High Performance Computing for HEP in the next decade

Dirk Hufnagel (FNAL)

LAWSCHEP in Mexico City

22nd November 2019

Introduction to HPC

- A high-performance computer (HPC system) is a tool used by computational scientists and engineers to tackle problems that require more computing resources or time than they can obtain on the personal computers available to them.

- Traditionally: HPC is for complex and large computations that span many nodes and rely on inter-node communication (and therefore require fast node interconnects).

    - Inter-connects differentiates it from “just a big cluster of nodes”

    - Usually also has high performance shared storage

Introduction to HPC

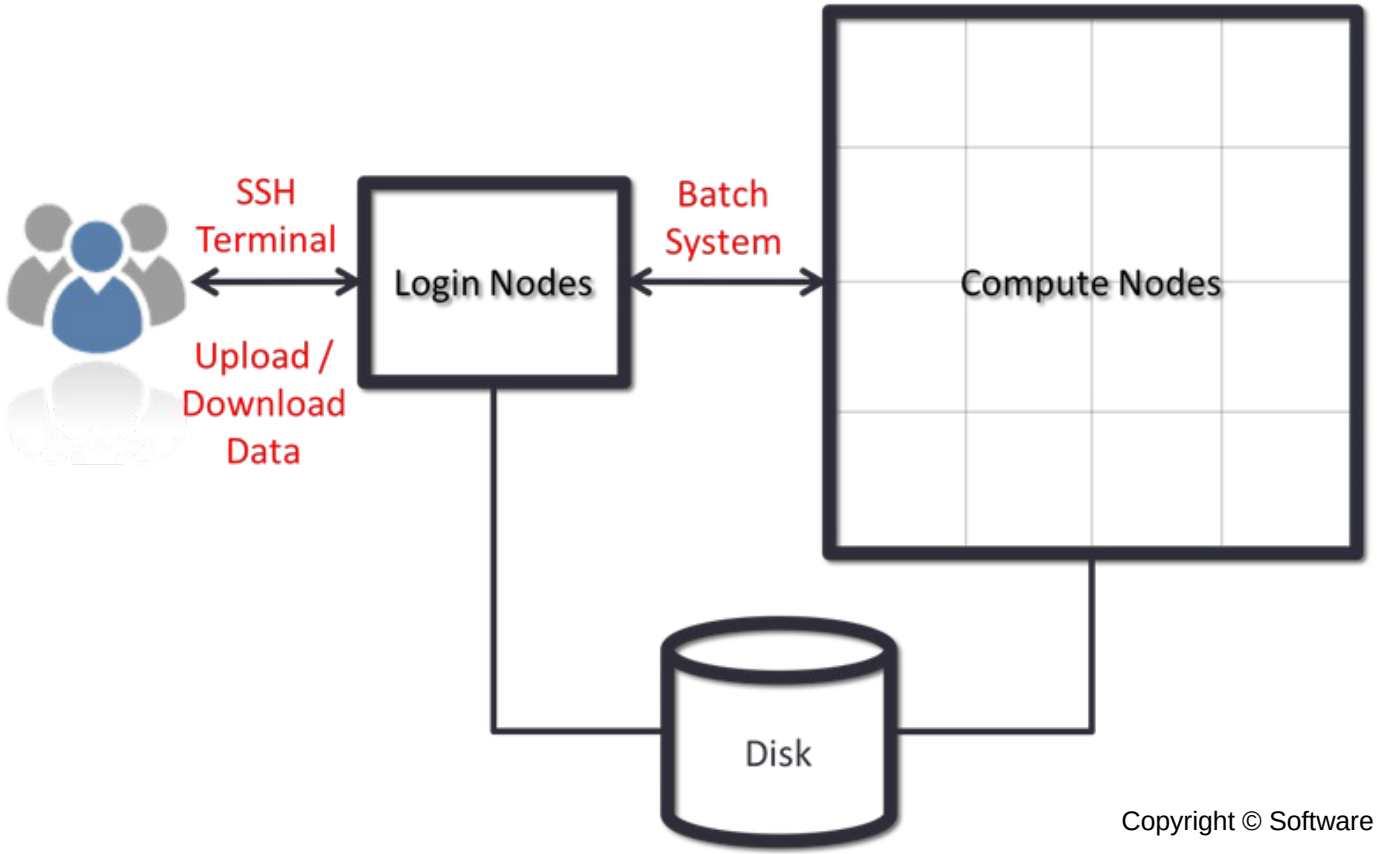

Introduction to HPC

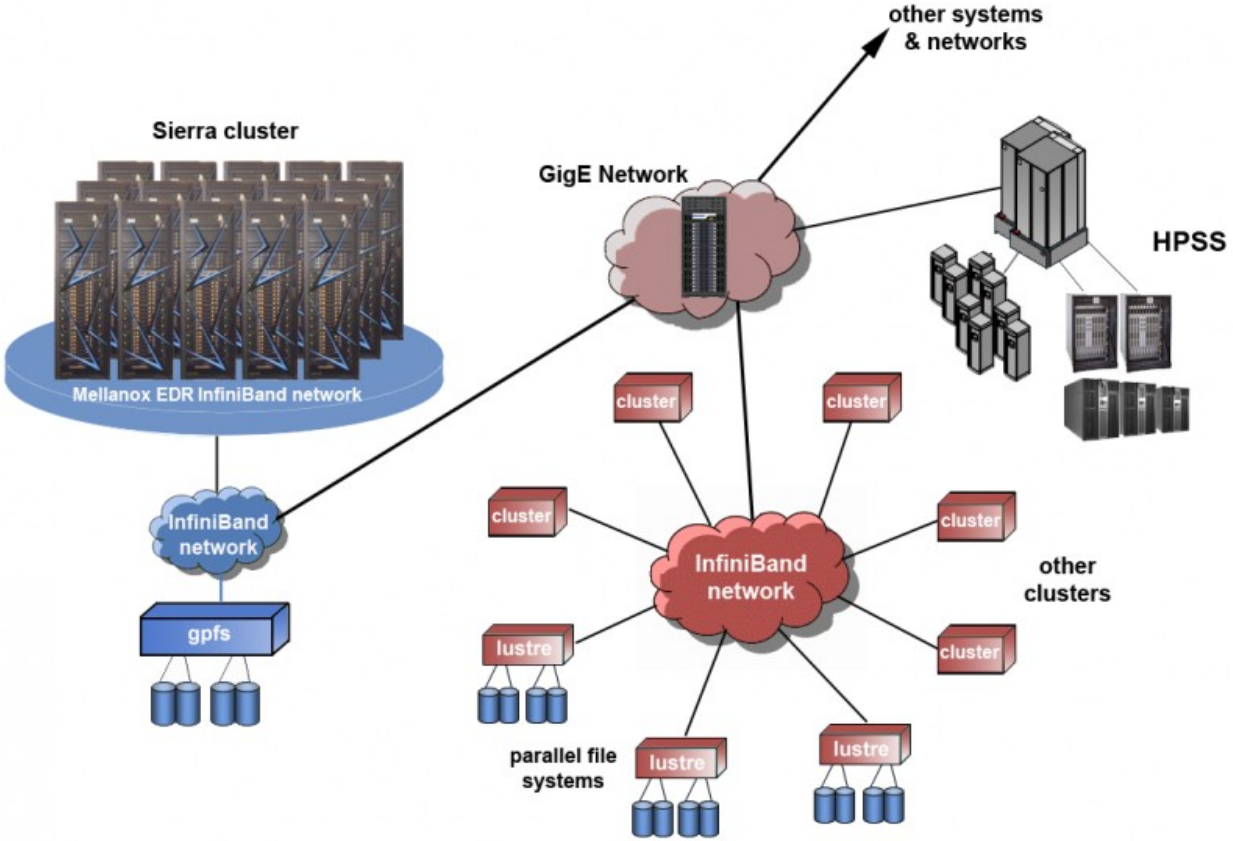

Fermilab

Introduction to HPC

- HPC in the Cloud?
- Amazon and Google sell “HPC”
- Microsoft also offers HPC with nVidia GPU (and even fast interconnect)
- This presentation targets the “traditional” HPC clusters at universities and labs (allocations, not pay-per-use)

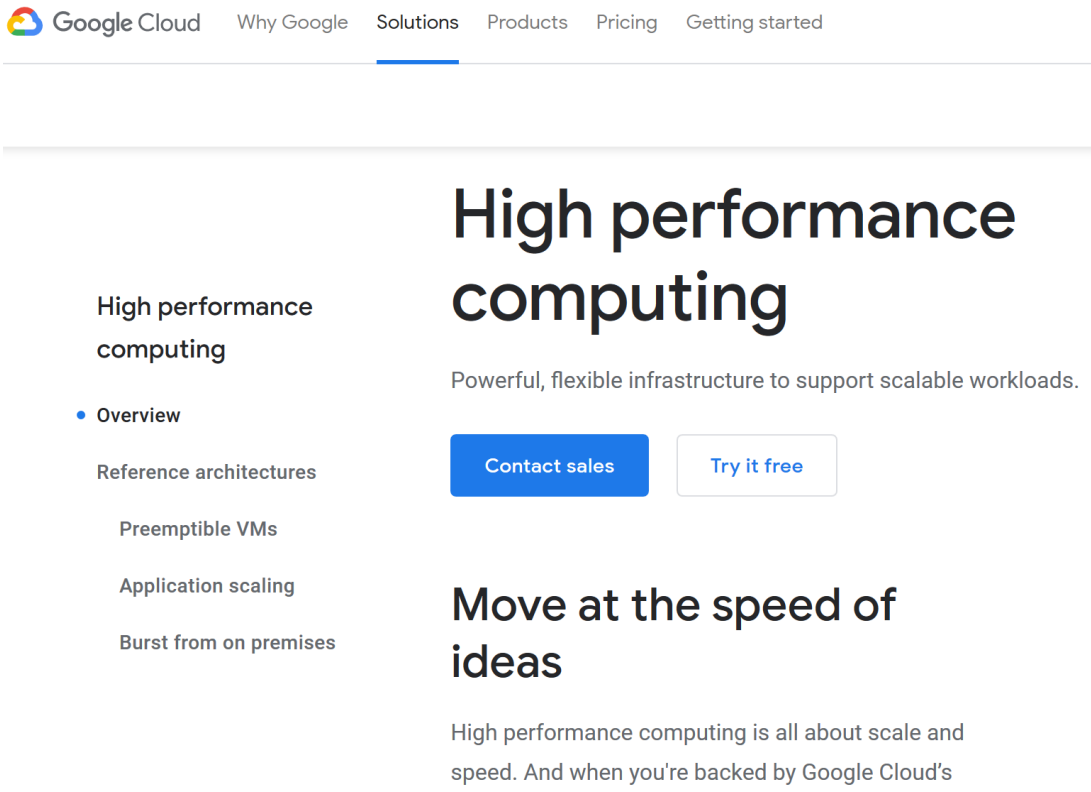

Fermilab

HPC Worldwide - Top500

- https://www.top500.org/
- Tracks 500 most powerful HPC since the 1990s (updated twice a year)
- Based on LINPACK benchmark
  - Solves a set of linear equations
  - Output is a FLOPS rating (floating point operations per second)
- Aggreate: Top500 1.65 exaFLOPS
- Lowest: 1.14 petaFLOPS

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband (/system/179397), IBM DOE/SC/Oak Ridge National Laboratory (/site /48553) United States	2,414,592	148,600.0	200,794.9	10,096
2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband (/system/179398), IBM / NVIDIA / Mellanox DOE/NNSA/LLNL (/site/49763) United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway (/system/178764), NRCPC National Supercomputing Center in Wuxi (/site /50623) China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 (/system/177999), NUDT National Super Computer Center in Guangzhou (/site/50365) China	4,981,760	61,444.5	100,678.7	18,482
5	**Frontera** - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR (/system /179607), Dell EMC Texas Advanced Computing Center/Univ. of Texas (/site/48958) United States	448,448	23,516.4	38,745.9	

Fermilab

US HPC Landscape

- DOE (Department of Energy) funds NERSC and LCF (Leadership Computing Facilities)
- NERSC : Primary scientific computing facility for the DOE Office of Science to provide computational resources and expertise for basic scientific research.
- LCF (ALCF/OLCF at Argonne/Oakridge National Labs) : Significantly more powerful than systems typically used for scientific research and available to researchers from universities, industry, and government agencies.
- Other DOE HPC only for Nuclear Weapons Simulation.

US HPC Landscape

- NSF (National Science Foundation) funds HPC in the XSEDE (Extreme Science and Engineering Discovery Environment) project.

- XSEDE's Mission: Substantially enhance the productivity of a growing community of scholars, researchers, and engineers through access to advanced digital services that support open research; and coordinate and add significant value to the leading cyberinfrastructure resources funded by the NSF and other agencies.

- To summarize: XSEDE and NERSC are more general user facilities, while LCF are intended for the big problems that can’t be solved anywhere else.

Europe HPC Landscape

- Not an expert, but my understanding is that most of it country-by-country.

- The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high performance computing service for scientists and researchers from academia and industry in Europe. The computer systems and their operations accessible through PRACE are provided by 5 PRACE members (BSC representing Spain, CINECA representing Italy, ETH Zurich/CSCS representing Switzerland, GCS representing Germany and GENCI representing France).

- EuroHPC, a 1 billion Euro joint initiative between the EU and European countries to develop a World Class Supercomputing Ecosystem in Europe.

HPC developments in the last years

- Historically the Top500 has been dominated by the US. In the last years China has increased their share significantly, now having 227 entries compared to 118 for the US. The US still leads in aggregate FLOPS, 37.8% vs. 31.9% of the total.

- 2015 US National Strategic Computing Initiative for the development of exascale HPC, brought together existing programs.

- First exascale HPC expected in the US in 2021

- China : goal to have exascale HPC in 2020

US HPC outlook

- We will get exascale HPC in the next few years, average lifetime of such a machine is about 5 years

- Due to their sheer size and computing capacity they are very attractive targets for any kind of scientific computational needs

- In addition, DOE funds the exascale HPC but also funds a major fraction of HEP and computing for HEP in the US. Scale of HEP computing budgets questionable with the increases in deployed HPC capacity.

US HPC (LCF) current and future

- Available now:
  - OLCF Summit : Power9 CPU + nVidia GPU, ~150 petaFLOPS
  - ALCF Theta : Intel KNL, ~7 petaFLOPS
  - NERSC Cori : Intel Xeon + KNL CPU, ~14 petaFLOPS

- Available ~2021
  - OLCF Frontier : AMD EPYC CPU + Radeon Instinct GPU, >1.5 exaFLOPS
  - ALCF Aurora : Intel Xeon CPU + X^e (GPU), >1 exaFLOPS
  - NERSC Perlmutter : AMD EPYC CPU + nVidia GPU, ~100 petaFLOPS

- Remainder of talk focuses on the HPC/GPU since they provide the largest amount of resources, but targeting other (more accessible) HPC as well.

Challenges to use the LCF

- HEP scientific computing has mostly standardized on x86-64 Linux systems, since they are cheap, readily available and widely used in industry and the internet economy

- Completely different HPC system architectures pose challenges

- Different CPU architectures ”in principle” usable with a recompile, but use of GPU requires rethinking algorithms, applications or even whole approach on how to go from experimental data to physics insight

- Different GPU architectures between the LCF doesn’t help

Challenges I won’t talk about here

- Data management, how to get data into HPC and out of them
- Workflow Management, how to integrate HPC into existing workflow management systems that work with other resource types (grid)
- Dealing with locked-down HPC, where nodes don’t have internet access (and how that affects science applications that rely on central services)
- Dealing with science application that require special system software
- Yearly allocation process and how this affects multi-year computational plans for long-running scientific experiments

HEP Computing use of HPC/GPU

- Exascale timeline ~matches with some big computational challenges in HEP
- HL-LHC starting 2026
- DUNE expected to start taking data in 2026
- Have some years to get ready for these, but large scale changes in software take time

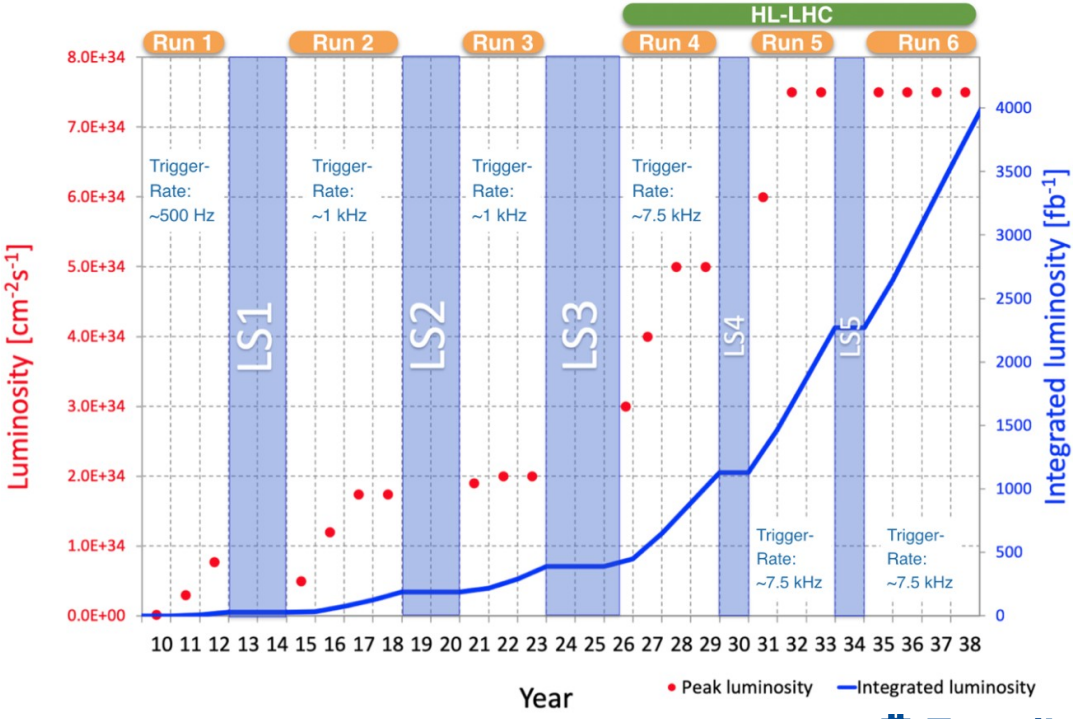

Fermilab

CMS, ATLAS and Neutrino (DUNE, Nova, others)

- Non-exhaustive summary of HPC/GPU efforts for different experiments
- Not saying much about Machine Learning here, separate talk for this

CMS

- CMS consists of many subdetectors whose data can be processed concurrently
  - E.g. local reconstruction of tracker and calorimeters is independent
- Many computations also depend on information from multiple subdetectors
  - E.g. particle flow

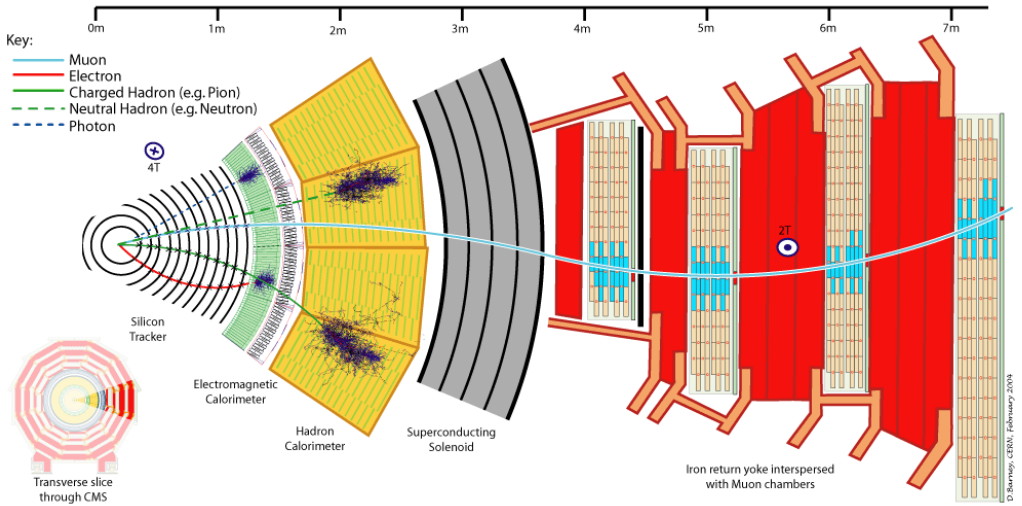

CMS data processing software, CMSSW

- CMSSW is multithreaded with Intel Threading Building Blocks (TBB)
  - In production since 2015 (as the first of LHC experiments)
- CMS reconstruction application consists of ~1300 modules
  - “Module” is a piece of code that does some transformation on the data of an event
  - “Module” is the smallest piece of work scheduled to be run by CMSSW framework
- CPU time profile is close to flat wrt modules
  - Historically track reconstruction is the most outstanding piece
    - Pattern recognition (13 modules) takes ~25 % of reconstruction time (2018 conditions)
      - Pattern recognition is based on Combinatorial Kalman Filter
    - Other components included (~170 modules) takes ~50 % of reconstruction time
- Challenges for code porting to non-CPU architectures
  - Many modules to port
  - Gain from porting a single module is relatively small

Efforts in tracking (incomplete list)

- Patatrack
  - Pixel local reconstruction, and pixel tracking and vertexing on NVIDIA GPUs
  - Pioneering direct GPU usage from CMSSW
  - Main goal is HLT for Run3
    - Concrete proposal for GPUs in HLT farm was presented to CMS in the summer
- mkFit
  - Parallelized Kalman Filter tracking
    - Vectorized math, internally multi-threaded
  - Main targets are many-core CPUs, Intel Knights Landing (now discontinued)
  - Exploring GPU implementation
  - Main goal is HLT for Run3 (running on CPUs)
- Exa.TrkX: Charged Particle Tracking using Graph Neural Network
  - More exploratory

Facilities

- GPUs are being considered for HLT in Run 3
  - Current porting R&D sufficient for null-sum change from CPU to GPU
  - But unable to fully load the GPUs
  - Forward-looking, GPUs have “free capacity” for
    - more algorithms to be ported to GPUs
    - existing algorithms to be run more often
      - e.g. pixel tracking for each L1-accepted event

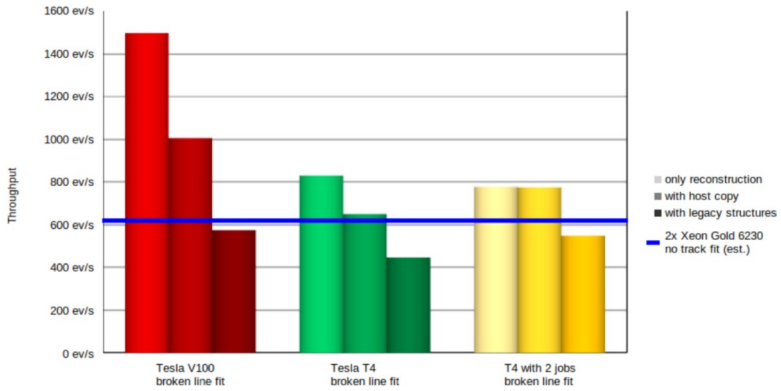

- Many of top supercomputers have most of their FLOPs from GPUs
  - In US, Summit in production since 2018; Perlmutter, Aurora, Frontier coming in 2020-21
  - May not be allowed to run unless GPUs are utilized

Porting status

- Current prototypes for GPU algorithms cover about
  - 5 % of offline reconstruction time
  - 24 % of HLT reconstruction time
- Algorithms include
  - Pixel local reconstruction, pixel tracking, pixel vertexing
  - Parts of ECAL and HCAL local reconstruction (pulse height fit)
- Work ongoing to
  - Bring the prototypes to production
  - Port more algorithms
    - Work on strip local reconstruction is ongoing

Summary

- We started the direct GPU effort with CUDA as a first step
- We are looking into performance portability solutions like Kokkos, Alpaka, SYCL
- Our main philosophy is to view a GPU as an additional computation resource
  - In principle already a GPU algorithm slower than its CPU version would be beneficial, if the CPU has some other work to do while the GPU is running

ATLAS and HPC/GPU

- Long experience with GPU for online trigger, initially decided not to deploy GPU for Run3 HLT, but that decision is being re-evaluated (but the decision is still no).

- Several ongoing projects for offline processing
  - Using ML for tracking (seeding)
  - Implementing current non-ML tracking on GPU
  - Porting parts of ACTS to GPUs
  - Re-writing gVEGAS and other parts of madgraph for GPUs
  - Fast (parametrized) simulation of the Calorimeter

- Working on how to integrate offloading of tasks to GPU into framework and both synchronous and asynchronous scheduling of CPU/GPU tasks
- Looking at “performant portability solutions” like Kokkos, SyCL, OpenMP, etc.

Neutrino Science and HPC/GPU

- Liquid argon detectors ideal for accelerated signal processing
  - well suited for GPU

- "traditional" HEP reconstruction
  - events are huge (time window), algorithms scale less than n-square
  - but large detector with independent detector regions
    - different regions can be independently reconstructed in parallel
    - can apply a lot of resources in parallel to process single event

Neutrino Science and HPC/GPU

- Statistical parameter estimation
  - Many fits across lattice, very well suited to MPI
  - Near-perfect strong scaling from 1 to 300k ranks
    - Computation/communication ratio very high
    - Reduction steps do no dominate
  - HPC allows denser grids and higher dimension in parameter space

- Inversion problem
  - Find initial (invisible) neutrino from reconstruction output
  - Needs lots of memory, benefits of whole node scheduling at HPC

Neutrino Science and HPC/GPU

- Analysis on HPC
  - Take advantage of high bandwidth shared storage for n-tuple processing
  - Many parallel analysis jobs accessing the same large HDF5 file
  - Easily repeatable, allows exploration of option space

- Example: hours to run analysis on grid with standard MB files input
  - HPC with single very large HDF5 file took seconds

- Preload many analysis jobs into a single HPC many-node batch job

Summary

- Large increase in computational demand in HEP in the next 5 years
- Large increase in deployed HPC capacity due to the race to exascale

- Both opportunity and challenge

- Porting current algorithms to GPU and increase throughput by offloading these to GPU while in parallel using the CPU ongoing. Unclear if this approach will lead to enough offloading to make efficient use of next-gen HPC.

- Fully exploiting these GPU-heavy next-gen HPC likely requires new approaches in how we go from experimental data to physics insight (Machine Learning?).