HEP Forum for Computational Excellence (HEP-FCE)

Vision and First Activities

Salman Habib & Rob Roser
(Co-Directors)

HEP Computing Frontier

- Computing is an essential enabling and empowering component of almost all aspects of HEP science
- HEP computing has a long history (~60 years) including notable contributions to High Performance Computing (HPC), High Throughput Computing (HTC), and large-scale Data Science
- Substantial resources are devoted to computation and data science as an essential aspect of HEP's scientific enterprise
- Winds of Change: New challenges posed by hardware evolution and rapid increase in data rates/volumes in an era of flat/ declining budgets
- Wide recognition that concerted response leveraging HEP community expertise and best practices from the outside world is needed (HEP Topical Panel, Snowmass, P5)

From the P5 Report…

- **Recommendation 29** : "Strengthen the global cooperation among laboratories and universities to address computing and scientific software needs, and provide efficient training in next-generation hardware and data-science software relevant to particle physics."
- Driver for a set of HEP-FCE goals in collaboration and community-building efforts, R&D to focus on building cross-cutting dedicated expertise in architecture and data science evolution
- "Investigate models for the development and maintenance of major software within and across research areas, including long-term data and software preservation. (p. 21)"
- Driver for a set of HEP-FCE goals in R&D, cross-cut software, and data preservation

What is the HEP–FCE?

- Official DOE HEP response to P5 recommendations, Snowmass, and the HEP Topical Panel on Computing and Simulation (Dec 2013): "Computing Frontier" as key strength of HEP
- Cross-cut exchange to promote excellence in HEP computing, simulation, data transfer/storage/preservation and promote the flow of information and best practices across HEP Frontiers
- Identify and undertake focused R&D tasks addressing frontier challenges posed by rapidly evolving computer architectures, next-generation data-intensive computing, and development of shared software elements
- Undertake information gathering tasks – community-based studies and surveys
- Access point for agency/office partnerships (ASCR, NSF, NASA…) and HEP connections (CERN, DPHEP, G4, HSF, OSG, USQCD...)

Why NOW and WHY HEP-FCE?

Given flat/declining budgets, it is increasingly important to leverage work that has already been done and to use expertise that already exists

Present computing model in HEP handles virtually all computing within experiments/frontiers

- Need to strengthen and expand horizontal component to best utilize HEP resources, **avoid duplication,** optimize technology, **develop external resources & strengthen partnerships**……

Rapidly evolving computer architectures and exponentially increasing data volumes REQUIRE effective responses & cross-cut solutions

- Evolution of commodity hardware and rise of cloud computing
- Enhance opportunities of interfacing with ASCR
- Understanding HEP's core competencies and leveraging that expertise moving forward

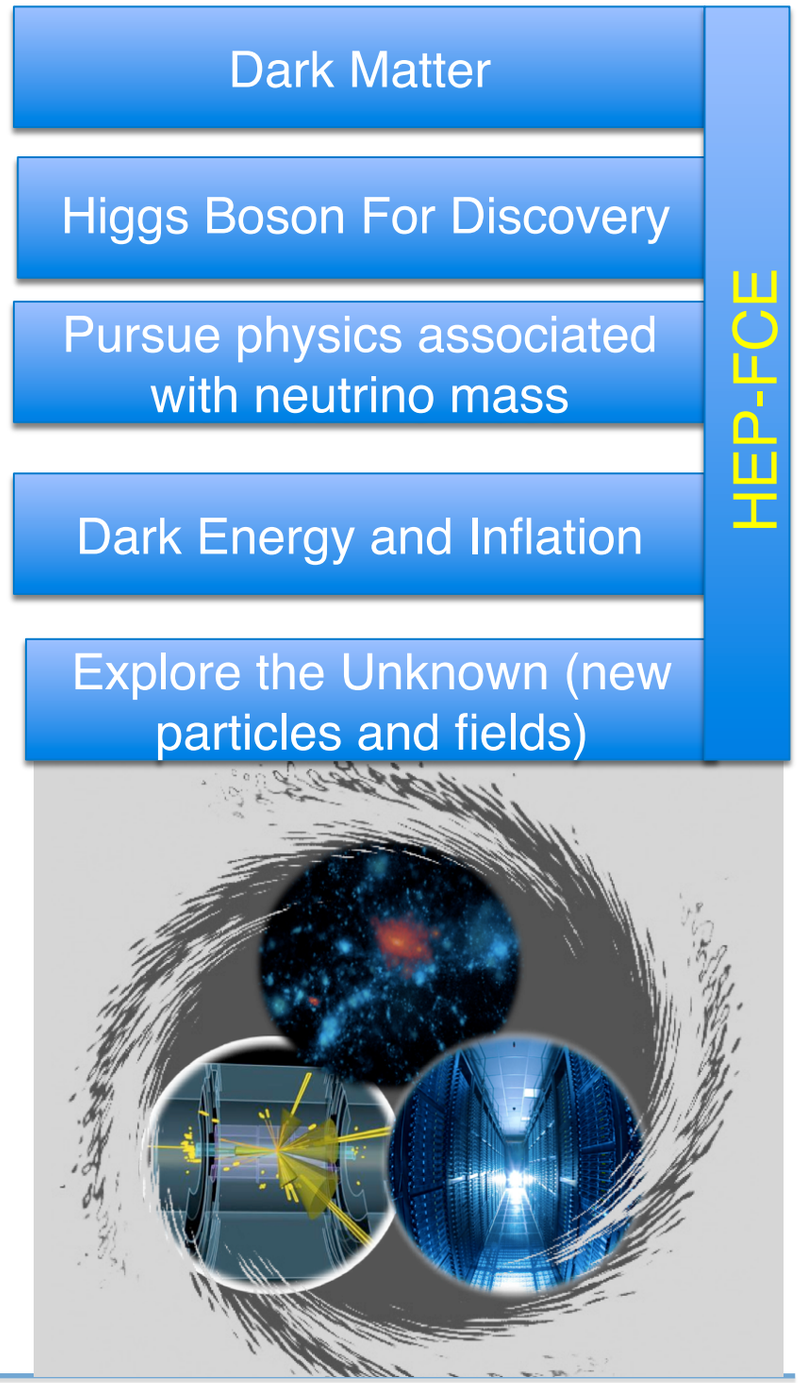

Roles that the FCE can play…

- Computing-related HEP community planning
- Coordinate HEP-FCE R&D portfolio focused on cross-cutting initiatives and tasks
- Advice on computational support for smaller/new experiments
- Provide cross-cutting knowledge/capability base
- Provide Training/Workshop activities
- Engage/reward junior researchers (training, travel/visiting awards)
- Disseminate information re data archiving/curation/management
- Strategic connectivity for DOE HEP program; serve as a community hub
- Work with existing groups to leverage capabilities – OSG, SciDAC teams, and others

Note: HEP-FCE is an intellectual provider, not a service provider; main roles above relate to facilitation and incubation

As these activities mature – we can expand the role further

HEP-FCE: Where we are today I

- HEP-FCE is still in its early days
- Success of HEP-FCE is tied to the success of collaboration and leveraging
- Embarked on a pilot project in mid-2014
- Started out by launching 3 Working Groups: Based on Snowmass, HEP Computing Topical Panel, and P5 reports, along with community input, the WGs were tasked to identify concrete, actionable opportunities in the HEP computational landscape:
    - Systems (hardware and systems software)
    - Tools and Libraries
    - Applications Software
- Working Group Reports: Draft reports in hand (public release in April 2015); major general conclusions:
    - Systems: Focus on 1) Data storage and data access technologies, 2) Future computer architectures
    - Tools and Libraries: 1) Adopt sharing across experiments, 2) Step beyond local solutions, 3) Avoid systemic shortcomings
    - Applications Software: 1) Top-down consolidation not required, 2) Need for recognition of software importance — on par with major detector R&D

HEP-FCE: Where we are today II

- Community Support and Collaboration Development: Web presence established at http://hepfce.org; initial team of experts and volunteers in place
- Next Step: First HEP-FCE workshop in first half of 2015; coordinating exascale requirements activity with DOE ASCR

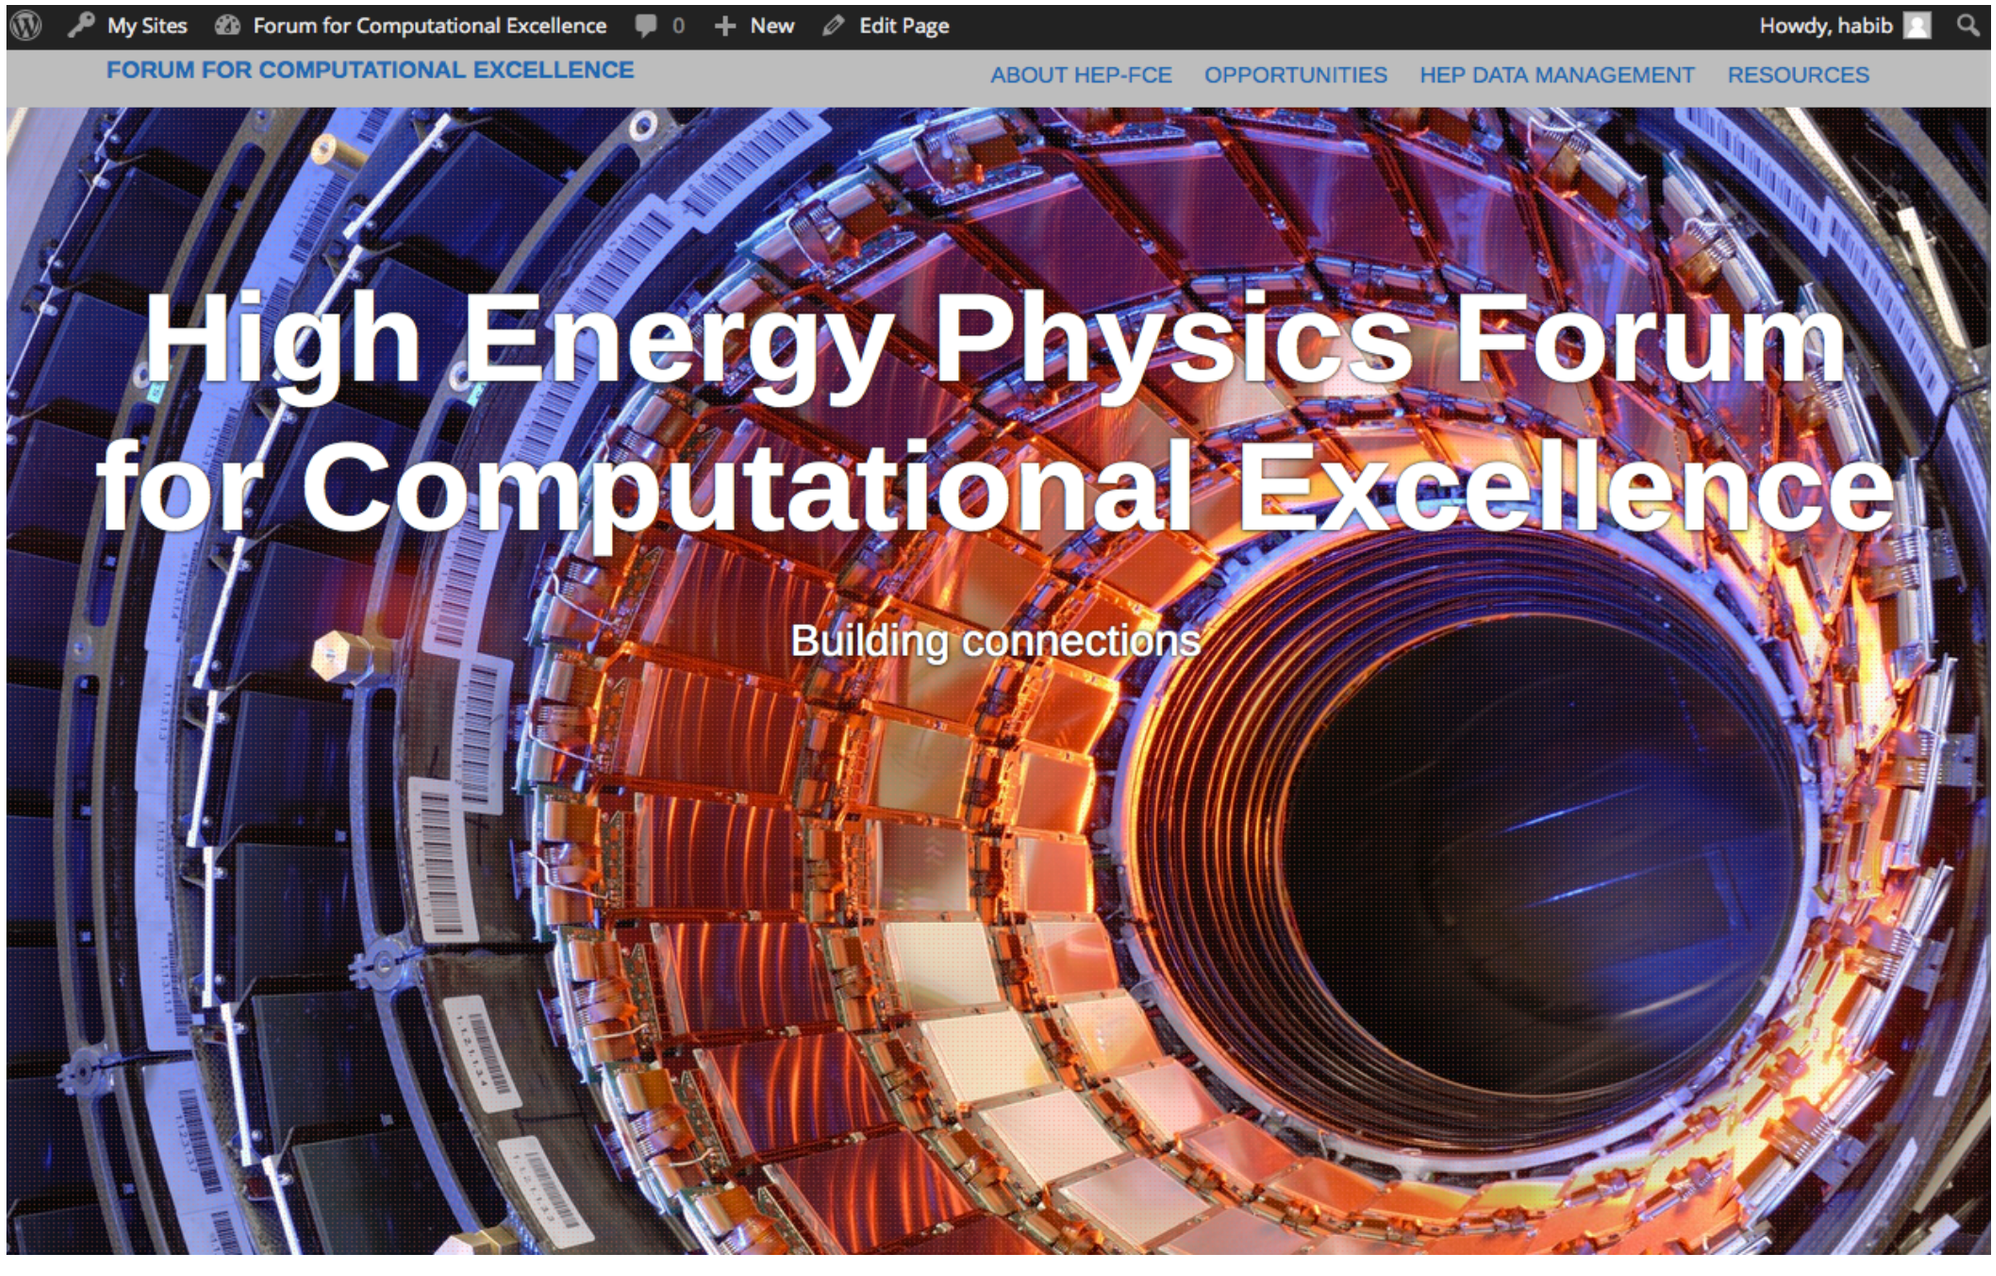

HEP-FCE: People

- Advisory Panel: David Asner (PNNL), Paul Avery (Florida), Amber Boehnlein (SLAC), Andy Connolly (Washington), Steven Gottlieb (Indiana), Craig Tull (LBNL), Torre Wenaus (BNL)
- HEP-FCE Working Groups: Systems – P. Calafiura (LBNL), A. Dart (Esnet), O. Gutsche (FNAL), T. Izubuchi (BNL), A. Lyon (FNAL), P. Nugent (LBNL), D. Petravick (NCSA); Libraries and Tools – A. Borgland (SLAC), P. Elmer (Princeton), M. Kirby (FNAL), S. Patton (BNL), M. Potekhin (BNL), B. Viren (BNL), B. Yanny (FNAL); Software – M. Asai (SLAC), L. Bauerdick (FNAL), H. Finkel (ANL), S. Gottlieb (Indiana), S. Hoeche (SLAC), T. LeCompte (ANL), Z. Marshall (LBNL), P. Sheldon (Vanderbilt), J.-L. Vay (LBNL)

J. Amundson, P. Avery, S. Bailey, D. Bard, L. Bauerdick, D. Benjamin, T. Bhattacharya, A. Borgland, J. Borrill, P. Calafiura, P. Canal, M. Carrasco-Kind, J.T. Childers, D. Daniel, E. Dart, V.D. Elvira, C. Faham, H. Finkel, N. Frontiere, G. Garzoglio, S. Gottlieb, O. Gutsche, S. Habib, M. Hance, K. Heitmann, S. Hoeche, B. Holzman, R. Illingworth, H. Ito, J. Kowalkowski, T. LeCompte, L.T. Lo, Z. Lukic, A. Lyon, Z. Marshall, B. Messer, V. Morozov, R. Mount, C.-K. Ng, P. Nugent, G. Perdue, G. Perumpilly, A. Pope, M. Potekhin, A. Rangel, R.D. Ryne, M. Sanchez, H. Schellman, M. Schram, C. Sewell, E. Sexton-Kennedy, J. Simone, P. Spentzouris, J. Strube, R. Thomas, T. Uram, J.-L. Vay, B. Viren, J. Woodring

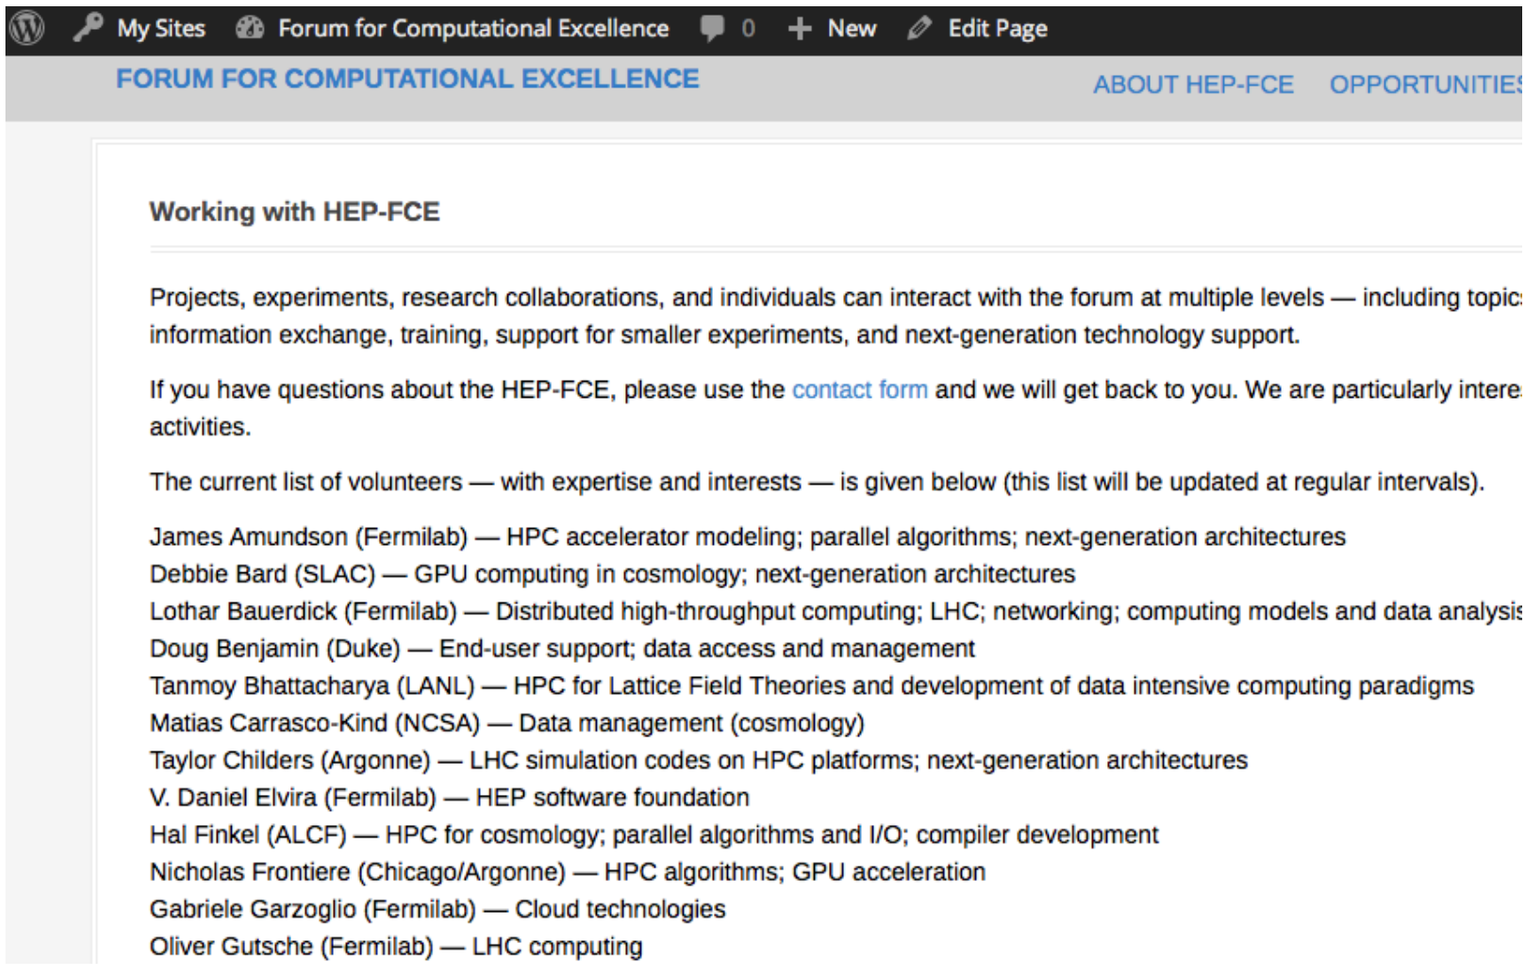

Enhance and Expand the HEP/ASCR Connection I

1. SciDAC (Scientific Discovery through Advanced Computing)
   - Lattice QCD
   - Accelerator modeling (parallel computing came to this community via a DOE Grand Challenge followed by SciDAC)
   - Computational cosmology
   - All three of these cases have solid success stories and are ongoing…
2. ESnet – partnership essential for Energy Frontier, also for Cosmic and Intensity Frontiers
3. ALCC awards – Cosmic Frontier End-Station, LHC Event Simulation, Theory
4. INCITE awards – Lattice QCD, Accelerators, Computational Cosmology
5. Other partnerships – showing directions forward but not yet at the same level as the ones above – include work on refactoring HEP software and building new data-intensive capabilities
6. Partnerships include work at all three major ASCR facilities: ALCF, NERSC, OLCF

Enhance and Expand the HEP/ASCR Connection II

7. Leveraging HEP capabilities –
    1. In the vein of "do what we do best" but learn/adopt best practices from the outside world
    2. Help promote the success of OSG and identify further opportunities
    3. Use our own internal expertise and external connections (with ASCR) to encourage and help others in HEP. Examples include:
        1. Hoeche et al., work on Sherpa
        2. LeCompte et al., running ATLAS event simulation on a production basis on ANL HPC machines
        3. Dark Energy Survey data management pipelines running at NERSC
8. "Mini-Apps" – benchmark and improve our core software apps; help HEP and ASCR make intelligent decisions for resource acquisition
9. Joint ASCR/HEP exascale requirements planning meeting in June, 2015

The Impact of HPC on HEP

- Not all problems can be solved using HPC systems, but many can (accelerators, cosmology, event generation/simulation, QCD…)
- Next generation of ASCR HPC machines (staging begins 2016, ends in 2018) will sum to ~500 petaflops of compute capability
- If HEP experiments use just 5% of that, i.e. 25 petaflops, it is ~25 times what the Grid will provide
- Learning how to leverage these resources to seamlessly supplement/enhance current capability is important
- New possibilities opened up by HPC platforms will offer unique computational opportunities

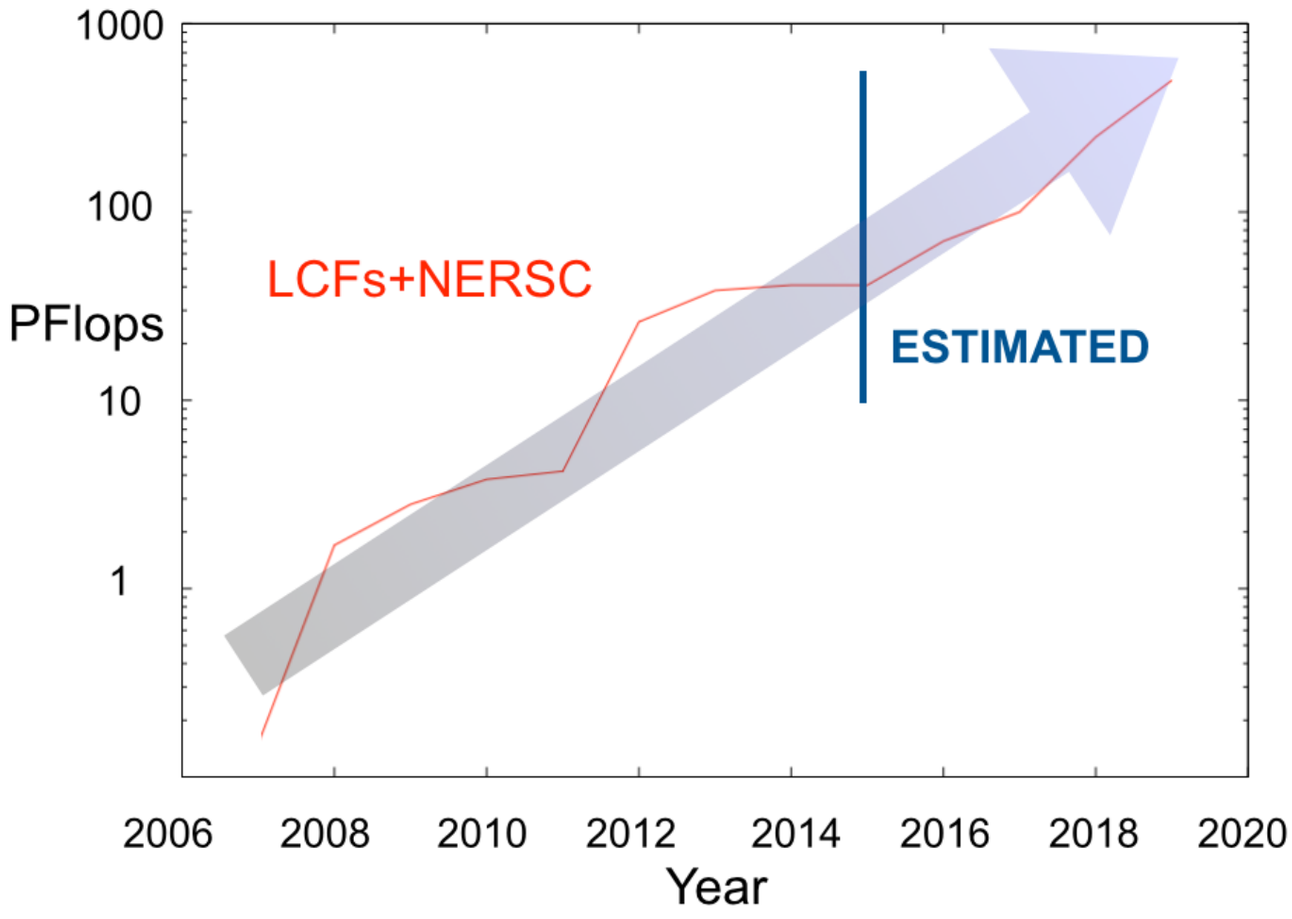

HPC for HEP Experiments

ATLAS has about a year and a half experience of production HPC use (Comp. HEP pilot project)

- ALCC award equivalent to several percent of the grid
- Used to run event generators
- Highlighted areas for code improvement
- Running at scales of 100K-1M parallel processes (close collaboration between HEP/ ALCF/code authors)

HPC fraction can grow, path easier for new experiments (e.g., IF)

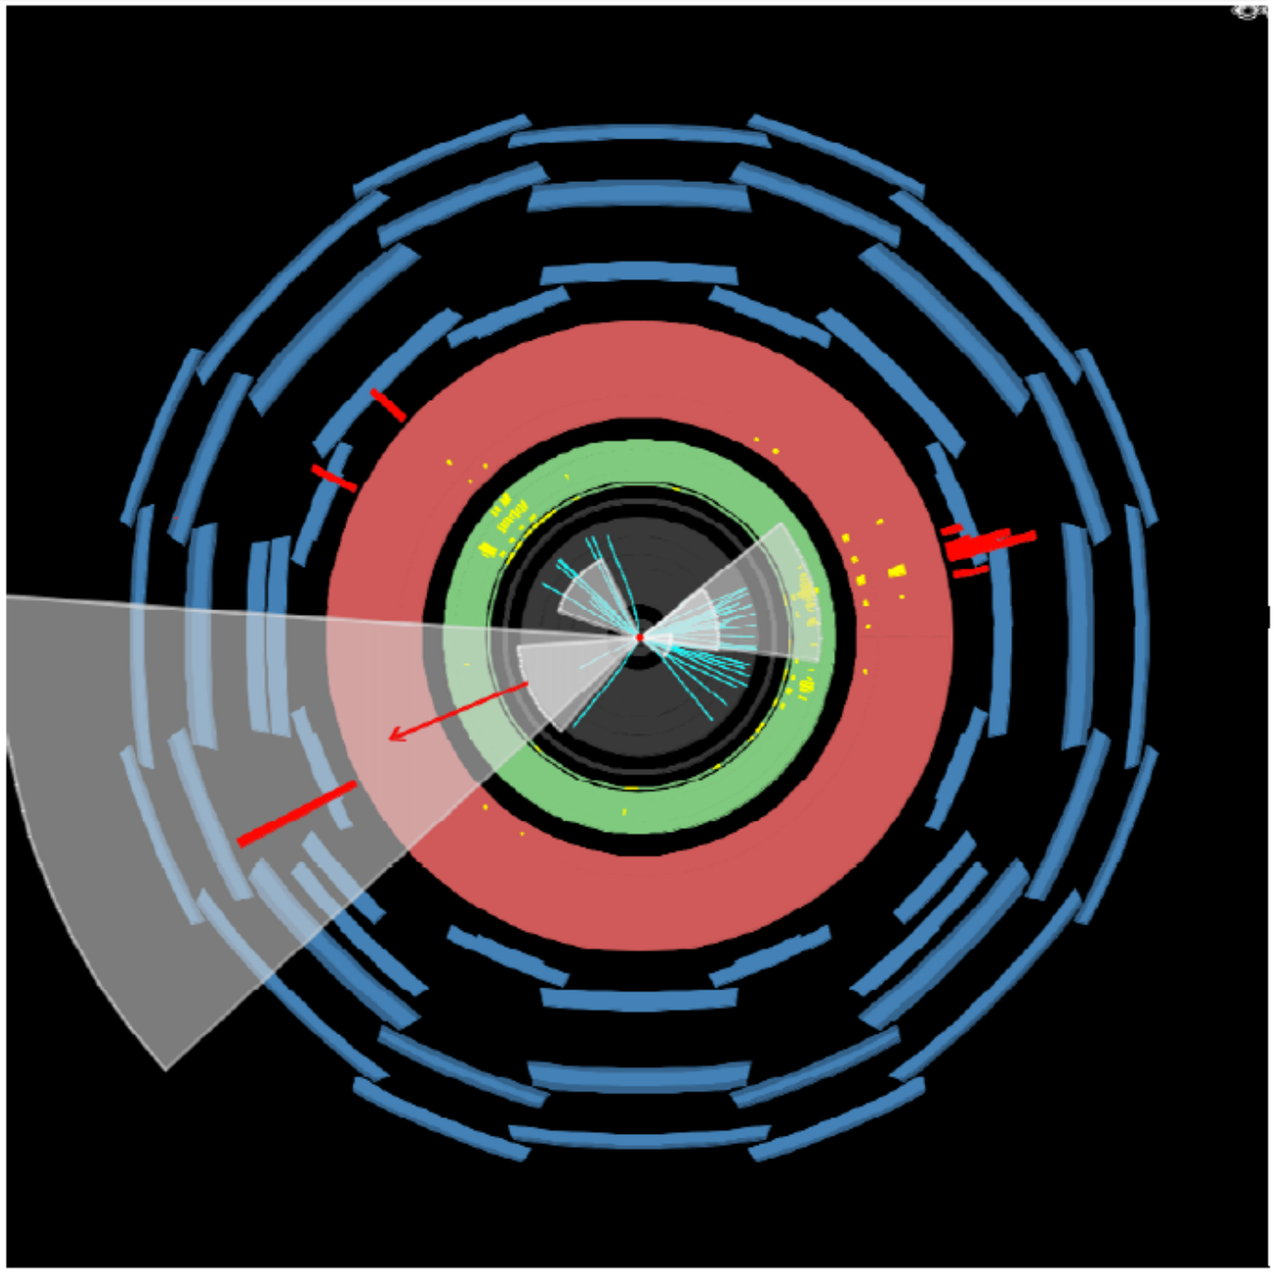

Event display for a simulated z+5 jets event: Generated on Mira (ALCF) with Alpgen, simulated on Edison (NERSC) with Yoda, reconstructed at CERN

Next-Generation Architectures: HEP Mini-Apps

- Real applications usually too large and complex to work with in design studies
- Kernels alone are often inadequate
- Mini-apps (Heroux/Barrett): Compact, self-contained proxies that adequately portray the computational loads and dataflow, but are not meant to be scientifically realistic
- HEP mini-app collection (thousands of lines of code each) will be a good way to develop a design suite for next-generation architectures
- Mini-app instances can cover all HEP application areas

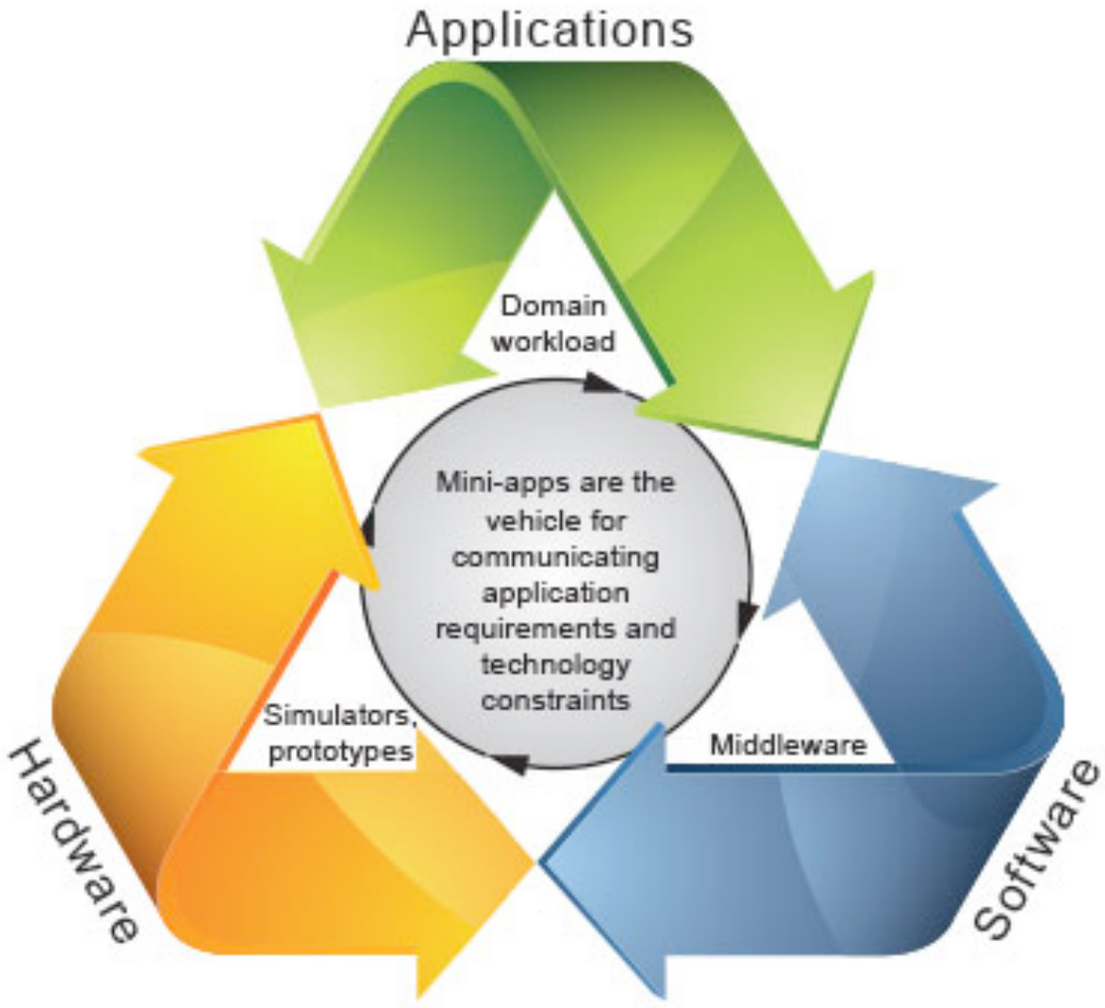

Co-design loop using mini-apps to approximate the workload of the parent application (LLNL STR 2013)

The Realities of HEP Computing Moving Forward

- Due to funding constraints, we will have to optimize computing resources for "average" demand and not peak
- Need to find creative solutions for those instances where we need more than we have and need it fast ("pledged" vs. "non-pledged" resources)
- Importance of leveraging HPC facilities – already demonstrated at level of several percent of ATLAS production computing, equivalent to all of Spain's contribution (LeCompte et al.)
- Computing in the Cloud will be an important player in future
- Helping each other across Frontiers/projects – sharing core competencies essential
- Finding ways to not just develop key software – but to maintain and evolve it going forward

Parting Thoughts…

- The HEP Forum for Computational Excellence has been launched
- While still in its infancy, it is gathering support from the HEP community and has a well defined path forward
- The plan over the next year or so will be to firmly establish the HEP-FCE in the evolving computing landscape and to promote the importance of computing within the broader HEP community
- The Forum will help the community seamlessly transition from today's paradigm to one that is based on best use of emerging architectures and programming models
- Computing has always been a strength of the HEP community – the FCE will help leverage that competency across the various scientific enterprises and take advantage of the existing expertise to solve outstanding problems

BACKUPS

FCE Next Steps

- Next-Generation Architectures/Supercomputing Applications
  - FCE will help foster and build collaborations to enable a new set of supercomputer users within HEP
- Data-Intensive/Cloud Computing
  - Virtualization and software containers offer elasticity of compute resources and handle complex software stacks
- Cross-Cut Software Development
  - Developing a strong cross-cut software base, by functioning as a helpful go-between across experiments, institutions, and frontiers
- High-Speed Networking
  - Engagement with entities such as ESnet, HEP and ASCR facilities, and software teams that are focused on managing data movement
- ASCR/HEP Interactions and Workshops
  - Joint meetings with HEP and ASCR to shape future collaborative activities
- HEP-FCE Infrastructure Support and Community Development
  - Develop and maintain a HEP community infrastructure for sharing ideas, algorithms, codes and evolving sets of best practices.

Those Involved in the Forum

Others Engaged or Engaging

- Argonne National Laboratory: Taylor Childers, Hal Finkel, Katrin Heitmann, Vitali Morozov, Adrian Pope, Tom Uram
- Brookhaven National Laboratory: Hironori Ito, Maxim Potekhin, Brett Viren
- Fermi National Accelerator Laboratory: James Amundson, Lothar Bauerdick, Philippe Canal, Robert Illingworth, Jim Kowalkowski, Adam Lyon, Liz Sexton-Kennedy, James Simone
- Lawrence Berkeley National Laboratory: Stephen Bailey, Julian Borrill, Paolo Calafiura, Carlos Faham, Zarija Lukic, Zach Marshall, Rollin Thomas, Jean-Luc Vay
- SLAC National Accelerator Laboratory: Debbie Bard, Anders Borgland, Stefan Hoeche, Cho-Kuen Ng
- Universities and Other Institutions: Paul Avery (Florida), Doug Benjamin (Duke), Tanmoy Bhattacharya (LANL), David Daniel (LANL), Eli Dart (ESnet), Steven Gottlieb (Indiana), Bronson Messer (ORNL), Mayly Sanchez (Iowa State), Heidi Schellman (Oregon State), Malachi Schram (PNNL)