Micro-Retreat: CMS R&D (an overview without claiming to be complete)

Computation

Many-Core

- CPU like hardware, not necessarily Intel (for example PowerPC, ARM, but primarily x86)
- Trend to more and less-powerful cores with less memory: 1000 CMS physicist program thread safe now (success)
- Challenges:
    - Physics validation of software for diverse platforms and more and more architectures
    - Increase thread-counts (> 100 threads, …) with limited memory
- Projects
    - Continue to maintain and evolve the CMSSW framework
    - ROOT I/O together with DIANA http://diana-hep.org
    - IPCC (Intel Parallel Computing Center) on ROOT code modernization (https://ipcc-root.github.io)
    - Investigation of future processor architectures together with DIANA: http://diana-hep.org
    - DOE NESAP grant to optimize software for HPC centers (http://www.nersc.gov/users/computational-systems/cori/nesap/)

Accelerators <-> Algorithms

- GPU, FPGA, (more exotic hardware)
- Techniques ow to use these accelerators: vectorization/SIMD vs. Machine Learning
- Challenges:
    - Social engineering: getting people to change the way they think about programming
    - Bridge: physicists developing algorithms <-> computer professionals optimizing code for advanced architectures
- Projects:
    - Parallel Kalman Filter Tracking
        - R&D to develop parallelized and vectorized algorithms to implement the analog of the iterative Kalman Filter tracking used today in CMS. Focused on multicore Xeon, Xeon Phi, KNL and GPUs.
        - http://trackreco.github.io
    - HEP.TrkX

        - DOE HEP/ASCR pilot project funded to explore the use of non-traditional algorithms for tracking, e.g. Machine Learning, Deep Learning, ... (FNAL/LBNL/Caltech)
        - https://heptrkx.github.io
    - HEP Event Reconstruction with Cutting Edge Computing Architecture
        - Recently funded 3 year project (FNAL, UOregon) covering LHC and Neutrino program, collaborates with Tracking R&D projects (previous slide)
        - http://computing.fnal.gov/hepreco-scidac4/ , partner with FNAL Scientific Computing Division and others in SciDAC-4 DOE program
        - new computing architectures in high-energy physics (HEP) event reconstruction
    - Patatrack: towards heterogeneous computing in CMSSW
        - CMS project at CERN to use NVidia GPUs in the high level trigger
        - Part of the pixel track reconstruction re-implemented on CUDA
    - FPGA's for ML inference or more
        - See talk of Nhan at https://indico.fnal.gov/event/16923/
    - GeantV
        - Vectorization and parallelization project for Geant simulation: http://geant.cern.ch/content/about-geantv
        - transport engine research is funded by DOE COMP-HEP (FNAL, SLAC),
        - FNAL contributes to geometry, EM physics, random number generation, as well as prototyping and testing in the context of the CMS experiment framework
        - SLAC and FNAL collaborate on a project to modularize hadronic physics and explore the feasibility of its vectorization
        - CMS deployed VecGeom, the vectorized geometry library developed in the context of GeantV, in February 2018. VecGeom has been integrated to Geant4 in scalar mode and yielded speedups of 4-13%

Storage

Store Data

- Interested in underlying storage architectures and storage systems
- Challenges:
    - Concurrency
    - Wide area streaming access
- Projects:
    - CEPH investigations to replace T2 storage systems (currently HDFS)
    - Xrootd-caching deployments at T2 sites

    - NANOAOD (goal: ~1 kB/event analysis format for 50% of the collaboration's analyses)

Organize Data

- How to distribute data and lower the complexity and effort for the distributed infrastructure
- Challenges:
    - Need to reduce effort to run distributed infrastructure (less sites, less storage elements)
- Projects:
    - Data Lakes (more automation)
    - Data Management Review in 2018 (replace PhEDEx): candidates are RUCIO and Dynamo

Analyze Data

- Change paradigm of analysis (ntuples to mini ntuples to plots might not be sustainable in the future)
- Use industry tools
- Challenges:
    - Social engineering
- Projects:
    - Striped database (LDRD: http://ldrd.fnal.gov/subdir/FNAL-LDRD-2016-032-D4.pdf)
    - Apache Spark
        - Toolkits to access ROOT files in Spark and bridge between HEP ecosystem and industry tools together with DIANA http://diana-hep.org
        - Industry collaborations: CERN openlab/Intel/Fermilab/Princeton CMS Big Data Project (https://cms-big-data.github.io)
    - HEP Data Analytics on HPC (SciDAC-4, http://computing.fnal.gov/hep-on-hpc/
        - CMS involved in generator tuning: exploit HPC facilities (compute, memory, and storage) to tune Pythia, updating workflow and data management especially for HPC machines
    - Provide access to GPUs through OSG (including developing ML benchmarks for GPUs)

Services

Elasticity

- Need ability to scale out dynamically
- Challenge:

    - Communication / Planning
- Projects:
    - HEPCloud (http://hepcloud.fnal.gov), coordinated access to all forms of resources: Grid, commercial cloud, HPC; and elastic scale out
    - CMS Resource Provisioning and Workload Management Review in 2018

Connectivity

- Network is of great importance to our science!
- Challenge:
    - Network capacity not infinite anymore (at least that is what everyone is telling us) ➜ scheduling, reservation will be needed
- Projects and interested in:
    - SENSE: three-year project to address multi-domain end-to-end SDN only
    - Recently introduced to Big Data Express (https://bigdataexpress.fnal.gov/)
    - Working with ESnet and USATLAS to determine network needs for HL-LHC: ESnet6
    - CMS collaborators involved in Pacific Research Platform (http://prp.ucsd.edu/)
    - Security: Work on non-X509-based authorization mechanisms for resource and WLCG storage access

Orchestration

- How to deploy diverse applications and services over distributed resources
- Challenges:
    - Hide complexity
- Interested in:
    - Kubernetes and Containers
    - SLATE (http://slateci.io/)
    - CERN swan to replace interactive login clusters and provide Jupyter notebook access to analysis resources: http://swan.web.cern.ch/