Scientific Computing at Fermilab

Rob Roser, Head of Scientific Computing Division
Fermilab S&T Review
September 5, 2012

U.S. DEPARTMENT OF ENERGY

Fermilab

Outline

- Introduction
- Experiment Computing
- Community Tools and Collaborations
- Conclusion

INTRODUCTION

The Organization Chart in Words

- Future experiments and programs
  - Simulations (accelerator and experiment)
  - DAQ
  - Frameworks
  - Engineering (instrumentation)
- Science Programs
  - CMS
  - Cosmic Frontier
  - Experiment Support
- Scientific Computing Facilities
  - High Performance Parallel Computing (Lattice Support)
  - Data Movement and Storage
  - Grid and Cloud Computing
  - Facilities Administration

Our “Mission”

- Provide computing, software tools and expertise to all parts of the Fermilab scientific program including theory simulations (e.g. Lattice QCD and Cosmology), and accelerator modeling
- Work closely with each scientific program – as collaborators (where a scientist/staff from SCD is involved) and as valued customers.
- Create a coherent Scientific Computing program from the many parts and many funding sources – encouraging sharing of facilities, common approaches and re-use of software wherever possible
- Work closely with CCD as part of an overall coherent program

A Few points

- We are ~160 strong made up of almost entirely technically trained staff
- 26 Scientists in the Division
- As the lab changes its mission, scientific computing is having to adapt to this new and more challenging landscape.
- Scientific Computing is a very “matrixed” organization. While I will concentrate on those aspects associated with the fermilab facility, I will touch on other areas in order to give a more complete picture on how it all fits together

Fermilab Computing Facilities

- Feynman Computing Center (FCC)
  - High availability services – e.g. core network, email, etc.
  - Tape Robotic Storage (3 10000 slot libraries)
  - UPS & Standby Power Generation
  - ARRA project: upgrade cooling and add HA computing room - completed

- Grid Computing Center (GCC)
  - High Density Computational Computing
  - CMS, RUNII, Grid Farm batch worker nodes
  - Lattice HPC nodes
  - Tape Robotic Storage (4 10000 slot libraries)
  - UPS & taps for portable generators

- Lattice Computing Center (LCC)
  - High Performance Computing (HPC)
  - Accelerator Simulation, Cosmology nodes
  - No UPS

Facilities: more than just space power and cooling – continuous planning

FermiGrid and other central facilities are used across the lab for accelerator design, experiment design and simulation, cavity design, theoretical calculations, etc. and that these are supported as part of the SCD mission.

EXPERIMENT COMPUTING STRATEGIES

Run II Computing Strategy

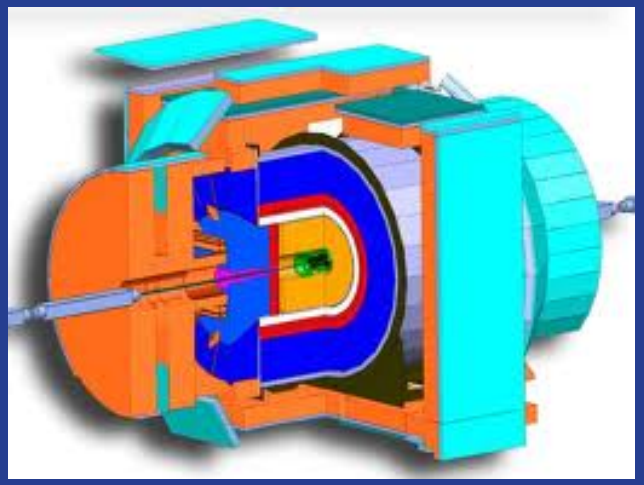

- With the cessation of beam operations in Sept 2011 – we are in the end game of Run II analysis
  - Data Reconstruction of entire data sample is complete
  - Data re-processing is also now complete
  - Some Monte Carlo production is still going on
- Analysis computing capability for at least 5 years, but diminishing demand after end of 2012
- Continued support for up to 5 years for
  - Code management and science software infrastructure
  - Data handling for production (+MC) and Analysis Operations
- Curation of the data: > 10 years with possibly some support for continuing analyses (more on this)

Tevatron Publications – A measure of Success

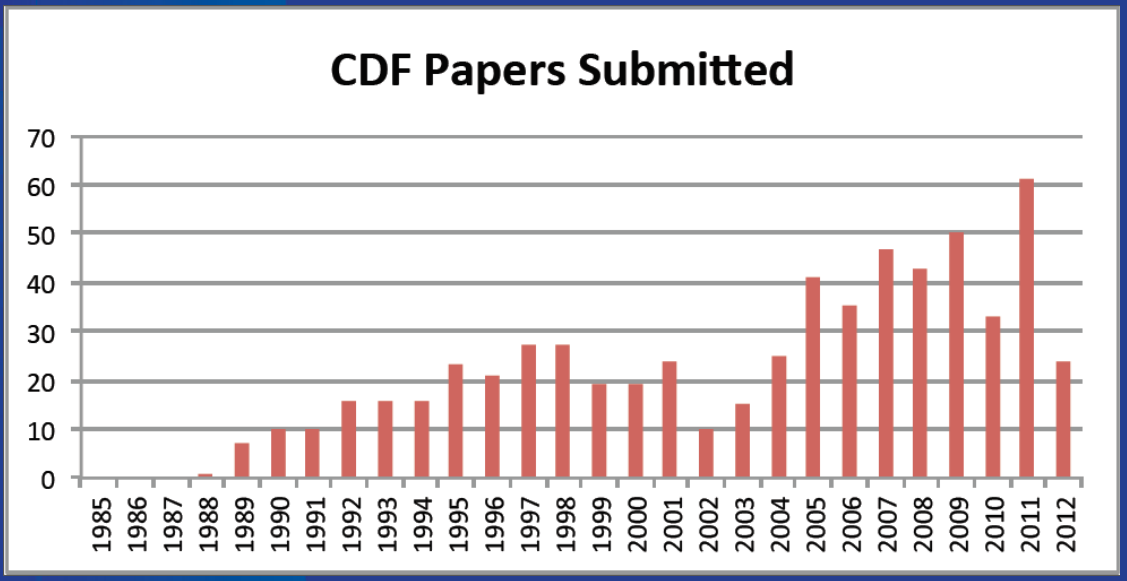

CDF has 603 papers to date

Goal of 50+ papers in 2012

D0 has ~300 papers to date

Goal of 45 papers in 2012

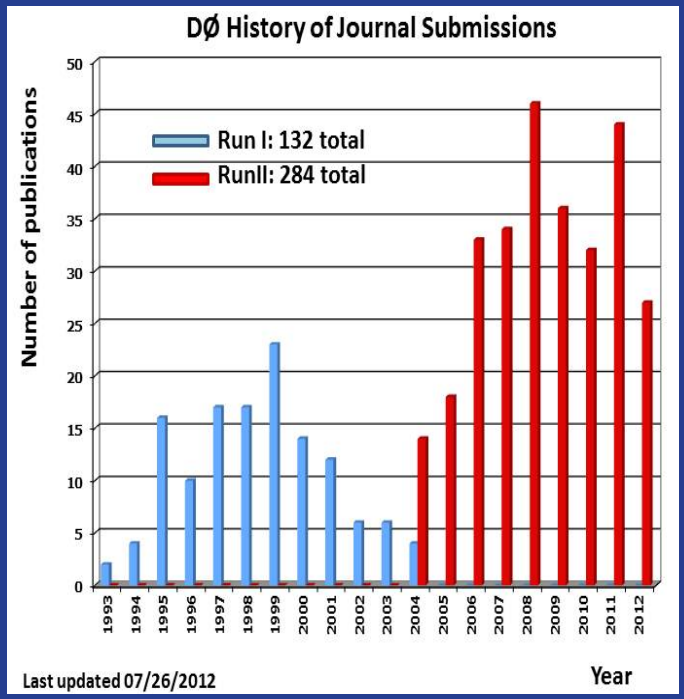

Final Tevatron Higgs Result – A measure of Computing Success!

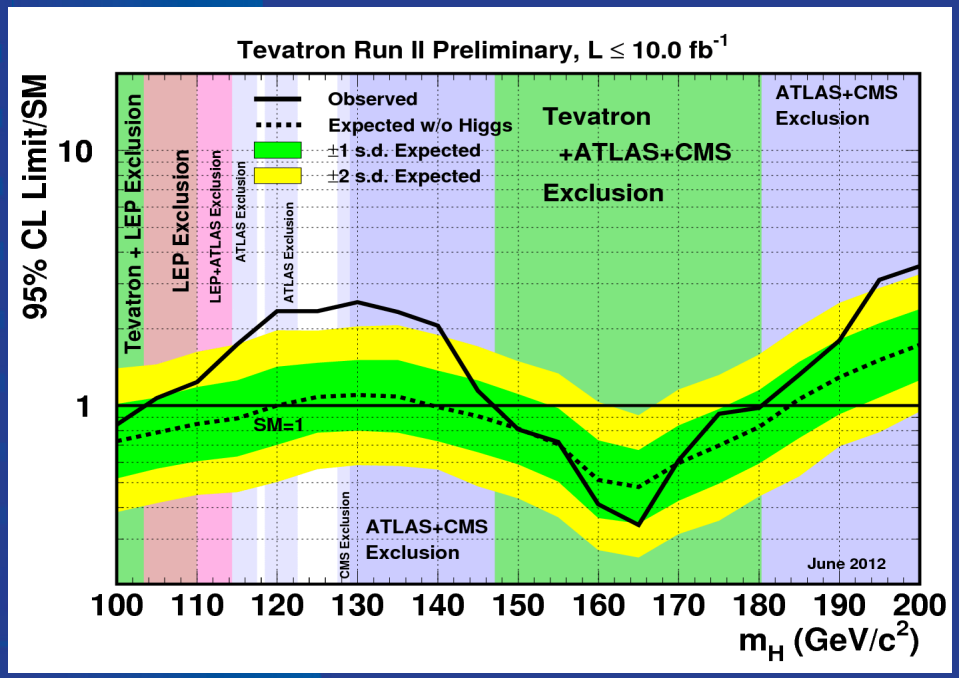

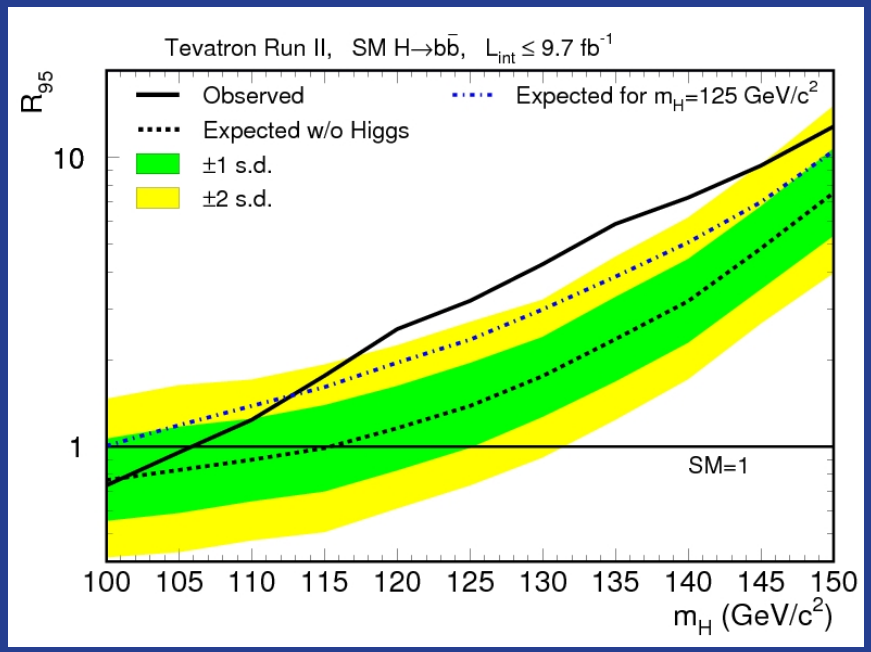

- Tevatron data are incompatible with background only hypothesis
  - For full combination of searches p-values are 3.0s local or 2.5s with LEE factor
  - For Higgs to bb channel p-values are 3.3s local or 3.1s with LEE factor
    - *evidence paper accepted by PRL in 24 hours*

"Data (Knowledge) Preservation" for Tevatron

- Members of international DPHEP organization and paying attention to what has been done world-wide
- Job is much larger than just archiving
- Archival Process
  - Data will be stored and migrated to new tape technologies for ~ 10 years. That has started – using T10k tapes
  - Eventually the 20 PB of data will seem modest
- Short term, we will maintain the ability to reprocess and do analysis on the data in the current environment
  - Code, access to databases, libraries, I/O routines, Operating Systems, documentation…..
- We need to move away from these current custom systems to being able to run on generic future platforms

Knowledge Preservation - 2

- Maintaining full capability will be a combined and coordinated effort between SCD and experiments
- Successful FWP funded to hire TWO domain knowledgeable scientists to lead the effort on each experiment. Offers to go out in < month
- Knowledge Preservation
  - Need to plan and execute the following…
    - Preserving analysis notes, electronic logs etc
    - Document how to do analysis well
    - Document sample analyses as cross checks
    - Understand job submission, db, and data handling issues
    - Investigate/pursue virtualization
- Try to keep CDF/D0 strategy in synch and leverage common resources/solutions

Knowledge Preservation - 3

- Minimal effort being made to make thie archived “publically” analyzable
- Experiments will each address their unique issues re data preservation in their presentations later today
- Both have assembled teams/task forces to establish plans/requirements, identify the people and determine milestones etc.

Intensity Frontier Program (Diverse)

- Many experiments in many different phases of development/operations.
  - MINOS
  - MiniBooNE
  - SciBooNE
  - MINERvA
  - NOvA
  - MicroBooNE
  - ArgoNeuT
  - Seaquest
  - Mu2e
  - g-2
  - LBNE
  - Project X era expts

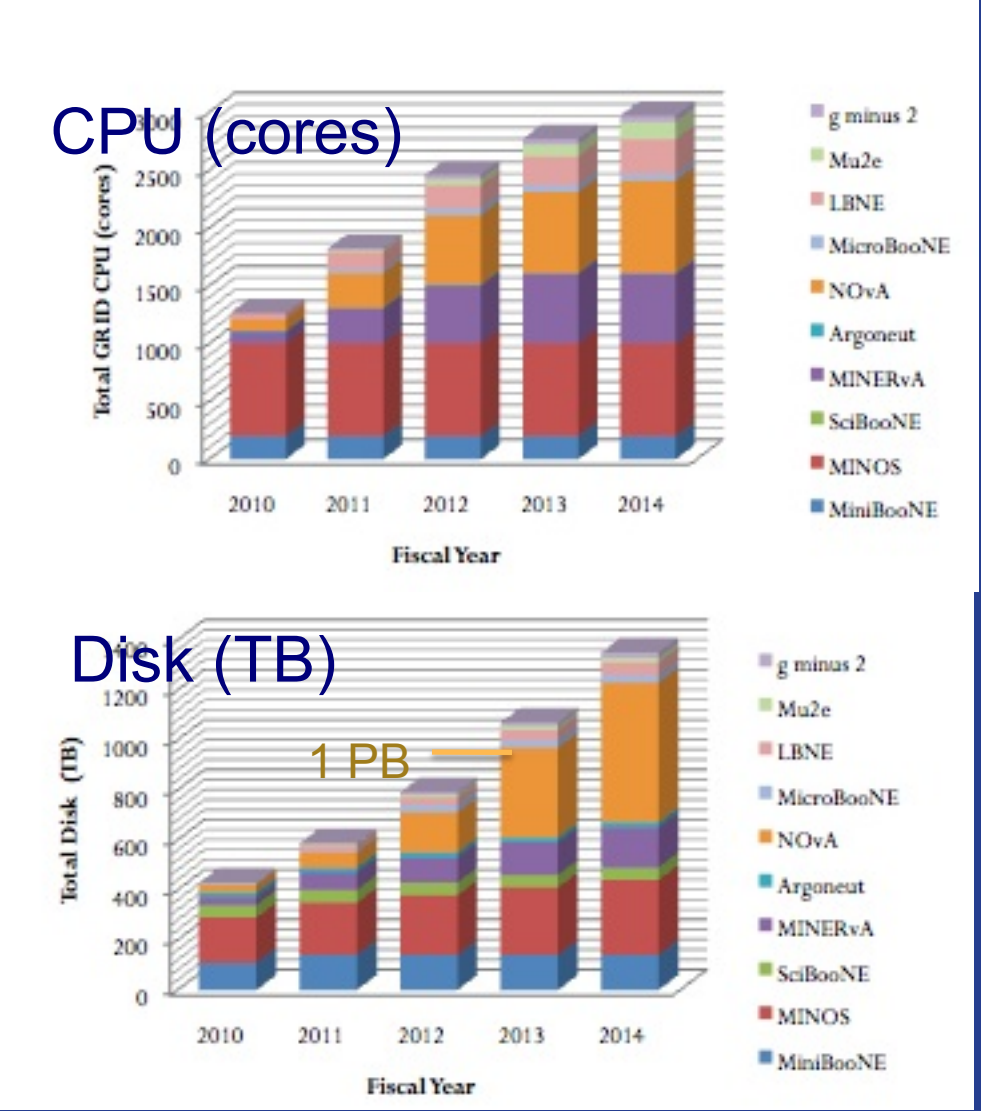

Intensity Frontier Strategy

- Common approaches/solutions are essential to support this broad range of experiments with limited SCD staff.
- SCD is establishing a liaison between it and experiment to insure communication and understand needs/requirements
- In the process of establishing MOU's between SCD and experiment to clarify our roles/ responsibilities

Intensity Frontier Strategy - 2

- A shared analysis facility where we can quickly and flexibly allocate computing to experiments
- Continue to work to “grid enable” the simulation and processing software
  - Good success with MINOS, MINERvA and Mu2e
- All experiments use shared storage services – for data and local disk – so we can allocate resources when needed
- Perception that intensity frontier will not be computing intensive which is wrong

CMS Tier 1 at Fermilab

- The CMS Tier-1 facility at Fermilab and the experienced team who operate it enable CMS to reprocess data quickly and to distribute the data reliably to the user community around the world.

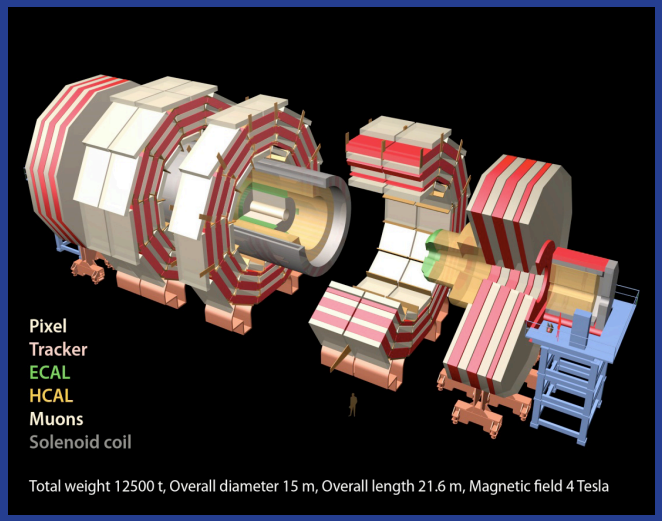

Fermilab also operates:

- LHC Physics Center (LPC)
- Remote Operations Center
- U.S. CMS Analysis Facility

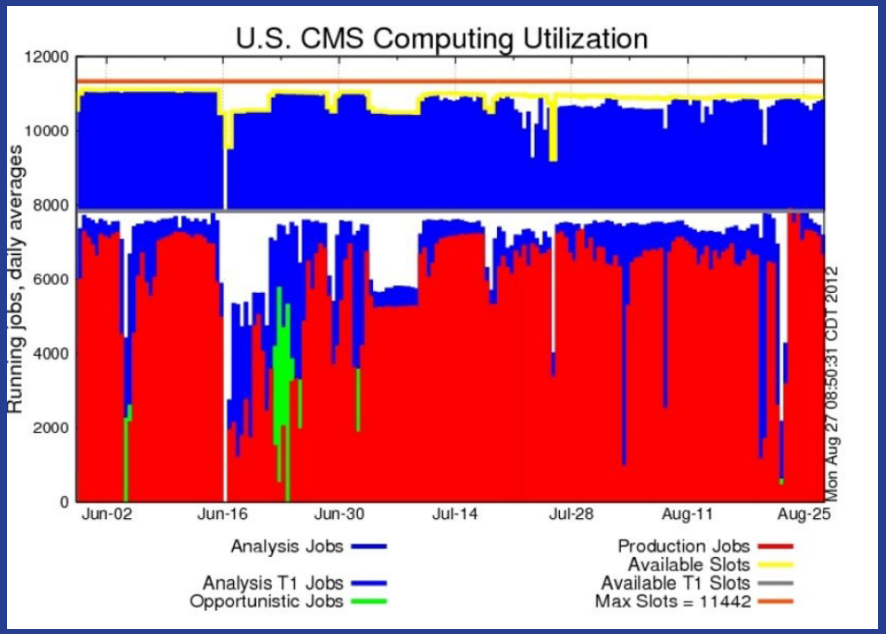

Fermilab

CMS Offline and Computing

- Fermilab is a hub for CMS Offline and Computing
  - Ian Fisk is the CMS Computing Coordinator
  - Liz Sexton-Kennedy is Offline Computing Coordinator
  - Lothar Bauerdick is head of US CMS Software and Computing
  - Leadership roles in many areas in CMS Offline and Computing: Frameworks, Simulations, Data Quality Monitoring, Workload Management and Data Management, Data Operations, Integration and User Support.
- Fermilab Remote Operations Center allows US physicists to participate in monitoring shifts for CMS.

Computing Strategy for CMS

- Continue to evolve the CMS Tier 1 center at Fermilab - to meet US obligations to CMS and provide the highest level of availability and functionality for the dollars spent
- CMS Tier-1 serves data to the 7 US CMS Tier-2s as part of its mission
- Continue to ensure that the LHC Physics Center and the US CMS physics community is well supported by the Tier 3 (LPC CAF) at Fermilab
- Plan for evolution of the computing, software and data access models as the experiment matures – requires R&D and development
  - Ever higher bandwidth networks
  - "AAA" Any data, anytime anywhere
  - Frameworks for multi-core

Cosmic Frontier

- Continue to curate data for SDSS
- Support data and processing for Auger, CDMS and COUPP
- Will maintain an archive copy of the DES data and provide modest analysis facilities for Fermilab DES scientists.
  - Data management is an NCSA (NSF) responsibility
  - Helping NCSA by “wrappering” science codes needed for 2nd light when NCSA completes its framework.
  - Working to understand the Science Computing needs of DES and are able to host what is required
- DES use Open Science Grid resources opportunistically
- Positioning now for a role on LSST and Darkside 50

COMMUNITY TOOLS AND COLLABORATIONS

Open Science Grid (OSG)

Open Science Grid

- The Open Science Grid (OSG) advances science through open distributed computing. The OSG is a multi-disciplinary partnership to federate local, regional, community and national cyberinfrastructures to meet the needs of research and academic communities at all scales.

The US contribution and partnership with the LHC Computing Grid is provided through OSG for CMS and ATLAS

- Total of 95 sites; ½ million jobs a day, 1 million CPU hours/ day; 1 million files transferred/day.
- It is cost effective, it promotes collaboration, it is working!

Geant4

Workshop Held

- between HEP and ASCR
- Discussed how to transform GEANT4 to run efficiently on modern multi-core computers
- Workshop chairs were Robert Lucas (USC) and RR.

Locally: Agorithmic development to be able to utilize multi-core architectures and are porting G4 sections to the GPUs)

Transforming GEANT4 for the Future

Report from the Workshop on Transforming GEANT4 for the Future,
Rockville Maryland, USA
May 8-9, 2012
DRAFT

Robert Lucas and Robert Roser, Editors and Workshop Chairs

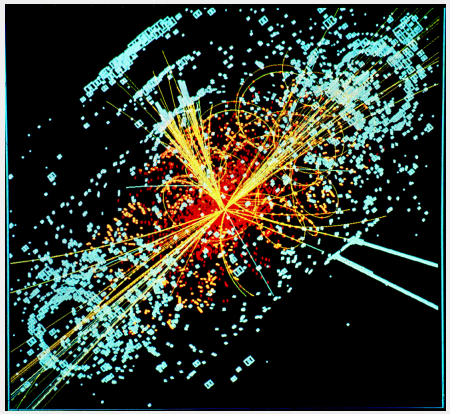

High Performance (parallel) Computing

Needed for

- Lattice Gauge Theory calculations (LQCD)
- Accelerator modeling tools and simulations
- Computational Cosmology:

Dark energy, matter | Cosmic gas | Galaxies

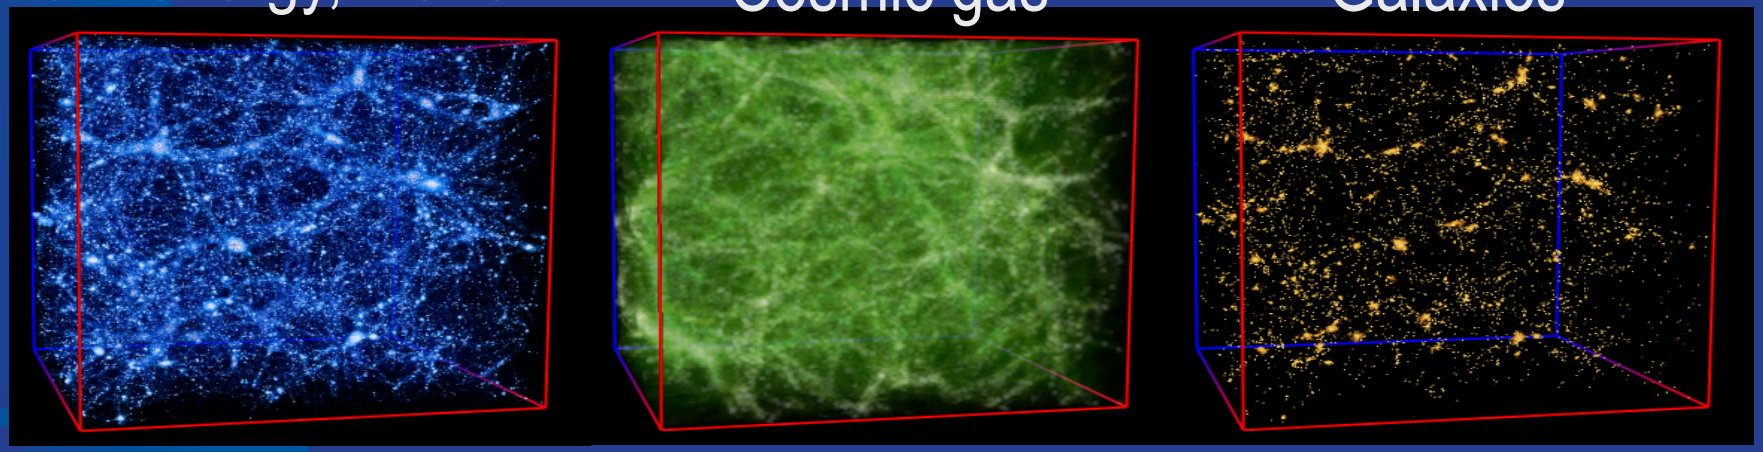

Simulations connect fundamentals with observables

Strategies for Simulation Science Computing

- Lattice QCD is the poster child
  - Coherent inclusive US QCD collaboration (HEP and Nuclear)
    - Fermilab Scientist leads. This allocates HPC resources.
  - LQCD Computing Project (HEP and NP funding)
    Fermilabmember holds the role as is the Project Manager
  - SciDAC II project to develop the software infrastructure
- Accelerator modeling
  - Multi-institutional tools project COMPASS – Fermilab Scientist is the PI
  - Also accelerator project specific modeling efforts
- Computational Cosmology
  - Partnering wth Argonne on Galaxy Simulation Project

Accelerator modeling tools at Fermilab

- Fermilab is leading the ComPASS Project:
  - Community Petascale Project for Accelerator Science and Simulation — a multi-institutional collaboration of computational accelerator physicists
  - Fermilab scientist is PI for ComPASS
  - ComPASS goals are to develop High Performance Computing (HPC) accelerator modeling tools
    - Multi-physics, multi-scale for beam dynamics; "*virtual accelerator*"
    - Thermal, mechanical, and electromagnetic; "*virtual prototyping*"
- Development and support of CHEF
  - general framework for single particle dynamics developed at Fermilab

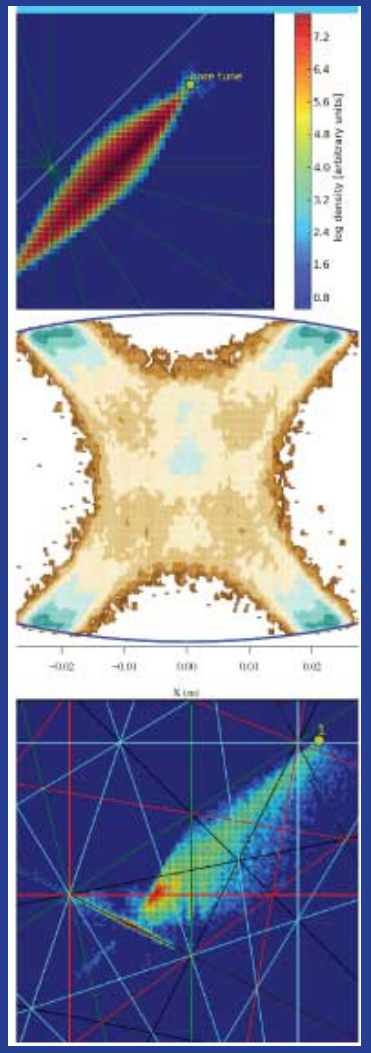

Framework development

- Modern HEP experiments need sophisticated software infrastructure to enable analysis of massive data sets by 100s of collaborators
  - The “Framework”: data model, persistency, configuration, geometry and conditions data management, interaction with detector simulation, build and release management, work-flow management
- We have a long (successful!) history developing tools that can be adapted for many different scientific applications
  - RunII, CMS, and now Intensity Frontier and future cosmic frontier activities
- Heavily leveraged activity
  - Work in close partnership with users

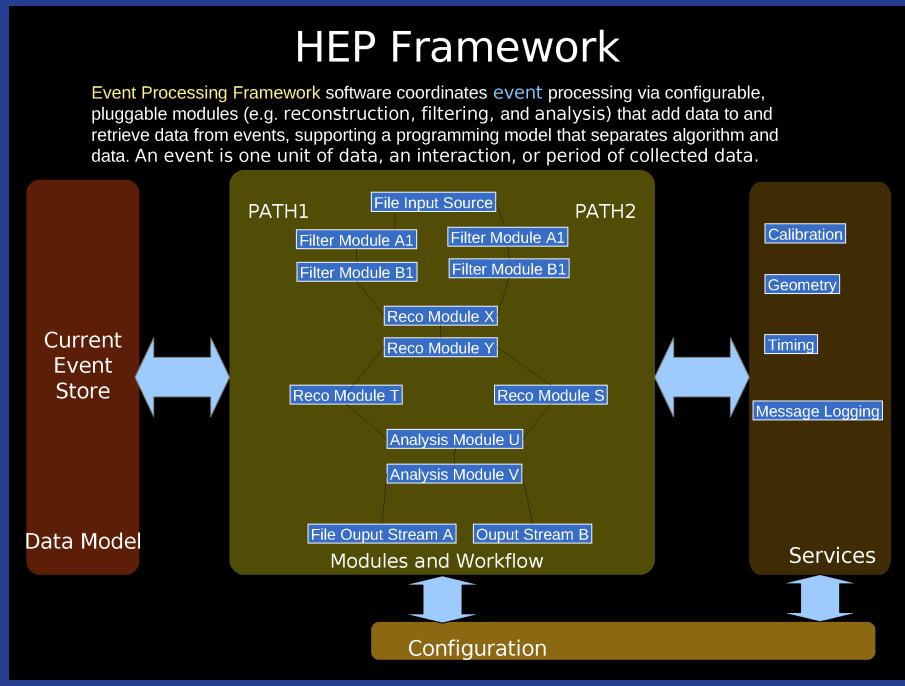

Framework Applications

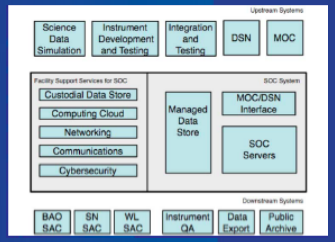

LAr

NOvA

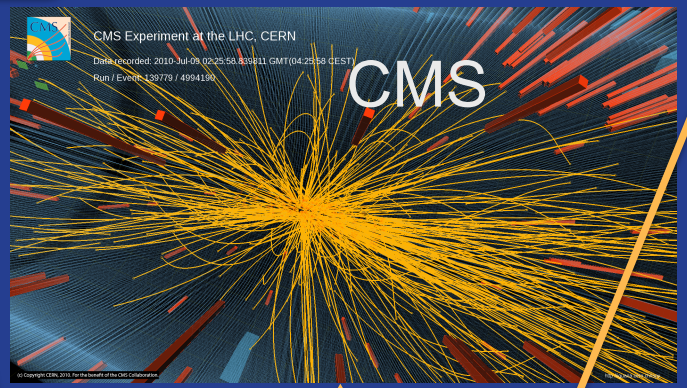

Mu2e

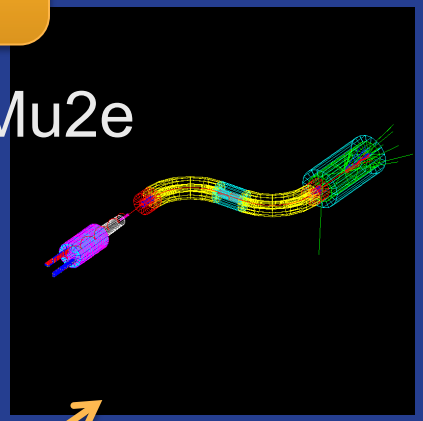

LQCD software

Framework

RunII Offline infrastructure

MiniBooNE

- **Success: specific application (RunII) leads to community tool and continuing requests for framework applications from new projects**
- **Success: high-quality implementations (most recently, CMS framework)**

Frameworks and the Intensity Frontier Experiments

- Buy-in from the Intensity Frontier Experiments exists
- This Framework has the advantage that Universities can load just what they need on their home clusters to do analysis
- Input from each experiment will insure solution that meets their demands/wishes (regular stakeholder meetings)

Strategy

In these tough economic times, trying to bring in additional funding based on projects that are well aligned to our core expertise.

- GEANT4 HEP/ASCR initiate – expect workshop will turn into a "call" this fall. Will look to partner with SLAC on this.
- Tevatron Data Preservation – money from energy frontier and computational high energy physics for two 2-year term appointments (~750k over 2 years)
- Partnering with Argonne (Salman Habib, PI) on Galaxy Simulation project (~200k for one year)
- Awarded 3 grants from U of Chicago for joint Argonne/FNAL/ Chicago partnerships
- Plan to leverage these successful collaboraitons into future joint initiatives

SciDAC 3 Awards

- Office of Science (SC) Program - Third Cycle
- SciDAC (Scientific **Discovery through Advanced Computing)**
- Joint HEP-ASCR Funding Opportunity Announcement Posted Sept 2011, Closed Jan 2012

  Research to advance the HEP mission by fully exploiting DOE SC leadership class computing resources.

  Projects funded in the areas of:

  Cosmic Frontier Scientific Simulations,

  Lattice Gauge Theory Research, and

  Accelerator Science Modeling and Simulation

SCIDAC Awards to FNAL

SciDAC 3 awards						
ComPASS	225	90	315	945	Spentzouris	Other funds are from ASCR and go to PI. Needs to be called Compass 2 or 3.
Lattice Gauge Theories	392	0	392	1176	Mac Kenzie	
Dark Universe	155	0	155	465	Dodelson	
total SciDAC 3	772	90	862	2586		KA140103 all funding this page
SciDAC budget code but NOT sciDAC 3 project						
OSG	1265	25	1290		Pordes/Lothar	OSG is funded for 5 years with co funding from NSF to univs, review in Fy14. Other refers to NP here.
PDACS						Partners with Habib's NERSC pilot project and ANL hardware.

Software – collaborative efforts

- ComPASS – Accelerator Modeling Tools project
- Lattice QCD project and USQCD Collaboration
- Open Science Grid – many aspects and some sub-projects such as Grid security, workload management
- Grid and Data Management tools
- Advanced Wide Area Network projects
- Dcache collaboration
- Enstore collaboration
- Scientific Linux (with CERN)
- GEANT core development /validation (with GEANT4 collaboration)
- ROOT development & support (with CERN)
- Cosmological Computing
- Data Preservation initiative (global HEP)

Conclusion

- We have a coherent and evolving scientific computing program that emphasizes sharing of resources, re-use of code and tools, and requirements planning.
- Embedded scientists with deep involvement are also a key strategy for success.
- Fermilab takes on leadership roles in computing in many areas.
- Greater emphasis on common tools as we transiiton to the Intensity Frontier

BACK-UP