Software R&D Strategy

Adam Lyon

2nd ICAC Meeting

15 October 2019

Reminder: Software & Computing Research and Development

Triggers (the why):

A. Requirements from experiments based on upcoming needs
B. Forward thinking to keep up with evolving computing landscape
C. Useful technologies that scientists adopt and needs support
D. Fruitful collaborations

Drivers (the what):

A. CMS in the HL-LHC era and DUNE
B. New computing architectures/accelerators and the Exascale High Performance Computing Era
C. Machine Intelligence's impact on HEP reconstruction and analysis
D. Specific funding calls (e.g. SciDAC from DOE-ASCR)

Guides (the how):

- Physics goals (of experiments and scientists)
- Software and Computing requirements from CMS and DUNE
- Community White Papers (HEP Software Foundation and IRIS-HEP)
- Goals of SciDAC and ECP
- Strive for common tools where possible and common principles for moving forward

There is overlap, of course

Reminder: R&D Activities Overview - A broad program

- Physics and detector simulations with advanced architectures and techniques
- Accelerator Modeling on HPC
- Evolution of Infrastructure Frameworks (CMS, DUNE) and Root
- HPC, Advanced architectures/accelerators, multithreading
  - Containerization
  - HEP Data Analytics
  - Reconstruction
  - Spack & SpackDev [HPC compatible packaging]
- Machine Intelligence
- Data Acquisition
- Advanced networking (BigData Express)
- Workflow (HEPCloud)
- Astro (CCD/MKIDs)
- QIS now has its own program and I won't discuss, but some personnel comes from SCD (myself included)

Funding comes from many sources

- DOE-OHEP (CompHEP)
- USCMS Software and Computing (S&C) Operations Program
- SciDAC-4 [DOE-ASCR] $17.5M awarded total
  - 5 yr and 3 yr projects started in FY18
- Fermilab LDRD (Lab Directed R&D)
- Exascale Computing Project (ECP)
- HEP-CCE (Center for Computational Excellence)
  - Promote excellence in HPC and R&D
  - Enhance connection to ASCR
  - FNAL, ANL, BNL, LBNL
- Other experiment projects & Detector R&D (KA25)
  - e.g. CMS Outer Tracker, Mu2e TDAQ
- We supplement with SCD funds

Personnel may be matrixed across projects

This time compared to last time

- Last time
  - I listed our R&D efforts
  - It was a long list

- This time
  - Explain our R&D Strategy
  - Address recommendation
    - Focus on priorities, clearly identify topics
  - Give examples of efforts within
  - Explain the funding sources
  - Main focus is above, *not* project status

Reorganization

- You heard from Jim about the reorganization
- More emphasis on projects and project management
- R&D is spread across the departments
- Department of project leaders that cross-cut across the division
  - Some R&D project leaders are in this department
  - Each R&D project has a leader or PI
- Overall stewardship of R&D is at the division head level
  - We will leverage synergies between CMS and LAr Neutrino as much as we can (e.g. CCE & SciDAC's)
- For some R&D projects (e.g. Quantum Computing), SCD provides effort

Software R&D Strategy

- Topics
  - AI
  - Evolving Computing Architectures
  - Compute (aka HEPCloud)
  - Storage
  - Networking
  - Analysis Facility Concepts
  - Supporting Quantum Computing (Fermilab Quantum Institute)

- These topics are aligned with strategic directions of the division and the lab

Artificial Intelligence

- Major strategic topic for the lab and SCD
- Nhan Tran moving to SCD to head **AI**-centered group within the **Artificial Intelligence** and Software for Physics Applications (AISP) department
- Goal is a large role in a to-be-competed DOE AI Center
  - Must show success in AI efforts

- Many efforts
  - DOE Early Career Award (new!) - Deep learning for boosted Higgs program
  - CompHEP: High Velocity AI (Fast inference, Distributed training, Uncertainty Quantification) [with ORNL and ANL] (new!)
  - LDRD: Accelerator Control with AI (new!)
  - LDRD: Graph Neural Networks for Calorimetry and Event Reconstruction (new!)
  - LDRD: Modeling Physical Systems with Deep Learning Algorithms

- New Computational Science Seminar series to explore collaboration opportunities with Fermilab, Argonne, and U of Chicago
  - First topic is AI. First talk is in November by Nhan

Evolving Computing Architectures

- **HPC Computing for the next five years (new announcements since last time)**
  - **OLCF: Summit** — IBM CPUs & 27K NVIDIA Volta GPUs (#1 supercomputer in the world)
  - **NERSC: Perlmutter** — AMD CPUs & NVIDIA Tensor GPUs (2020)
  - **ALCF: Aurora** — Intel CPUs & Intel Xe GPUs (early 2021) — first US **Exascale** machine
  - **OLCF: Frontier** — AMD CPUs & AMD GPUs (later 2021) - **Exascale**

- **GPUs are the big winners.**
- **Vectorization looks less critical**

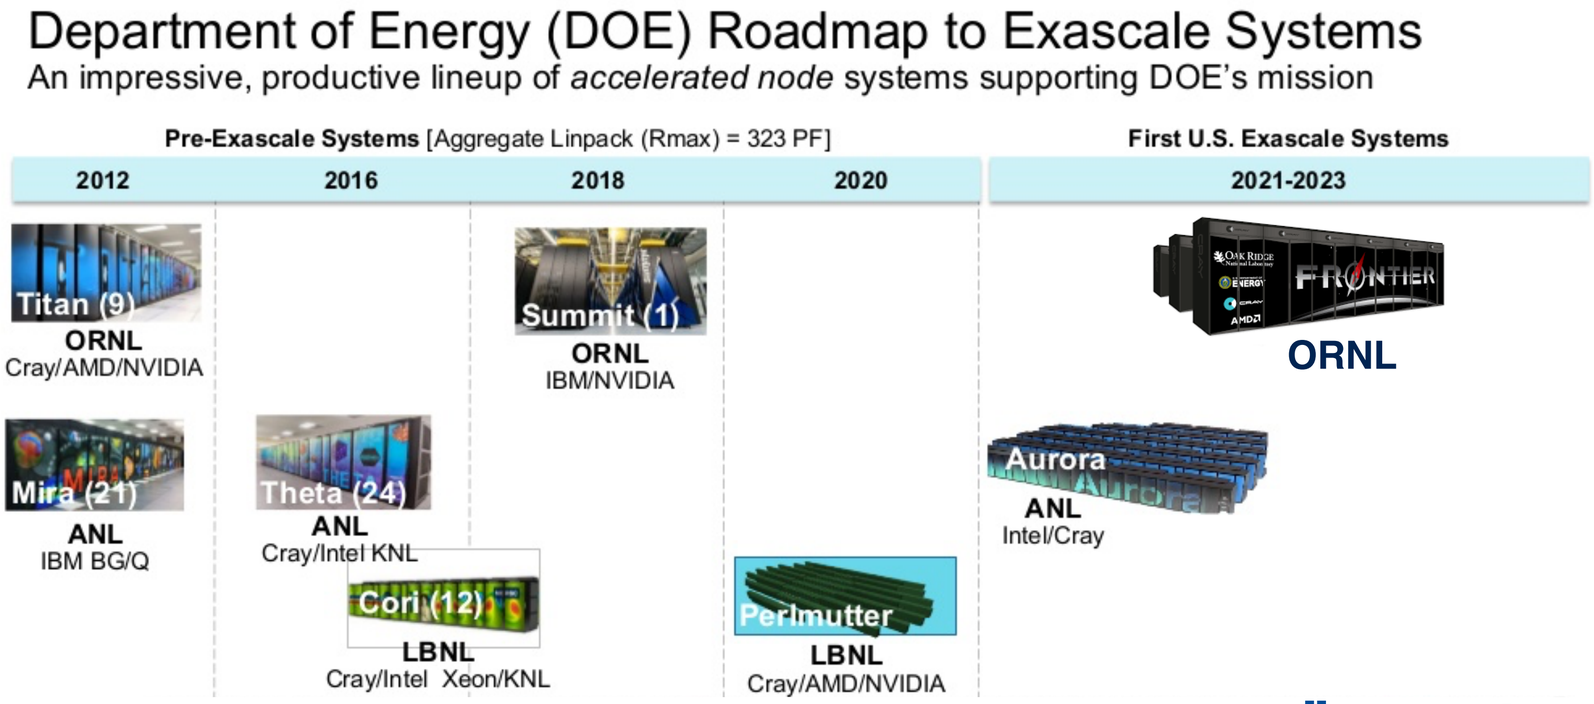

Fermilab

Evolving Computing Architectures

- Strategic direction is away from vectorization (e.g. KNL) and towards enabling acceleration on GPUs

- Efforts:
  - Exa.Trx: Algos for tracking/classification on HPC and TDAQ (FNAL, LBNL, SLAC, Caltech)
  - SciDAC: HEP Reco on Parallel Architectures
    - CMS Tracking and LAr Reconstruction
    - Working with ORNL experts on adapting algorithms to GPUs
  - SciDAC: HEP Data Analytics on HPC
    - Optimization and fitting frameworks, tuning for generators, exploring HDF5 and object stores, data parallel programming with python and HPC C++ tools
  - SciDAC: Accelerator Simulations (ComPASS 4) …

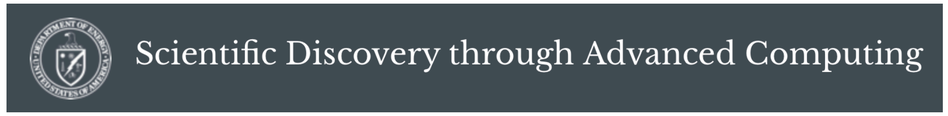

Synergia accelerator beam dynamics simulation

- Study emittance growth over 100,000 revolutions in GSI SIS18 accelerator
  - Effects of statistical noise are important
- Largest beam dynamics simulation ever

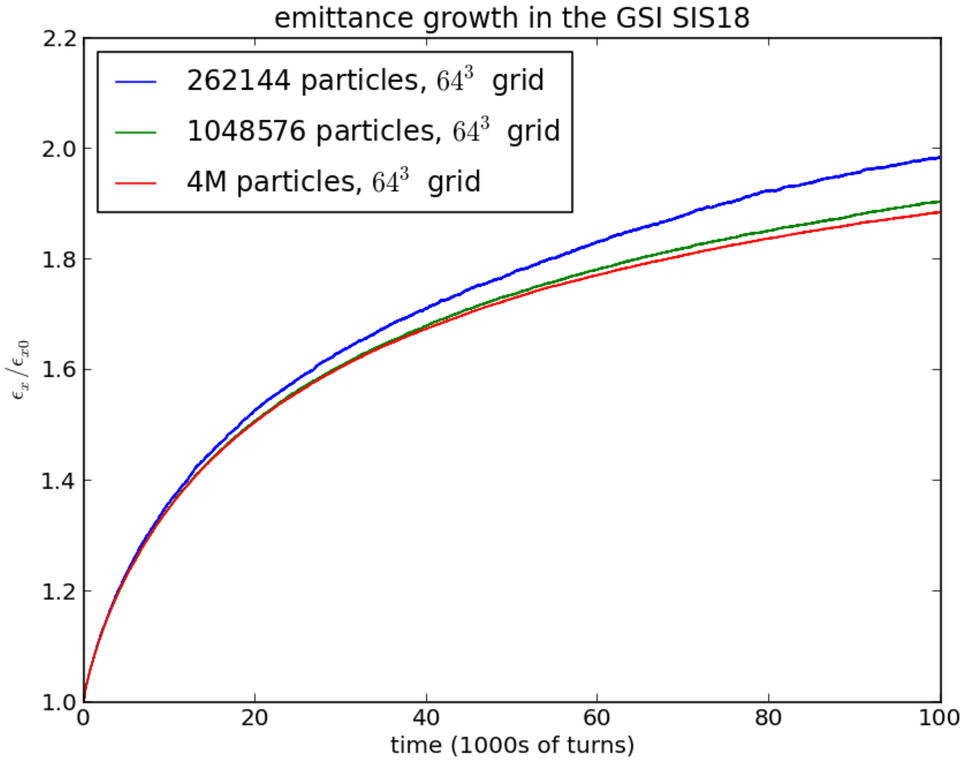

71 steps/turn
7,100,000 steps
4,194,304 particles
29,779,558,400,000 particle-steps
1,238,158,540,800,000 calls to “drift”

Yes, that's over a quadrillion

Fermilab

Using Synergia to explore GPU technologies

- Simulation on previous slide took 2 weeks on 144 cores
  - 12 12-core Intel machines
  - MPI
- Newly-developed Kokkos version
  - Synergia's drift calculation is near-perfect for Kokkos and GPU acceleration
    - Used to be half of all CPU time (other half is space charge calculation)
    - Floating point intensity is extremely high and memory utilization of L1 cache is near maximum
    - Moving Drift calculation to GPU is ~ 500x faster on one GPU than one node of CPUs
  - Working on space charge calculation
  - Expect to do same calculation on 4 GPUs in ~10 hours
  - Extraordinary!

- Synergia is a demonstration of leveraging GPUs with an HPC performance toolkit (Kokkos)
- Testing ground for HPC tools for HEP

Evolving Computing Architectures

- Center for Computing Excellence
  - Office of High Energy Physics Program with Fermilab, ANL, LBNL, and BNL
  - Proposal (see talk after this one)
  - For running applications for Dune, LHC and Cosmic on the DOE Leadership Class Facilities (Argonne and Oak Ridge)

  - Thrusts (all below for HEP on High Performance Computing)
    - Portable parallelization strategies
    - Fine-grained I/O and storage (includes data models and structures)
    - Event generators
    - Complex workflows

- CMS - offloading GPU work with framework

Compute

- HEPCloud
  - Make for common cause with OSG for heterogeneous scheduling
  - Concentration on support for leadership computing facilities starting with ALCF
  - Evolution of our facility, especially with Rucio
  - Application integration (CMS)

Storage

- Future Storage Strategy (see Bo's talk)
  - Software strategy involves Rucio
  - Rucio for CMS project - transition to Rucio for storage management
  - Rucio for DUNE and IF projects
    - Goal is to use Rucio for storage management and replace our current IF file and metadata catalogs
    - Integrate Rucio into Dune workflows
- Evolving Architecture successes involving I/O lead into the future storage strategy
- Storage technology (tapes)
  - Proposed National Consortium
  - CERN CTA

Networking

- Big Data Express
- Quantum Networking (led by FQI)

- Work with ESNet and OSG
  - Idea to blur networking, storage and compute

Analysis Facility Concepts

- How would one do analysis on an Exascale machine?

- Closer to home:
  - Deploying analysis nodes with Kubernetes (CMS)
  - Coffea (HEP analysis with python tools - originated at CMS)
  - SciDAC HEP Analytics (HEP analysis at HPC with python/C++)

Quantum Computing

- We support the Fermilab Quantum Institute
  - Two scientists and one engineer on FQI projects
    - Large Scale Quantum Device Simulation on HPC with HEP Analytics
    - Quantum Computer Controls
    - Event generation algorithms with QC

Summary

- R&D Strategy aligned with Lab and DOE
  - Heavy emphasis on AI and Evolving Computing Architectures

- Wide variety of funding under various programs
  - Internal LDRDs
  - OHEP (CompHEP, CCE)
  - OHEP & ASCR (SciDAC, Exa.Trx
  - CMS

- Leveraging external work (e.g. Kokkos)

- Leveraging internal work (e.g. Synergia)