Outline for this session

- In this session: discuss the Snowmass2013 process for Computing
  - ★ What are we supposed to deliver?
  - ★ Charges for each subgroup
  - ★ Gaps?

- outline for “report”: go through sub-group, discuss charge and work plan to prepare for CSS2013 meeting

- Come back at end of session ~2:30pm, what we’ve learned in other sessions/discussions, plan for presentation tomorrow, next steps
- subscribe to listserv CPF@fnal.gov
  - ★ send email to listserv@fnal.gov with no subject and the text SUBSCRIBE CPF Firstname Lastname

Charge To Computing

- ★ Computing has become essential to advances in experimental and many areas of theoretical physics. Research requirements in these areas have led to advances in computational capabilities. The participants in the Computing Frontier will address these issues:

- ★ 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?

Computing Technology Frontier?

- computing as tool to enable science (does it need to be a Frontier?)
  - focus on what computing/software is needed to get the science done
  - this is the main focus/deliverable of the user needs groups
- The infrastructure will project computing capabilities into the future and see how the user needs map onto the trends
  - identify where research is needed to meet computing needs
    - e.g. in OSG we identified a number of areas where computing research is needed to sustain the needs of communities for this decade
  - advances needed in computing research *and* advances and investments needed in infrastructure
  - identify where computing capabilities open physics opportunities
- approaching computing and software problems in a creative and innovative way might give leadership advantages to US groups
  - quite possibly US leadership position in LHC science is helped and enabled through computing and support for collaborative research
  - “out-innovate” Asia/Europe in computing?

Goal for CPM2012

- ★ At CPM2012, groups will present the scientific issues to be emphasized, experiments to be discussed, and strategies for implementation both in national and global terms.

- ★ for computing, define the scope of our study, and the process to prepare for CSS2013

- ★ Deliverable from CPM2012: short (few-page) summaries from each group, outlining what questions are to be addressed between now and CSS2013.

- ★ Ideally: “outline” of a report the contents of which would result from CSS2013

“Process” from CPM2012 to CSS2013

- ★ After CPM2012, subgroup conveners will formulate specific charges for their areas. These charges will clarify the physics questions to be discussed and the experiments to be given most attention. They will also detail choices made in treating areas-overlapping subgroups or linking high energy physics to other areas. In principle, these charges could evolve over the year in response to continued research, new physics results, and new proposals.
- ★ During the winter and following spring, each subgroup will hold meetings to develop and refine its ideas. We encourage groups interested in specific proposals or scientific topics to assemble white papers on their subjects.

CSS2013

- ★ CSS2013 will provide an opportunity for discussion, analysis, and arrive at conclusions for each area of the study. By the end of this meeting, each pair of conveners will have prepared an executive summary for their area, and overlap areas if necessary.
- ★ Each subgroup will produce a report answering its charge and summarizing the discussion of its area throughout the process. The ensuing electronic record, which may also contain contributed papers, will be an important resource for the community.

“Deliverable” of Our Group

- ★ Ultimately, we’re to deliver a report
- ★ current thinking: each subgroup would be invited to produce a writeup of up to 30 pages plus a ~4 page executive summary
  - ✦ some homogeneity is desirable so that various groups and subgroups do not produce documents of wildly differing length.
- ★ 30 pages per subgroup might be excessive for computing “frontier”?
  - ✦ should we aim for maybe ~ 30p user needs exp (10/exp frontier?) ~30p theory? ~30p infrastructure?
- ★ Should the documents include “strategic plans”
  - ✦ e.g. for networking, facility infrastructure developments, multi-core/fine-grain parallelism/GPUs, software libraries? can we even pull that off? Should we?
- ★ Conveners to produce few-page executive summaries of the subgroup executive summaries

“Outline” of Computing Report
Work Plan to Prepare for Report:

Walk-through

Sub-groups

CpF T1: Accelerator Science

- Estelle Cormier (Tech-X), Panagiotis Spentzouris (Fermilab); Chan Joshi (UCLA)
  - ★ Document the current computational and computing needs for supporting the optimization and design of accelerators, new accelerator concepts, and accelerator technologies.
  - ★ available tools, capabilities, computing and data resource needs
  - ★ What are the new capabilities required for for future accelerators and accelerator technologies?
  - ★ role of new computing technologies for parallel computing?
  - ★ How do existing models and algorithms w/ multi-core, multi-threaded, GPU, etc?
  - ★ need for common infrastructure such as analysis tools, data model, data portals?

CpF T2: Astrophysics/Cosmology

- Salman Habib (Chicago), Anthony Mezzacappa (ORNL); George Fuller (UCSD)

CpF T3: Lattice QCD

- Thomas Blum (UConn); Ruth Van de Water (FNAL); Don Holmgren (FNAL)
  - ★ lots of pre-existing documentation
  - ★ charge for this group to be defined
  - ★ e.g.
  - ★ “Can we imagine needing additional or more directed resources for particular areas of Theory?”

CpF T4: Perturbative QCD

- Stefan Hoeche (SLAC), Laura Reina (FSU); Markus Wobisch (Louisiana Tech)
- Provide a compact summary of current computing needs
  - ★ available tools and their respective computing and storage requirements
  - ★ exploiting these tools beyond their original scope?
- Describe advanced infrastructure needed to perform calculations required for precision SM and BSM measurements in the future
  - ★ parallel computing models? What are the benefits of inter-thread communication vs. inter- process communication? using GPUs on a large scale?
  - ★ consolidating resources? access from Grid environments? storage?
  - ★ software environment
  - ★ coordinate with other working groups, in particular the Energy Frontier working groups HE5, QCD, and HE3, Top Quark

General considerations for Experiment "User Requirements"

- ★ Is the current LHC computing model robust into the future? what is the limit e.g. if we will go to very high rates that scale beyond moore's law?
- ★ Can we imagine transformative breakthroughs
  - ✦ e.g. analysis 10X faster? (actually, CMS achieved factor ~10 in reco performance)
  - ✦ or, factor of 10 less cost for cloud computing?
  - ✦ or, all data access over network, no data placement, globally available datasets?
- ★ Are other compute models needed to support the other Frontiers, or can 'one size fit all'?
- ★ is the "distributed computing model" where one collects resources from many facilities making access ~homogeneous appropriate for the "smaller" communities of Intensity frontier experiments, or the very data intensive cosmic frontier experiments?
- ★ Are there new architectures that need to be developed to help with the DAQ/triggering etc?
- ★ Approach to resource estimates:
  should we document/recommend "best practices"?

CpF E1: Cosmic Frontier

- Andrew Connolly (U. Washington), Alex Szalay (Johns Hopkins); Salman Habib (ANL)
-

CpF E2: Energy Frontier

- Ian Fisk (Fermilab), Jim Shank (Boston)
-

CpF E3: Intensity Frontier

- Brian Rebel (Fermilab), Mayly Sanchez (Iowa St.); Stephen Wolbers (Fermilab)
-

CpF I2: Distributed Computing and Facility Infrastructures

- Ken Bloom (U.Nebraska/Lincoln), Sudip Dosanjh (LBL)
- field relies on distributed facilities, group sets out to answer questions:
  - How do computational problems map onto types of facilities? How need facilities be structured? one kind fits all? more specialization?
  - will required size of computing facilities be straightforward to achieve?or do we beed new targeted efforts? what research investments are needed?
  - will existing distributed computing models scale up, are new targeted efforts needed? what sort of research investments are needed?
  - Will national computing centers play a larger role in computations for particle physics? Will there be a role for computing on demand, i.e. cloud, facilities? If so, what will be the role of facilities at universities and particle physics-focused labs in such an environment?
  - What sort of coordination and services will be required across distributed facilities of this scale, and are new models of computing required for it?

CpF I3: Networking

- Gregory Bell (LBNL), Michael Ernst (BNL)
- work plan for group (charge):
  - describe architectures and quantify performance requirements
  - explore how HEP applications need to evolve to benefit from enhanced network capabilities (i.e., bandwidth, functionality).
  - ask how network infrastructures, capabilities, service models should be modified , so that networks can become key components of the next technology innovation cycle.
  - articulate HEP's request to network designers
- aim at a comprehensive report
  - Science and ultra-high-performance networking scenarios
  - HEP use cases that motivate need for enhanced network capabilities.
  - Attributes and features of such enhanced capabilities.
  - develop vision for opportunities for accelerated discovery made possible by the development of enhanced network capabilities.

CpF I4: Software Development, Personnel, Training

- David Brown (LBL)
- Peter Elmer (Princeton)
- Ruth Pordes (FNAL)

- my notes:
  - don’t forget code preservation, e.g. "old" code libraries for specific simulations or calculations, address questions on improving longevity of software, effort/funding needed for long-term maintenance of software, including “orphans”

CpF I5: Data Management and Storage

- Michelle Butler (NCSA),
- Richard Mount (SLAC);
- Mike Hildreth (Notre Dame)

★ my notes:

- ★ HEP has developed amazing technologies in storage and data management
  - mostly because in many areas database technologies didn't quite fit
  - we developed dcache/xrootd/rootIO before or while Google/Haddop did MapReduce and noSQL... but now we probably need to catch up
- ★ need for "innovation" or at least adoption of cutting-edge technology possibly higher than in "processing": storage, data management
- ★ scope should also include questions of data preservation

CpF I1: Computing, including special purpose hardware

- no conveners yet nominated, but do we need to? what charge?
- do we need to track computing technology/hardware trends (a la “Pasta” in the 2000’s?)
- probably need to address interface with DAQ/triggers? like GPUs, switches, very high core counts etc?
- need to coordinate with the Instrumentation frontier group

Next Steps

- coordinate/discuss within your subgroup, with related other frontiers
  - ★ come back with refined charge, a plan for addressing the charge, and possibly an outline of your section — if possible email docs
- lets take stock at 3:15pm today in this room
- we should plan for a phone meeting in the coming weeks to follow up on what we learn today and agree on the work plan for computing
  - CPF@fnal.gov listserv created -- please subscribe to stay in touch
  - also emails to conveners at allcomputingfrontier@denali.physics.indiana.edu
  - ★ make sure we recruit people we need where necessary
- need to decide what workshops and pre-meetings we need
  - ★ "user needs" subgroups should make sure to participate in pre-meetings of other frontiers, and bring back information to the whole of computing
  - ★ for meetings we can normally use phone/video, but maybe specific workshops, DOE has signaled they'd support us in this
- Please come back at 3:15!!

WH11NE (this room): more Infrastructure discussions?

- Assuming the “User Needs” groups will get together with their “frontiers”:

- maybe the “Infrastructure” people could stay in this room and talk
  - ★ about how to interact and organize cross-subgroup work
  - ★ about if we need workshops to hammer out some strategies
  - ★ etc