US Software and Computing R&D projects

Peter Elmer - Princeton
Oliver Gutsche - FNAL

USCMS collaboration

- Collaborators and colleagues at USCMS universities and laboratories drive R&D ➜ **People**
- Develop new and/or evolve established technologies ➜ **Innovation**
- Embedded in International CMS ➜ **Collaboration**
- Core Research Program and Operations Program, plus additional solicitations and funding opportunities from NSF and DOE ➜ **Funding**
- USCMS Software and Computing Operations Program
  - Supports R&D projects, often through partnerships with separately funded PI-initiated projects
  - Facilitates continuity for people at universities and laboratories
  - Provides infrastructure and operational support for integration and deployment
  - Provides long term maintenance and support

Software

Tracking Pattern Recognition

- Speed up dominant component of reco time, especially at HL-LHC pileup
- 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.
  - Began with 3 year NSF funded project (Princeton/UCSD/Cornell) + USCMS support and FNAL SCD & SciDAC collaboration. The NSF part nominally ends this year, but with no-cost extensions will probably run into 2019.
  - http://trackreco.github.io
  - Embedded in this is an effort to look at segment linking and optimization of the tracker layout for HL-LHC
- 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

Simulation

- GeantV
  - An ongoing FNAL contribution (DOE COMP-HEP funded in part)
  - Vectorization and parallelization project for Geant simulation
    - http://geant.cern.ch/content/about-geantv
  - First developments are already in production in CMSSW (vecgeom) ➜ CMS first adopter
- Pre-mixing
  - Was developed for LHC Run 1 and 2
  - Now driving implementation for Phase 2 upgrade simulations

Visualization & Geometry

- Visualization & ROOT 7 (UCSD/USCMS):
  - many changes in GUI components planned for ROOT 7: switch from QT based to web based visualization libraries ➜ big consequences for event display solution Fireworks
  - R&D conceptualization: client-server architecture
    - adapt to ROOT 7 changes
    - evolution and maintenance of EvE

- Geometry (USCMS/FNAL):
  - CMS geometry currently implemented in DDD in CMSSW (CMS' own development)
  - Proposal to switch to DD4HEP, community solution that is already used by LHCb
    - https://dd4hep.web.cern.ch/dd4hep/page/about/

Multi-threading and ROOT - USCMS Ops efforts

- Multi-threading in CMSSW
  - Core software team continues to increase parallel thread counts
    - Continuous monitoring and maintenance of multi-thread safety of physics algorithms
  - Current limitation to go to multi-hundred parallel threads ➜ ROOT I/O
- Core software team is looking into event processors
- ROOT I/O improvements
  - Leading ROOT I/O team
    - Partner with DIANA project
  - Improving performance in ROOT and CMSSW
- Also partner with IPCC (Intel Parallel Computing Center) on ROOT code modernization, in particular C++/Python interoperability (PyROOT)
  - https://ipcc-root.github.io

New Architectures - additional FNAL projects

- New architectures (at HPC centers)
  - DOE NESAP grant to optimize software (http://www.nersc.gov/users/computational-systems/cori/nesap/)
  - NSF Supercomputing Center optimization grants in preparation

- 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

- Bridge between physics algorithm development and core framework software
  - Partner with FNAL Scientific Computing Division to help physicists write algorithms that can exploit advanced hardware architectures

Analysis

DIANA/HEP

- Project focused on the development of analysis tools: interoperability, performance, collaborative analysis.
- 4 year NSF funded project (Princeton/U.Nebraska/NYU/Cincinnati). Nominally ends mid-2019, but may continue in part with no-cost extensions into 2020.
- Recent activities: ROOT I/O improvements, interoperability of HEP tools (ROOT) with Spark, ROOT I/O in Python w/o ROOT (uproot), Scikit-HEP toolkit, etc.
- Note in particular the frequent DIANA/HEP “topical” meetings on R&D topics: https://indico.cern.ch/category/7192/
- http://diana-hep.org

CMS Big Data Projects

- FNAL/Princeton/CERN(IT,Openlab)/Intel, USCMS Ops partnership
- Investigate industry “Big Data” tools for HEP analysis
  - https://cms-big-data.github.io
- Analysis Thrust
  - Full analysis in Apache Spark
- Intel Big Data Reduction Facility
  - “1 PB to 1 TB in 5 hours”
- Fermilab LDRD (lab directed R&D)
  - Striped data representation with multiple layers of caching
- Using developments from DIANA and CERN-IT
  - spark-root
  - spark-xrootd-connector
  - uproot

HEP Data Analytics

- New capabilities at ASCR computing facilities
  - exploit HPC facilities (compute, memory, and storage) to meet new HEP data analysis demands, enabling computationally expensive physics studies to be completed on time-scales not currently feasible
  - http://computing.fnal.gov/hep-on-hpc/
  - FNAL Scientific Computing Division, Argonne, LBNL, Colorado State, UCincinnati
  - SciDAC-4 DOE program, recently started 3 year(?) project
- Integration of new analysis tools in CMSSW ecosystem
  - Infrastructure facilitating R&D
  - Python analysis tools from the Data Science community, promote and enable python as analysis language for HEP (eg, Tensorflow now available from CMSSW)
  - Site tuned installations of toolkits provided to CMS collaborators by OSG & CMS connect (facilitating R&D on GPUs, etc)
  - New interest: FPGA for ML inference and beyond

Facilities

HEPCloud

- Portal to an ecosystem of diverse computing resources commercial or academic
  - http://hepcloud.fnal.gov
  - USCMS Ops partnership with FNAL Scientific Computing Division
- USCMS strategy for elastic scale out of compute resources
  - T1_US_FNAL endpoint scales into
    - Allocations at US supercomputer centers (HPC)
    - Commercial cloud providers
    - Opportunistic resources
  - Multiple demonstrations on AWS and Google Cloud
  - Intelligence to make decision where to send work
- Provide new capabilities for CMS
  - Use resources when work is available, get work done quicker
  - Request to WLCG to change pledging process to two components: constant + elastic

HPC

- Significant push by US funding agencies (especially DOE) to utilize HPC centers for CMS workflows
  - HPC center architectures range from x86-compatible resources to PowerPC + GPU based “exascale” systems
  - CMS does not have significant needs for generator-only or simulation-only work
  - Assume that HPC centers will have to be able to run GEN-SIM-DIGI-RECO and data RERECO (especially in HL-LHC)
- US
  - Enabling centers individually to run CMS work
  - Requesting generic or dedicated allocations for CMS work
    - For 2018: allocations at DOE and NSF supercomputers
  - HEPCloud key to facilitate HPC center usage for CMS
- Software efforts are needed to be able to exploit the capabilities of the new architectures
  - But also support for multiple architectures and validation of the software will be needed

Facility R&D - USCMS Ops efforts

- Ceph
  - Demonstrating the use of Ceph as a storage system for a Tier-2
- XRootD caching
  - Deploying increasingly-large XRootD caching prototypes until they are at the scale of a standalone Tier-2
- Security
  - Work on non-X509-based authorization mechanisms for resource and WLCG storage access
- Networking
  - SENSE: address multi-domain end-to-end SDN
    - Partners: ESnet, FNAL, ANL, NERSC, Caltech, & the University of Maryland
  - ESnet6: Partner with ESnet to evolve the current domestic and transatlantic network infrastructure
- Machine Learning
  - ML as a service
  - ML benchmark
  - Collaboration with U.Florida/CHREC/SHREC (http://www.chrec.org/), high performance reconfigurable computing

Infrastructure Software

Resource Provisioning, Workflow management and Data management

- glideinWMS and HTCondor as core of resource provisioning
  - USCMS Ops partnering with FNAL Scientific Computing Division, OSG and HTCondor
  - New: GPU grid access through OSG
- Data management
  - Dynamic Data Management, USCMS Ops partnering with MIT
  - Rucio evaluation: USCMS Ops partnering with FNAL Scientific Computing Division, OSG
- Data Lakes
  - See previous talk by Brian in this meeting series
- Workflow management
  - “Work unit” instead of jobs
  - Review in 2nd half of 2018 to define direction

Training

Training

- To support the R&D activities above, training is needed
- USCMS has participated in CMSDAS, of course, which is one eventual vector for introducing new tools for analysis to users.
- The U.S. has however lacked a HEP-focused school for advanced development training for software. Schools like ESC (INFN, Bertinoro) have trained many people involved in aspects of advanced development in CMS over the years. In the U.S. this has fallen a bit between the cracks of various programs and missions.
- NSF funded CoDaS-HEP (http://codas-hep.org) for 2017 and 2018 (applications soon!), which aims to provide a broad introduction: parallel programming, Big Data Tools, Machine Learning technology and methods, performance, development tools.
- We are exploring how to sustain this going forward.

Community

S2I2-HEP Conceptualization

- “Planning” project whose deliverables included enabling the process to produce the CWP document as well as a “Strategic Plan” for a possible NSF-funded software institute.
- 2 year NSF funded project (Princeton/UIUC/Cincinnati), nominally completed with these deliverables:
  - Strategic Plan for a Scientific Software Innovation Institute (S2I2) for High Energy Physics - arXiv 1712.06592
  - A Roadmap for HEP Software and Computing R&D for the 2020s (Community White Paper) - arXiv 1712.06982
- However some participant cost funds remain and these are being spent to support travel to CWP follow-on activities, e.g. HSF/WLCG and other small planning workshops.
- http://s2i2-hep.org