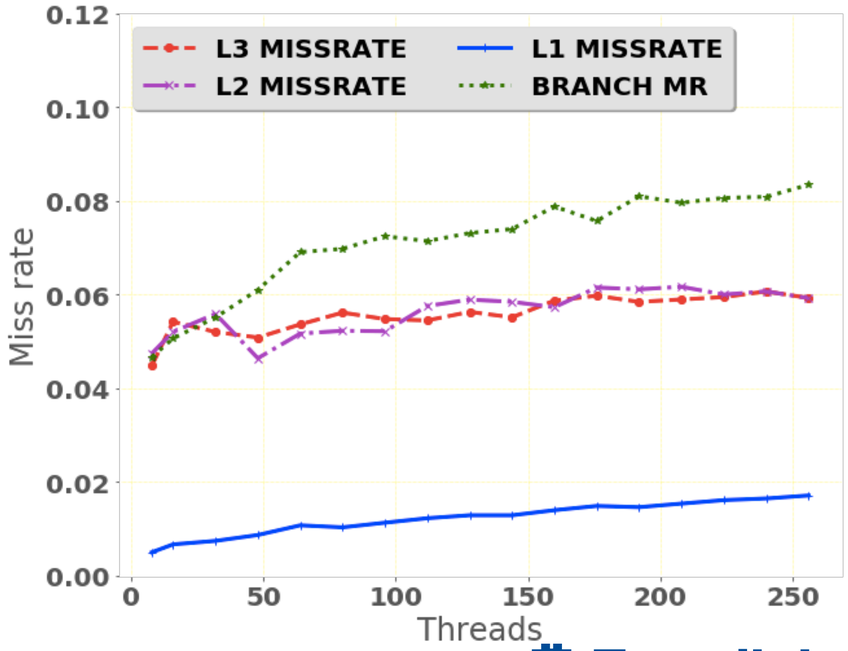

Fermilab

Hit Finding in LArTPC

- Reconstruction in LArTPC neutrino experiments is challenging, speedups are critical for future detectors (DUNE)
  - Today it takes O(100) s/event in MicroBooNE
- Hit finding algorithm promising for parallel computing since wires are processed independently - chosen for first feasibility study.
  - Signal from charged particles produces Gaussian pulses, hit finding is the process of identifying such pulses and determining their properties (peak position and width)
  - Currently a non negligible contributor to timing, takes ~15% of the MicroBooNE reconstruction workflow

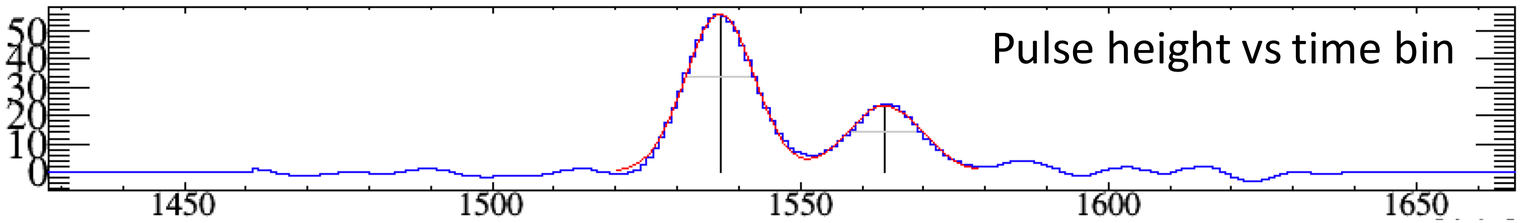

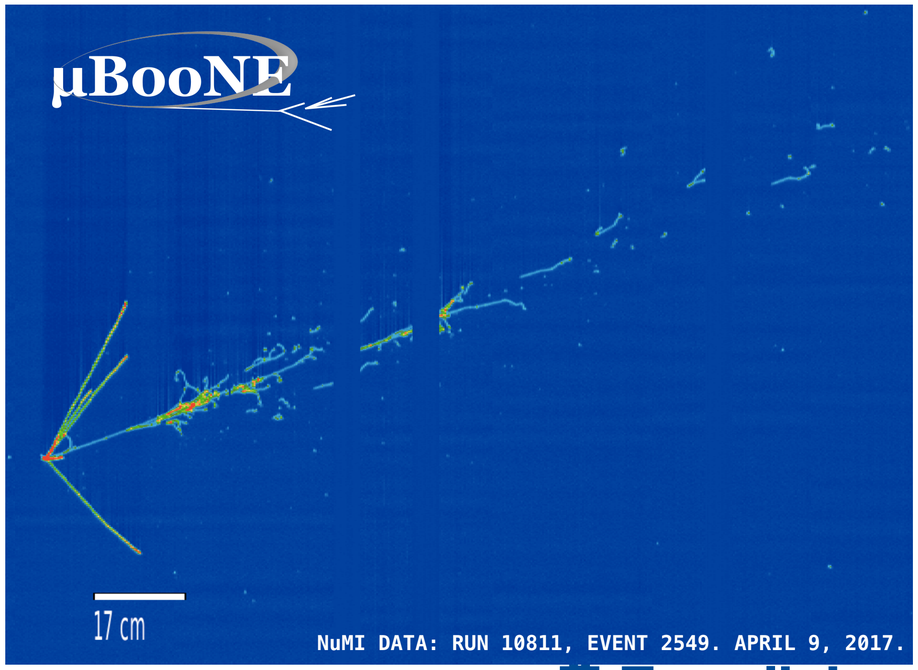

Fermilab

Hit Finding First Optimization Tests

- Produced a standalone implementation with same hit results but 8x faster than LArSoft version
- On KNL, achieved further speedups of 2.7x from explicit vectorization and of >10x from thread-level parallelization
  - Similar qualitative results on Skylake
- More work ongoing, but first tests already demonstrate potential for large speedups in LArTPC reconstruction using highly parallel architectures!
- Next steps include deploying the optimized version of hit finding and expand the work to more sophisticated/time consuming algorithms

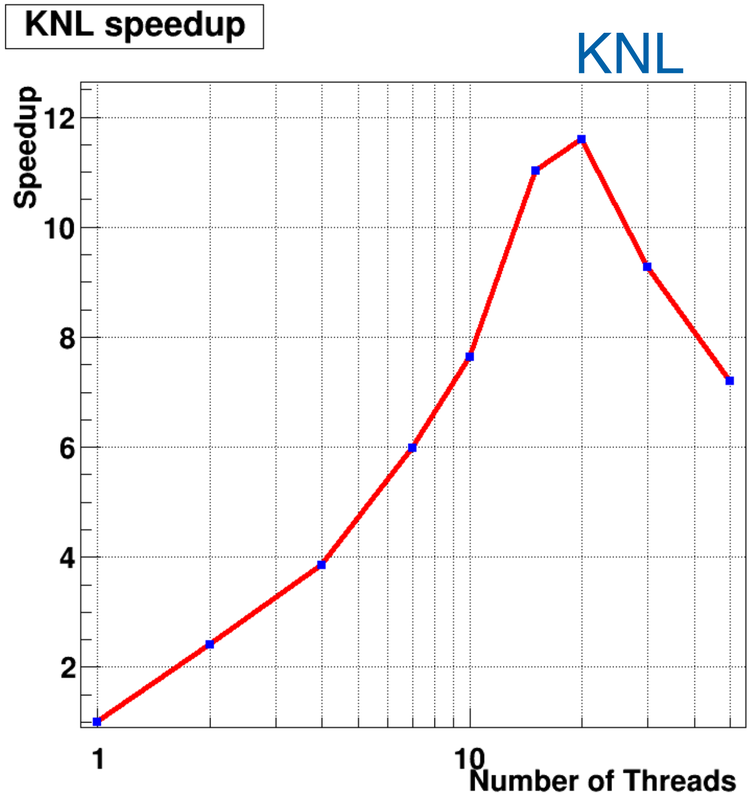

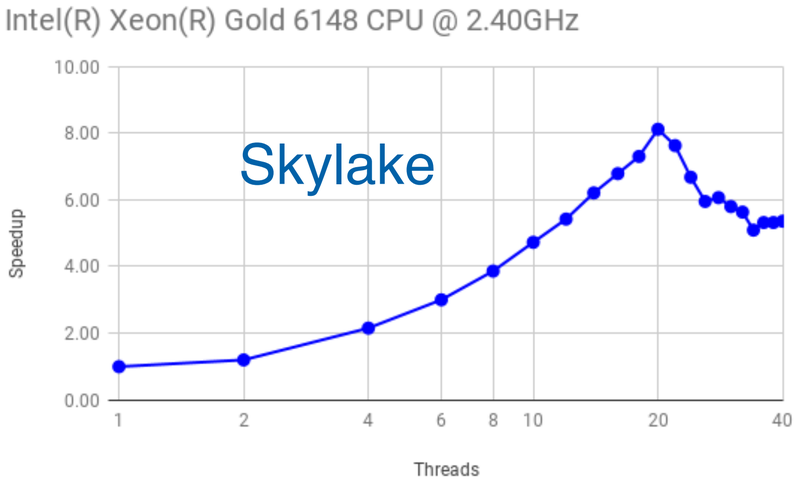

Thank you!


U.S. DEPARTMENT OF
ENERGY
Office of
Science


Fermilab


UNIVERSITY OF
OREGON