Snowmass Input to P5

Computing Frontier Study Group

L.A.T. Bauerdick, S.Gottlieb,
for the Computing Frontier Group

Snowmass Computing Frontier Study

Conveners: Lothar Bauerdick (FNAL), Steven Gottlieb (Indiana U)

The charge:

- What are the computational **requirements for carrying out the experiments** that will lead to advances in our physics understanding?
- What are the computational **requirements for theoretical computations and simulations** that will lead to advances in our physics understanding?
- What **facility and software infrastructure** must be in place in order to meet these requirements, and what **research investments** does it require in computing, storage, networking, application frameworks, algorithms, programming, etc. to provide that infrastructure?
- What are the **training requirements** to assure that personnel are available to meet the needs?

Subgroups for "user needs" **Computing needs of physics**

- CpF E1 Cosmic Frontier
  - Alex Szalay (Johns Hopkins), Andrew Connolly (U Washington)
- CpF E2 Energy Frontier
  - Ian Fisk (Fermilab), Jim Shank (Boston University)
- CpF E3 Intensity Frontier
  - Brian Rebel (Fermilab), Mayly Sanchez (Iowa State), Stephen Wolbers (Fermilab)
- CpF T1 Accelerator Science
  - Estelle Cormier (Tech-X), Panagiotis Spentzouris (FNAL); Chan Joshi (UCLA)
- CpF T2 Astrophysics and Cosmology
  - Salman Habib (Chicago), Anthony Mezzacappa (ORNL); George Fuller (UCSD)
- CpF T3 Lattice Field Theory
  - Thomas Blum (UConn), Ruth Van de Water (FNAL); Don Holmgren (FNAL)
- CpF T4 Perturbative QCD
  - Stefan Hoeche (SLAC), Laura Reina (FSU); Markus Wobisch (Louisiana Tech)

Subgroups for "infrastructure" **Technical capability mapping to needs**

- CpF I2 Distributed Computing and Facility Infrastructures
  - Ken Bloom (U.Nebraska/Lincoln), Sudip Dosanjh (LBL), Richard Gerber (LBL)
- CpF I3 Networking
  - Gregory Bell (LBNL), Michael Ernst (BNL)
- CpF I4 Software Development, Personnel, Training
  - David Brown (LBL), Peter Elmer (Princeton U.); Ruth Pordes (Fermilab)
- CpF I5 Data Management and Storage
  - Michelle Butler (NCSA), Richard Mount (SLAC); Mike Hildreth (Notre Dame U.)

… and participation from the wider community

Computing Frontier Results of the Summer Study

- Draft Report of Computing Frontier Study (Oct-2013)
  - http://www-public.slac.stanford.edu/snowmass2013/docs/preliminarypublic10-30/ComputingFrontierOct30.pdf

- Sub-group report coming out on arXiv

- Summary presentation at Snowmass Meeting (Aug-2013)
  - https://indico.fnal.gov/getFile.py/access?contribId=13&sessionId=1&resId=0&materialId=slides&confId=6890

- T.Wenaus' presentation at DPF2013 (Aug-2013)
  - https://indico.bnl.gov/getFile.py/access?contribId=24&sessionId=0&resId=0&materialId=slides&confId=603

Computing Challenges at the Physics Frontiers

✦ Energy Frontier

- ★ Programs will likely have **10-fold increases in trigger rate, complexity**
  - ✦ computing costs will constrain data rate, and thus choices on triggers, analyses
  - ✦ LHC raw data: ~15 PB now; ~130 PB in 2021
  - ✦ however, total dataset sizes are 10x larger, including processing steps and MC
  - ✦ in future, data management must become much more efficient
- ★ **Storage is cost driver**, as disks get cheaper only more slowly
  - ✦ not all data needs to be on disk — $10 puts 1M additional events on tape (CMS)
- ★ To control processing cost, **we must track Moore's law effectively**
  - ✦ Adapting to new processors is much more challenging than in the past

✦ Intensity Frontier

- ★ Smaller, but significant!, computing challenges, and **a large diversity** of experiments (N.B., Belle II estimates LHC Run 2-like data rates)
- ★ Our **survey of experiments** suggests **convergence on common computing model; commonality**, also w/ Energy Frontier
- ★ Avoid duplication, **profit from common expertise and approaches**

Computing Challenges at the Physics Frontiers

✦ Cosmic Frontier

- ★ Very large computing & data needs: sky surveys, radio telescopes, also for simulation
- ★ >1 PB data today, 50 PB total at end of decade, 400 PB/yr in 10-20 yrs
- ★ Advanced databases and data management, HPC, need for data analytics

✦ Accelerator Science

- ★ Simulation important for achieving high gradients for EF, low losses for IF
- ★ Very high computing needs, HPC, required to produce end-to-end designs, wide range of modeling scales, from particle bunch to accelerator complex
- ★ R&D on advanced algorithms to utilize new processor architectures

✦ Non-Perturbative and Perturbative QCD

- ★ pQCD, Lattice QCD crucial for EF, IF experimental programs: interpreting data requires theoretical predictions with commensurate precision
- ★ Continued reliance on HPC facilities; "capability" and "capacity" computing, USQCD project manages allocations, among the largest in the US
- ★ Benefit from new processor technologies and algorithms, such as GPUs

Computing Infrastructure for HEP
Grids, HPC Centers, Clouds

- ✦ Challenging resource needs require **efficient and flexible use of all resources**
  - ★ **Distributed High-Throughput Computing** (Grids)
    - ✦ OSG provides 800M hours/year to EF, IF experiments and others
  - ★ **High-Performance Computing** (mostly theory/simulation/modeling)
    - ✦ HPC is used and required by a number of projects, in 2013 total was 1,400M hours/year
  - ★ programs need to continue to fund these resources, both dedicated and shared
- ✦ **Sharing** and **opportunistic use** help address resource needs, from all tiers of computing, eventually including commercial clouds etc
  - ★ More communities need **data intensive computing**, including at HPC, e.g. CF for data analytics, combining simulations and observational data etc.
- ✦ Cloud Providers
  - ★ **Commercial clouds still too costly** to replace dedicated resources
    - ✦ recent study (T.Wong/BNL, shown at HEPIX): Cost of computing/core at dedicated data centers compare favorably with cloud costs $0.04/hr (at BNL RACF) vs. $0.12/hr (amazon EC2)
  - ★ More existing HEP Grid and HPC resources are moving to cloud interfaces
  - ★ Cloud approach allows peak responses, dynamic resource provisioning
  - ★ **Significant gaps and challenges** exist
    - ✦ managing virtual machines, workflows, data, cyber-security, and other areas

R&D Needs

- The Snowmass discussions identified main challenges that drive the need for specific near-to-midterm R&D. These require efforts beyond the program-funded R&D at labs and universities, and need help in addition to the computing and software operations programs
- Computing is “unique” in that although we don't have to develop our own hardware, we do have to constantly re-analyze our needs, adapt to a rapidly changing environment and maintain and develop complex and deep software stacks
- Just to keep going requires continuos R&D and engineering
- We are not building/buying hardware that is going to last 20 years, but our computing projects may well last that long. HEP software and systems may have quite a long lifetime, and will require constant development, verification and deployment, while “hardware” upgrades may only happen every several years
- To enable this across the frontiers requires **strategic investments in several R&D areas**

Chip Architecture

- There is a major shift in the nature of processors
  - performance of single sequential applications has roughly stalled due to limits on power consumption
- For HEP processing performance to stay on the Moore's law curve, need to proactively make full/better use of **advanced architectures:** multi-threading, GPU environments, low-energy CPUs w/ "small" cores
  - With the need for more parallelization the **complexity of software and systems** continues to increase:
    frameworks, workload management, physics code
  - Important needs for **developing and maintaining expertise** across field, including re-engineering of frameworks, libraries and physics codes, adapting key software tools
- Unless corrective action is taken we could be **frozen out of cost effective computing solutions** on a time scale of 10 years.
  - There is a large code base to re-engineer
  - We currently do not have enough people trained to do it

Data Storage, Management, Access

- The growth in HEP data drives the need for continued R&D investment in **data management, data access methods, networking**
  - Storage cost are not scaling with huge increases in future demand
  - Need continued evolution to take advantage of network capabilities
  - Ensure efficiency and robustness of the global data federations
  - **Contain the level of effort needed for operations**
- Rotating disk will suffer marked slowdown capacity/cost.
  - Computing models must attempt to optimize roles of tape, rotating disk, solid-state storage, networking and CPU
- This requires to develop and adopt a number of **new approaches**
  - including more advanced distributed (federated) storage solutions, data caching across sites avoiding duplication (“move instead of store"), using tiered storage to make use of cost-effective tape archive solutions, and even preferring re-computing instead of storing derived data sets, trading computing cost for storage cost (“virtual data”)

Networks and Distributed Data Access

- ✦ **Data intensive distributed computing** is enabled by networks
  - ★ Network capabilities and data access technologies improve our ability to **use resources independent of location**, over the network
  - ★ Enables use of a **large spectrum of resources**: dedicated facilities, university compute cluster, sharing and opportunistic use across communities, use of commercial clouds, leadership-class HPC,...
- ✦ Distributed **data and processing management systems** unify resources into a usable system
  - ★ emerging solutions are based on **content delivery networks** approach, **dynamic data placement** across federated storage, **remote data access** over the network
  - ★ treat **networks as resource**, include capabilities in computing models
- ✦ All Frontiers depend on reliable, high-bandwidth, feature-rich networks
  - ★ HEP's objectives through 2020 require basic and applied network R&D
- ✦ These technologies require investments,

  but already pay big dividends, e.g. at the LHC!

Access and Usability

- ✦ With distributed computing more pervasive, access and usability must be improved...
  - ★ …and documented, with friendly tutorials!
- ✦ Improve ease
  - ★ to acquire access rights and authenticate
  - ★ to access, move, manage data
  - ★ to use processing resources
- ✦ Promising initiatives underway, Open Science Grid an important driver

- ✦ Even in the LHC community, most accustomed to grids, local computing remains an “extremely important” complement to large scale resources

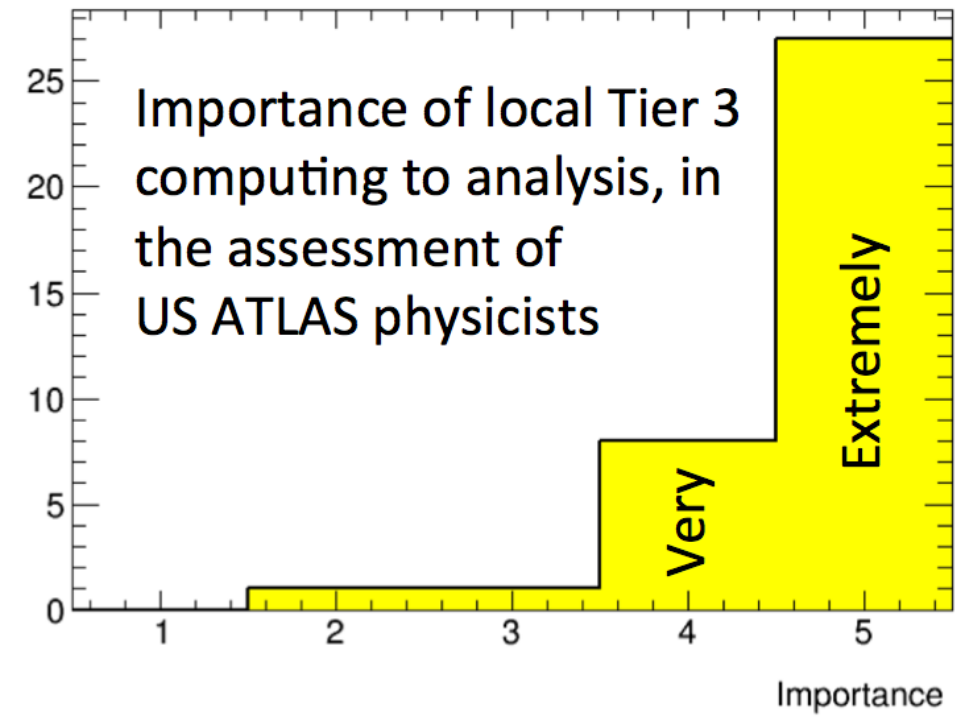

Need for Training and Career Paths

- Encourage and support training, as a continuing activity
  - Use certification to document expertise and encourage learning new skills
  - Use mentors to spread scientific software development standards
  - Involve computing professionals in training of scientific domain experts
  - Use online media to share training
  - Use workbooks and wikis as evolving, interactive software documentation

- We need to provide young scientists with opportunities to learn computing and software skills that are **marketable** for non-academic jobs

- We need training and **career paths** including tenure stream) for researchers who work at the forefront of computation techniques and science is critical

Computing Outlook

- ✦ Successes in HEP have always been closely tied to advances in computing
  - ★ New technologies and approaches have been and will be transformative
  - ★ Distributed computing, networks, parallelization, virtualization, GPUs etc
- ✦ Industry has caught up to us, and in cases has surpassed us
  - ★ Should take advantage of progress in industry and advances in other sciences
  - ★ There's much to leverage! Big Data is good news, Clouds are good news…
- ✦ Investing in community planning, solve common problems in partnerships
  - ★ Investments in common development are important
  - ★ Several already happening with agency support and sponsorship
- ✦ Collaboration and partnerships are essential
  - ★ Certainly within and across the frontiers
  - ★ Distributed computing requires partnerships, between sites, science communities, with computer scientists, between agencies etc
  - ★ A Virtual Center for HEP Computing is being discussed in DOE
- ✦ A concerted program is needed to address the computing challenges at the physics frontiers!

Additional Slides

Need for a Program

- Needed advances in computing require targeted R&D in a number of areas, and a partial list drawing from the Snowmass discussions is given in the following pages. In the spirit of Snowmass, we are looking for common solutions across the field and collaboration between experiments, institutes, the different frontiers, with computer science and other disciplines.

- As the problems are common across the field, realizing investments in common developments is important. A Virtual Center for HEP Computing is being discussed in DOE, and that could help in formulating common needs, identify gaps and propose target areas for specific R&D, foster collaboration for projects to address these needs, and be a center for training and sustaining knowledge in these areas.

- The following is a (possibly incomplete) list of areas that require an influx of innovative approaches, frameworks and technologies, and of ideas and projects that have been mentioned and put forward as part of the Snowmass Computing Frontier study.

Software Architecture, Engineering and Training

- To stay on the Moore's law curve we need to proactively make full use of advanced architectures. This requires to develop new algorithms and to re-engineer software for fine-grained parallelism, to learn to effectively use GPU and other co-processor environments and to make optimal use of low-power CPUs. This requires to develop and maintain a new level of expertise on parallel processing across the field, and to pull off a concerted effort for re-engineering of frameworks, libraries and physics codes, on adapting key software tools etc.

- With the need for more parallelization the complexity of software and systems and the need for specialized software know-how continues to increase: frameworks, workload management, physics code, increasing the need for training. Competence centers could address specific engineering needs.

Workflows

- HEP experiments in particular at the Energy Frontier have ever increasing CPU demands. But also the Intensity Frontier experiments are looking for more sharing and opportunistic use to help address peak resource demands, tapping into all tiers of computing, eventually including commercial clouds.

- Distributed data and processing management systems like GlideinWMS and PanDA unify resources into a usable system. The OSG is working with EF experiments and labs to aggressively pursue opportunistic use of CPU cycles wherever they become available, and on optimizing throughput. We need to strengthen a targeted program of making opportunistic resources fully useable, across frontiers, and of including HPC centers into HEP workflows where it makes sense. This requires work on the workload management systems, on dynamic resource provisioning, on virtualization of application environments, and on dealing with security and identity management issues, etc.

- With remote data access becoming available across the network, work in this area also addresses the need to enable distributed data intensive computing, including at HPC centers, for data analytics, combining simulations and observational data e.g. for Cosmic Frontier science, etc.

Storage Systems, Data Management and Big Data

- ✦ The growth in HEP data drives the need for continued R&D investment in data management, data access methods, and networking. Innovative data management solutions are based on content delivery network approaches, dynamic data placement across federated storage systems, and remote data access over the wide-area network. Moving forward, HEP should investigate advanced use of data discovery algorithms and other Big Data technology.

- ✦ We need to ensure efficiency and robustness of the global data federations. The EF experiments (and in future the large CF surveys) are putting in place federated systems that are consisting of a diverse set of WAN-connected storage technologies and database/data discovery systems, requiring R&D on network and systems monitoring, optimizing event I/O and data formats, etc. New functionalities needed are e.g. data archive and tape library facilities as a service within the data federation.   Storage systems need to be coupled to the networks which requires continued evolution to take advantage of network capabilities.

- ✦ Advances in data management will contain the level of effort needed for operations, making these solutions usable for the smaller IF and CF frontier experiments.

Networks

- ✦ Networks have become powerful enablers of the data intensive distributed computing that all frontiers require to solve their computational needs. Network capabilities and data access technologies improve our ability to use computational and data storage resources independent of location, over the network, enabling to marshal a much larger spectrum of resources for any given project: dedicated facilities, university and lab computing systems including HPC centers, opportunistic use, commercial clouds, etc.

- ✦ The Snowmass report makes a number of recommendations for targeted networking R&D. These are aiming to treat networks as a managed resource, and to fully include network capabilities in the computing models. This requires closer collaboration between HEP and the providers of advanced networks.

Training

- ✦ With increasing complexity of software, data management and workflow systems the need for training plays an important role. Training needs to be encouraged and supported as a continuing activity. The Snowmass report has a number of concrete recommendations in this regard, how to make better use of existing opportunities and practices.

- ✦ The need to provide young scientists with opportunities to learn computing and software skills that are marketable for non-academic jobs implies the need to adopt modern and relevant technologies.