Report of Topical Panel on HEP Computing

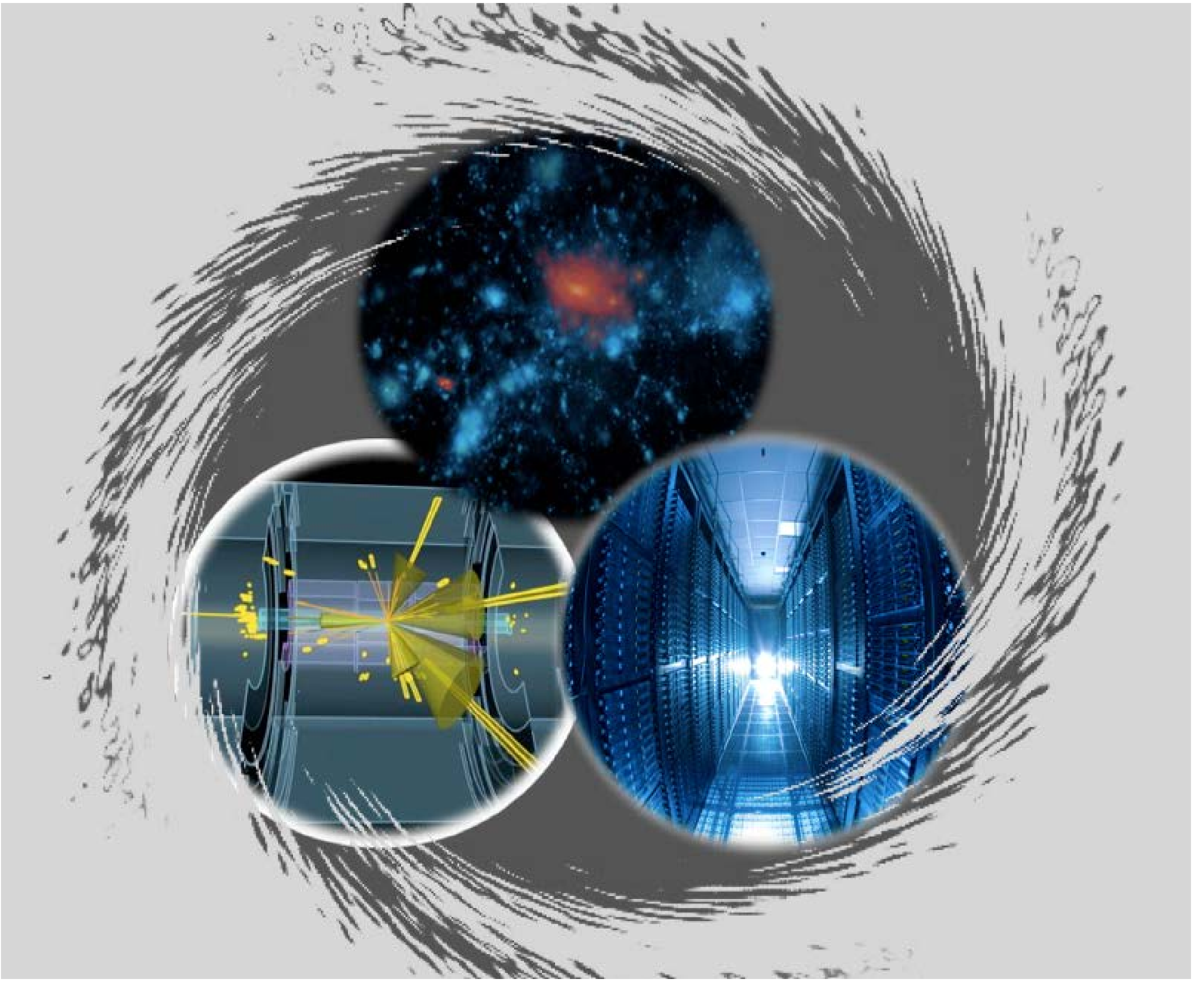

HEPAP Meeting
Bethesda, MD
Mar. 13, 2014

Paul Avery
University of Florida
avery@phys.ufl.edu

Salman Habib
Argonne Nat. Lab
habib@anl.gov

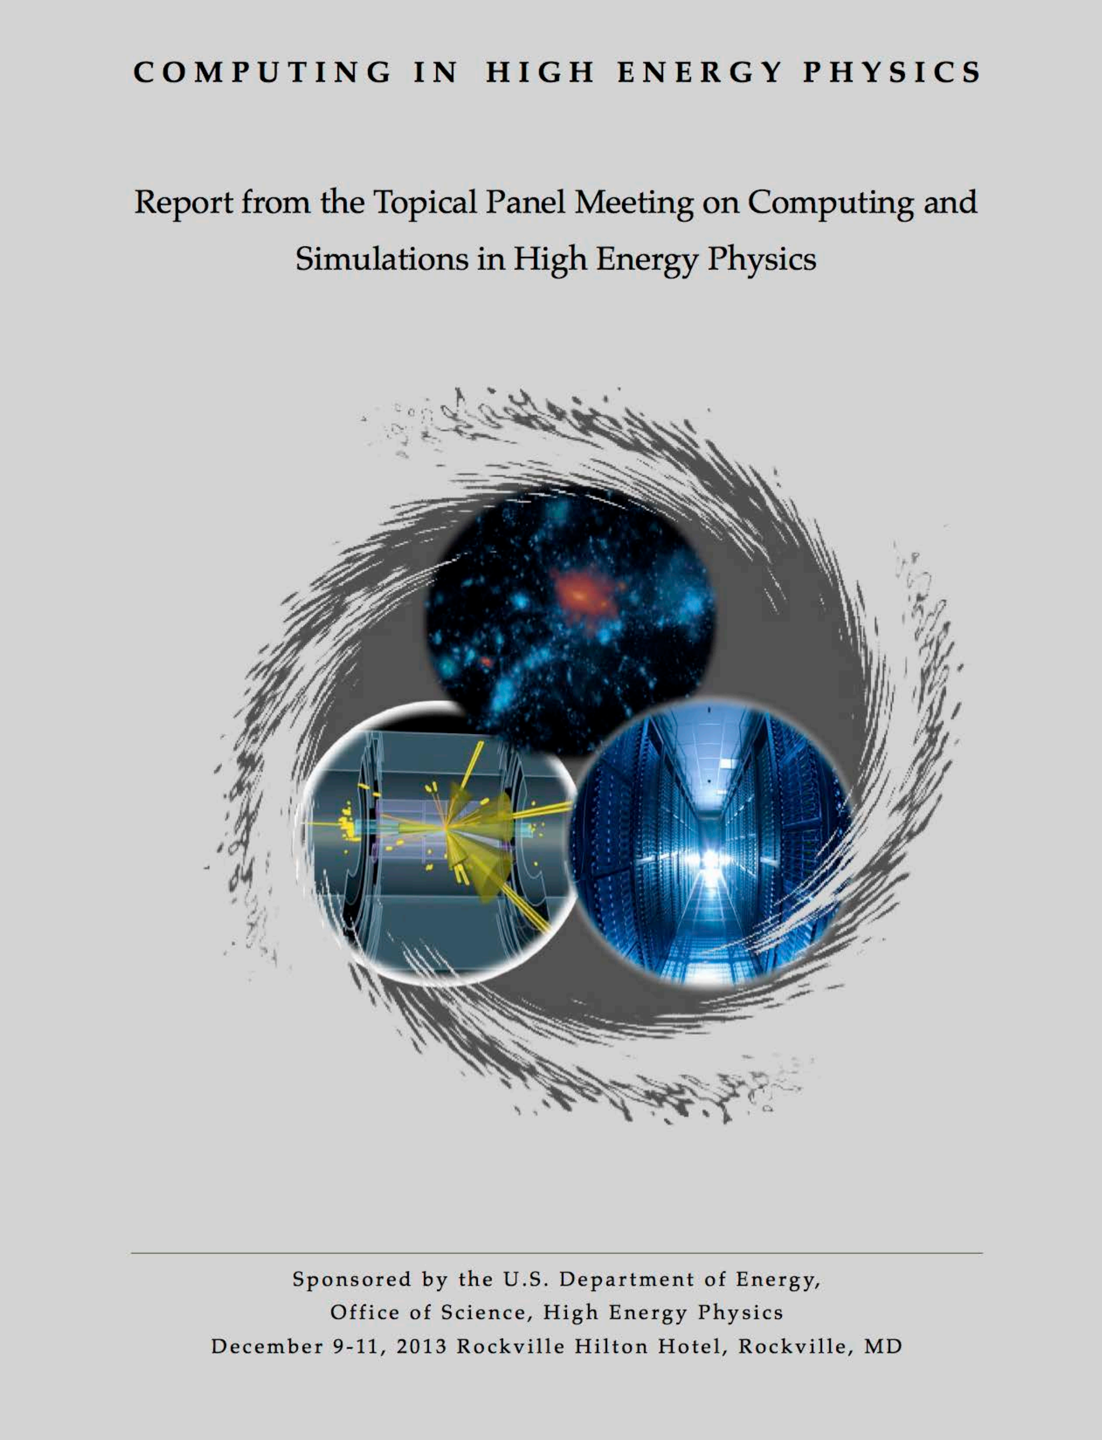
COMPUTING IN HIGH ENERGY PHYSICS

Report from the Topical Panel Meeting on Computing and Simulations in High Energy Physics

Sponsored by the U.S. Department of Energy,
Office of Science, High Energy Physics
December 9-11, 2013 Rockville Hilton Hotel, Rockville, MD

Topical Panel Meeting

Dec. 9-11, 2013
Washington, DC

Organization of Topical Panel

Steering Committee (8 members)

Paul Avery (co-Chair)	U Florida
Salman Habib (co-Chair)	Argonne
Amber Boehnlein	SLAC
Robert Roser	Fermilab
Stephen Sharpe	U Washington
Heidi Schellman	Northwestern
Craig Tull	LBNL
Torre Wenaus	BNL

30 Panel Members in 5 Areas

1. Cosmic Frontier
2. Intensity Frontier
3. Energy Frontier
4. Accelerators
5. Technology

Topical Panel Members

Cosmic Frontier
Anders Borgland (SLAC)
Andy Connolly (U of Washington)
Gus Evrard (U Michigan)
Cristiano Galbiati (Princeton)
Peter Nugent (LBNL)
Martin White (Berkeley)

Energy Frontier
Lothar Bauerdick (FNAL)
Chip Brock (Michigan State)
Kaushik De (UT Arlington)
Steven Gottlieb (Indiana)
Stefan Hoeche (SLAC)
Tom LeCompte (Argonne)
Dan Marlow (Princeton)
Harvey Newman (Caltech)

Technology
Paul Messina (ANL)
Inder Monga (ESnet)
Richard Mount (SLAC)
Don Petravick (NCSA)
David Skinner (NERSC)
Artur Barczyk (Caltech)

Intensity Frontier
Bob Bernstein (FNAL)
Milind Diwan (Brookhaven)
Paul Mackenzie (FNAL)
Mayly Sanchez (Iowa State)
Malachi Schram (PNNL)

Accelerators
Cho Ng (SLAC)
Robert Ryne (LBNL)
Tor Raubenheimer (SLAC/Stanford)
Panagiotis Spentzouris (FNAL)
Jean-Luc Vay (LBNL)

Topical Panel Timeline

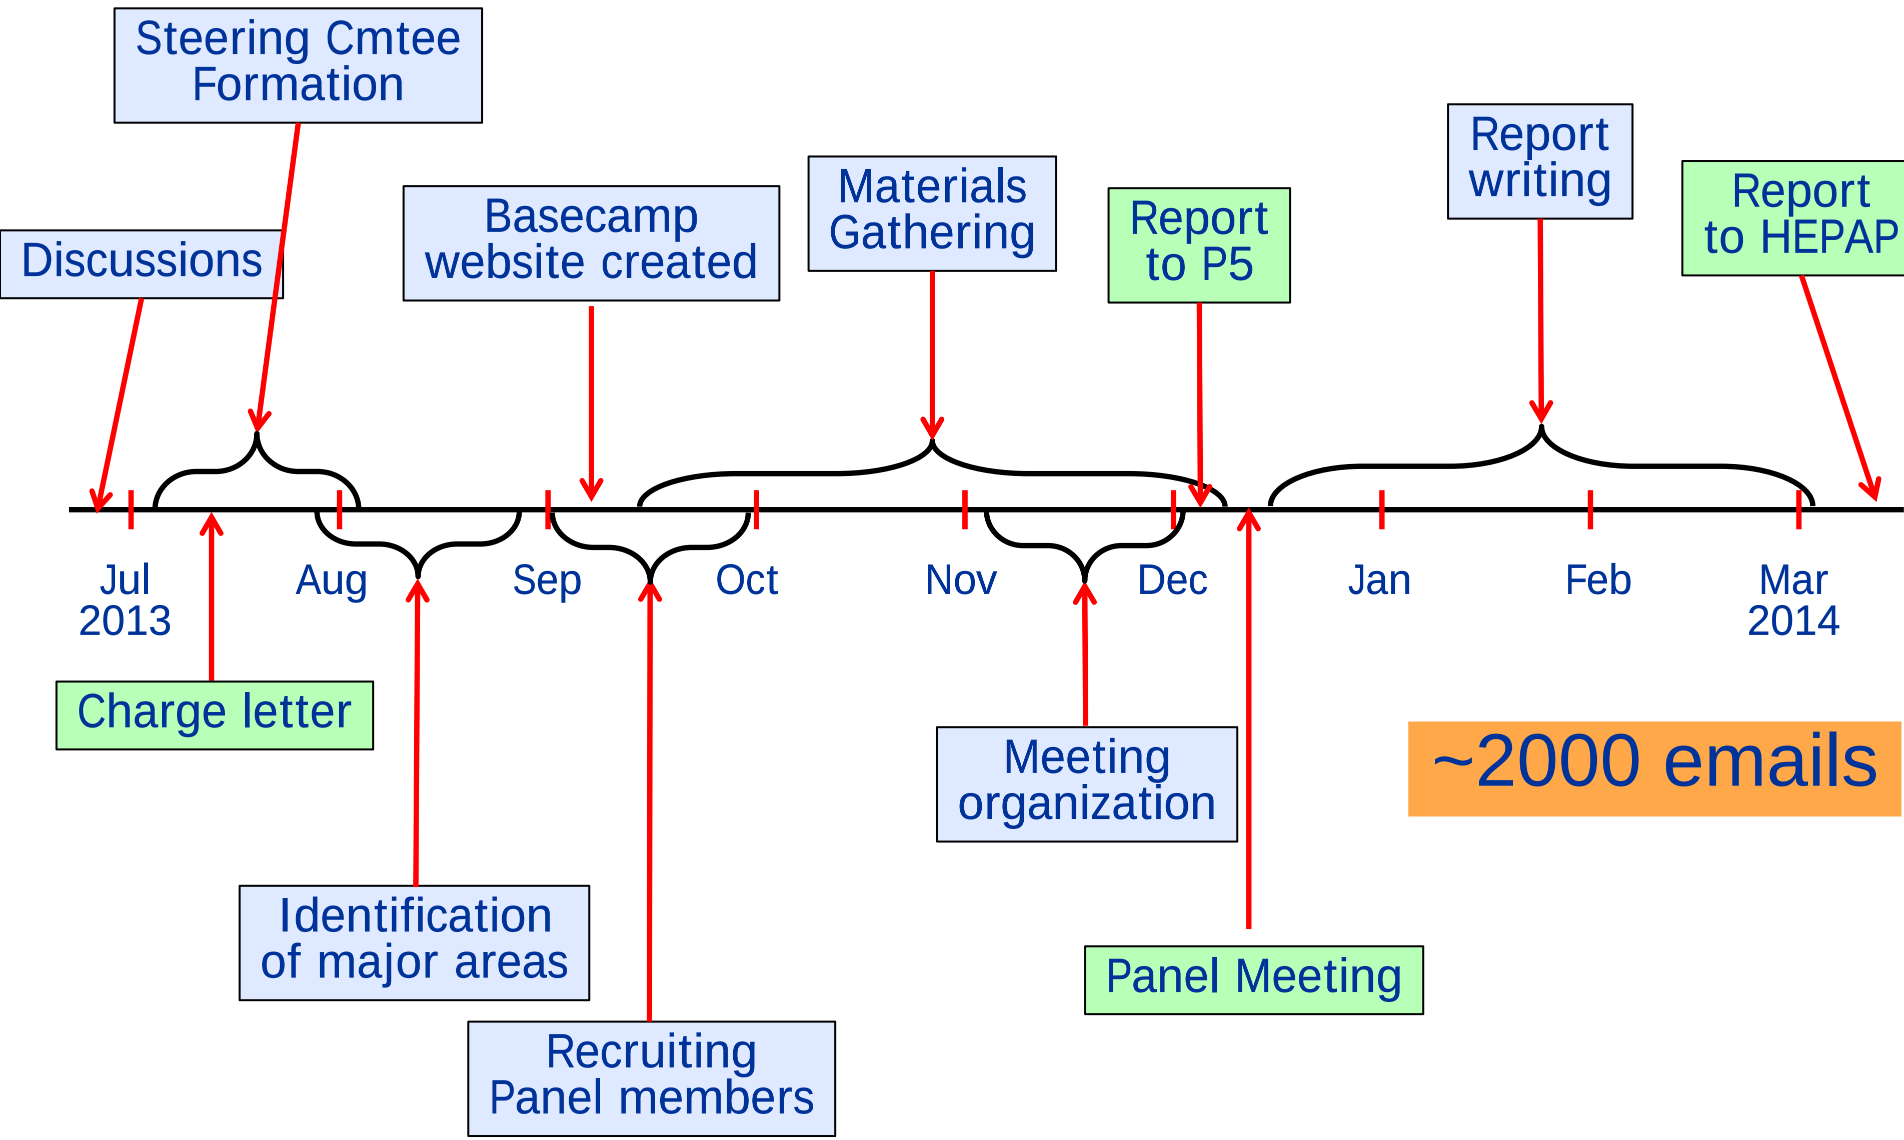

Topical Panel Charge (J. Siegrist)

- Key elements of charge letter
  - Identify cross-cuts across the HEP computing program that can benefit from common solutions
  - Identify opportunities for R&D with high programmatic impact, including international leadership
  - Survey HEP software and identify actions related to maintenance/updating, gaps, non-HEP partnerships, and lifecycle management
  - Survey current computing and data management practice across HEP: Can an improved structure accelerate progress?
  - Survey use of hardware, identify opportunities for increased efficiency, cost effectiveness, & application of the best technologies
  - Identify opportunities presented by establishing a (virtual/distributed) Center for HEP Computing Excellence

Some Observations

- Computing plays a fundamental role in HEP
  - While computing now underlies ALL of science, its role in HEP is pervasive and unique
  - Computing represents a significant fraction of experiment costs
- HEP computing is *highly* diverse across the field
  - The Dec. 9-11 meeting and materials gathering made this clear
  - Major variations between Cosmic, Intensity, Energy, Accel., LQCD
  - HPC vs HTC, use of storage, networks, …
- HEP computing has been influential outside the field
  - Leadership/innovation in scale, data handling, global operations, networking, Grid/distributed computing
  - Creating/operating Open Science Grid

Topical Panel Drivers

- Awareness of large investments: people & infrastructure
  - Both DOE + NSF
- Desire to increase effectiveness of computing investments
  - More ideas/best practices exchanged within/across HEP domains
  - “Global” vs. “local” optimizations
- Need for increased responsiveness to near-future drivers
  - Science focus changes
  - Changes in commodity computing and storage hardware
  - Era of limited funding
- Positioning HEP computing in the wider computing world
  - How to best leverage HEP expertise
  - How to best interact with other science domains and industry

Key Panel Activities

- Gather & organize information
  - Specific data & overall view of activities and infrastructure
  - Dozens of 1-pagers each providing brief summary of an area
  - XrootD, Root, Pythia, bbcp, Geant4, LQCD, OSG, data catalogs, workload management, HACC, CMB simulations…

(See examples)

- Use Basecamp project software (web-based)
  - Place to upload documents/reports, organize information
  - Communications (discussions, broadcasts), daily summaries
- Suggest possible ways to move forward
  - improve overall effectiveness of computing-based activities
- Create & maintain website with documents, links, etc.
  - Location TBA

1-Pager example: Energy Frontier

Heidi Schellman

Pythia

- Pythia is a **primary** tool for interpreting data at the Energy Frontier
  - Generates complete events of proton-proton collisions up to the detector simulation-level
  - Originated in 1980's from attempt to understand PETRA data
  - User manual was the most-highly cited HEP publication of 2012
  - Successful because of its responsiveness to experimental and theoretical developments

- Current State
  - "Best" models of soft physics that underlies hard events
  - Framework is highly adaptive to advances in Standard Model and Exotic theory
  - The primary event generator of all HEP experiments, but also used widely by CF, IF, and theorists
  - 1 FTE in US, 3-4 FTE in Europe
  - Developers at Lund, CERN, and Fermilab

Jean-Luc Vay

1-Pager example: Accelerators

Modeling of laser plasma acceleration

- Laser-plasma accelerators offer very high gradient for shorter future linacs
  - EM plasmas, gas dynamics, radiation/scattering
  - Very large space and time scales, e. g. l_{laser} (mm) << l_{accel} (meters)
    - full scale is needed for correct physics (emittance, focusing, …)
    - scaled or reduced models are used for faster turnaround still require HPC
  - Scaling to 100's kcore: EM plasmas modeled with parallel Particle-In-Cell/fluid codes
    - OSIRIS, VORPAL, WARP, ALaDyn, VLPL, INF&RNO, QuickPIC, REMP, …
- Current State
  - SciDAC ComPASS supported codes run at scale (OSIRIS, VORPAL, WARP, QuickPIC)
  - HEP requires and has implemented new algorithms to address accelerator physics
    - boosted frame, quasistatic, improved field and laser envelope solvers, noise mitigation, …
  - Support cutting edge experiments (BELLA, FACET, …) that produce GeV beams with low emittance and energy spread & develop path to future colliders
  - Codes developed and supported by ~5 FTEs in the US, supported by HEP
    - complementary development of general PIC by other offices (e.g. WARP, OSIRIS, VORPAL by fusion)

1-Pager example: Technology

Artur Barczyk
Harvey Newman

LHCONE
(LHC Open Network Environment)

- Started in 2010 to address/avert intercontinental connectivity bottleneck for Tier2 related LHC traffic
  - maintain high quality services for the LHC community
  - protect R&E network infrastructures against effects of large data flows
- Today has 2 main components:
  - transitioning to operational status: multipoint service similar to existing general internet, but private to the LHC community
  - under construction: point-to-point service with dedicated capacity and traffic separation
- Investigating use of SDN in the future
- Fits/supports ScienceDMZ model
- Multipoint service today connects 44 LHC computing sites (Tier0/1/2)
- Supported/provided by many NRENs in US and Europe, Asia, and soon South America
- Success yet to be evaluated
  - First feedback 'encouraging', but struggling with complexity in current deployment
  - Operational model yet to emerge
- Global engagement, hard to quantify, many fractional FTEs in NRENs and at LHC computing sites

Topical Panel Resources

- Snowmass 2013 community study
  - Final computing report issued Jan. 23, 2014
  - arxiv.org/abs/1401.6117
  - 3-page Conclusion included as Appendix 3 of our report
- Other documents & workshops (some examples)
  - SC/ASCR Data Crosscutting Requirements Review (Apr. 2013)
  - HEP-ASCR Data Summit (Apr. 2013)
  - HEP-NP Network Requirements (Aug. 2013)
  - FIFE Workshop (Jun. 2013)
  - HEP's influence on ESnet Development (W. Johnston, 2013)
  - LQCD influence on computing technology (S. Sharpe, Jan. 2014)
  - …

Meeting Agenda (Day 1)

- Monday December 9, 2013
  - Session 1: Opening Session
    - Opening Remarks and Expectations (HEP AD Jim Siegrist)
    - Guidelines, overview & details (Habib, Avery, Chatterjee, Price)
  - Session 2: Technology
  - Session 3: Energy Frontier
  - Session 4: Intensity Frontier (Part 1)
  - Session 4: Intensity Frontier (Part 2)
  - Session 5: Computing Issues 1 (Science Frontiers cross cuts)
    - Software sustainability models
    - Common needs in software and infrastructure
    - Ways to achieve better collaboration, increased flexibility & creativity, etc.

Meeting Agenda (Day 2, Morning)

- Tuesday December 10, 2013
  - Session 6: Accelerators
  - Session 7: Cosmic Frontier
  - Session 8: Computing Issues 2 (Next-generation infrastructure)
    - Adoption and design for next-gen hardware and infrastructure
    - Needs re storage, networks, compute 'cycles'
    - HPC for experiments, data-intensive computing with HPC-like nodes
    - New storage models, intelligent networks, etc.
  - Session 9: Computing Issues 3 (Data issues)
    - Data issues across the three frontiers, what is and is not common
    - Possibilities for common infrastructure usage, databases
    - New methods for large-scale data analytics, etc.

Meeting Agenda (Day 2, Afternoon)

- Tuesday December 10, 2013 (cont.)
  - Session 10: Computing Issues 3 (Data issues continued)
    - Presentation by SC Senior Advisor Laura Biven
    - New DOE rules (data management, long-term data storage, availability)
    - How can HEP connect to other offices, esp. BES, BER, NP, etc.
  - Session 11: Computing Issues 4 (Partnerships)
    - Presentation by ASCR Facilities Director Barb Helland
    - Discussion on ASCR-HEP research partnerships
    - Other partnerships
  - Wrap-up

Meeting Agenda (Day 3)

- Wednesday December 11, 2013
  - Session 12: Computing Issues 5
    - Summaries by scribes
    - Agree on what’s been done
    - Next steps, assignments
  - Close out
  - Wrap-up discussions & working lunch

Computing Topics Explored

- Areas covered
  - High Throughput Computing
  - Distributed / Grid Computing
  - High Performance Computing
  - High Performance Networking
  - Large-Scale Data Storage
  - Large-Scale Data Management and Analysis
  - Global Scale of Operations
- Areas not covered (“detector technologies”)
  - Online Data Processing
  - Data Acquisition

Overview of Major Areas (I)

- Cosmic Frontier
  - Cross-cuts with Intensity Frontier for dark matter experiments, data analysis and management pipelines
  - ‘Big data’ support, use of HPC and network resources
- Intensity Frontier
  - *Very* diverse set of requirements from IF experiments
  - Some experiments have connection to Energy Frontier activities
  - Opportunity for common approaches, new developments
- Energy Frontier
  - Next-gen hardware issues
  - Cross-cuts with Intensity Frontier
  - Use of HPC for Energy Frontier experiments, LHC Run II

Overview of Major Areas (II)

- Lattice QCD
  - Operates as facility, with SciDAC funding, competitive subgrants
  - A possibly useful model to emulate
- Accelerators
  - Sustainability of modeling software
  - Collaborations w/ other frontiers
  - Need for integration of diverse accelerator simulation codes
  - Conversion of code base to next-gen HPC
- Technology
  - Opportunities for collaborations with ASCR facilities
  - Next-gen hardware/software
  - Adoption of cloud technologies

Cross-Cuts Discussion Points

- Experiment timelines
  - Timelines of experiments must be respected / supported
- Development of common tools
  - Clear benefit: Reduce substantial redundancies
  - Possible cost: Sufficient flexibility to undertake new approaches?
- Future focus: Better integration w/ activities outside of HEP
  - Software practices
  - Data management
  - Hardware utilization
  - Training processes (students, postdocs, faculty, staff)

Table of contents

Topical Panel Report

- 2 page Exec. Summary
- 3.5 pages of Findings & Comments
- 2 pages of "Future Opportunities"

Findings: Hardware Evolution

1. Emerging computing architectures pose a major across-the-board challenge to HEP computing

   Less memory/core, more concurrency, reduced resiliency, etc. require new programming models & workflows. HEP follows commodity computing.

2. Evolution of data archiving, data-intensive computing, and storage will drive new computational strategies

   Unit storage costs will decline more slowly, requiring new balances of CPU, storage, networking, and advanced analysis strategies.

3. HPC platforms are an important resource for HEP science

   Includes DOE (NERSC, ALCF, OLCF) & NSF (SDSC, TACC, PSC). Some HEP groups (Accelerator, Cosmic Frontier, LQCD) already exploit these. HPC sites provide early access to new architectures.

Findings: Software Environment

4. Code maintenance and distribution lacks well-defined guidelines and support mechanisms

   No base mechanism to maintain, update, and distribute codes within HEP programs. Problematic given architecture evolution. Updates needed simply to maintain capability at its current level, especially for key software components such as Geant4, Panda, Root, etc.

5. Common tools and coding standards are insufficient

   "Stove-piping" inherent to project-oriented software development & utilization inhibits awareness of tools built outside of given experiment/frontier that could otherwise provide broader benefits.

6. Code diversity is a strategic issue

   Large diversity of codes, with significant overlaps in capability. Architecture changes will lead to different evolutionary paths, making it difficult to implement common strategies.

Findings: Resource Management

7. The scale of computational needs of smaller-scale projects is significant

   Ensemble of smaller projects in IF and CF require large aggregate computing support. Coordinated strategy (with enough expertise & manpower) difficult to implement without defined inter-project mechanisms.

8. The role of simulations will continue to grow in importance

   Driven by complexity of experiments and increases in computing capability.

9. Investment in data preservation and the formulation of an associated data policy is lacking

   Increased pressure to be able to analyze data from past experiments. Some efforts already exist within CF and IF, but any large-scale effort will require significant framework commonalities.

Findings: Organization of Computing Tasks

10. Computational R&D programs are needed

    Two classes: (1) evolutionary developments, specified by clear needs from experiments and (2) possibly revolutionary approaches targeting large gains in science impact and productivity. 2^{nd} type better suited outside of project boundaries.

11. More uniform interactions w/ external organizations are desirable

    Single entity functioning as the natural coherent point of contact for computation-related matters between the communities of HEP, ASCR, other DOE offices, NSF and industry. Natural role for the Center for Computational Excellence

12. Training in computational science is lacking

    Shortage of computationally well-trained students and junior researchers cited multiple times. Needs in programming paradigms, algorithms, software tools.

Findings: Center for Computational Excellence

13. The formation of a distributed Center for Computational Excellence is viewed as highly desirable
    - The CCE could address several of the concerns and opportunities identified in the above findings.
    - It would enable quick and effective responses to a number of identified gaps and well-defined tasks
    - It could initiate collaborations and mechanisms for cross-frontier activities and interactions with external entities (ASCR, NSF).

Important Caution!

- As currently constructed, our Topical Panel cannot give "advice" or make "recommendations".
  - Advising/recommending requires setting up under FACA rules
  - Federal Advisory Committee Act (1972)
  - http://www.gsa.gov/portal/content/100916

- Our report instead identifies "Areas of Opportunity" and we refer to "research communities", "particle physics", etc. rather than Agency entities throughout the document.

Computing Visibility: A Major Opportunity

- From the Topical Panel Report

  "… a fundamental observation is that bringing major computing-related activities to a higher level of visibility within the high energy physics community would present opportunities for a more coordinated and optimized approach, especially as it pertains to the software and resource management aspects of the findings listed above. The Center mentioned in the charge would be a powerful element of such an initial step."

"Areas of Opportunity" (1-2)

- 1. Code modernization, maintenance, and dissemination

  Well-defined mechanisms for the continued maintenance and development of a number of particle physics software frameworks and tools, especially those that cut across frontiers, can greatly benefit the HEP science effort.

- 2. Common tools & coding standards; reduced software footprint

  Husbanding software resources in high energy physics, given future funding and manpower limits, strongly argues for increased emphasis on the use of common tools across and within frontiers, and the development of a shared set of 'high energy physics computing best practices' to optimize resources.

“Areas of Opportunity” (3-4)

- 3. Resource support models for smaller-scale projects

  Smaller-scale projects are resource-starved in their ability to exploit complex tools and to develop new ones. To address this need, a support model that cross-cuts across all frontiers would significantly aid these projects.

- 4. Data preservation policy for the HEP community

  Establishment of data preservation policies within the community that are consistent with DOE–HEP and with broader DOE-SC Data Management Requirements would promote the long-term value and integrity of HEP science.

"Areas of Opportunity" (5)

- 5. Distributed Center for Computational Excellence (CCE)

  The CCE would provide opportunities for a cross-cutting R&D program that includes compute and data intensive elements addressing the challenges of next-generation architectures, innovative approaches to the use of high performance computing for HEP science, and data-intensive collaborations with the ASCR community, including networking and data transport, as well as partnerships with industry when appropriate.

"Areas of Opportunity" (6)

- 6. Multi-level computer and computational science training activities

Activities aimed at directly impacting the particle physics science agenda and advancing next-generation software development are expected to have significant scientific impact.

“Areas of Opportunity” (7-8)

- 7. Community-based expert group for HEP computing

  Such a group would provide opportunities to continue the community-based approach to address issues in computing and simulation and address technical details.

  An expert group meeting on a predictable schedule would provide additional visibility for HEP computing activities.

- 8. Expansion of current interactions with researchers in external disciplines, particularly those in DOE-ASCR community

  New opportunities exist in key areas such as data-intensive science, with emphasis on ASCR computational and data facilities; in initiation of ‘triangular’ collaborations with researchers from other program offices and industry, in collaboration with the ASCR community; in establishing collaborations with research communities supported by other agencies (e.g., NASA, NSF), as relevant.