Computing Frontier

Lothar Bauerdick, FNAL
Steven Gottlieb, Indiana University
Co-conveners

Community Planning Meeting
Fermilab
October 11, 2012

Three Frontiers

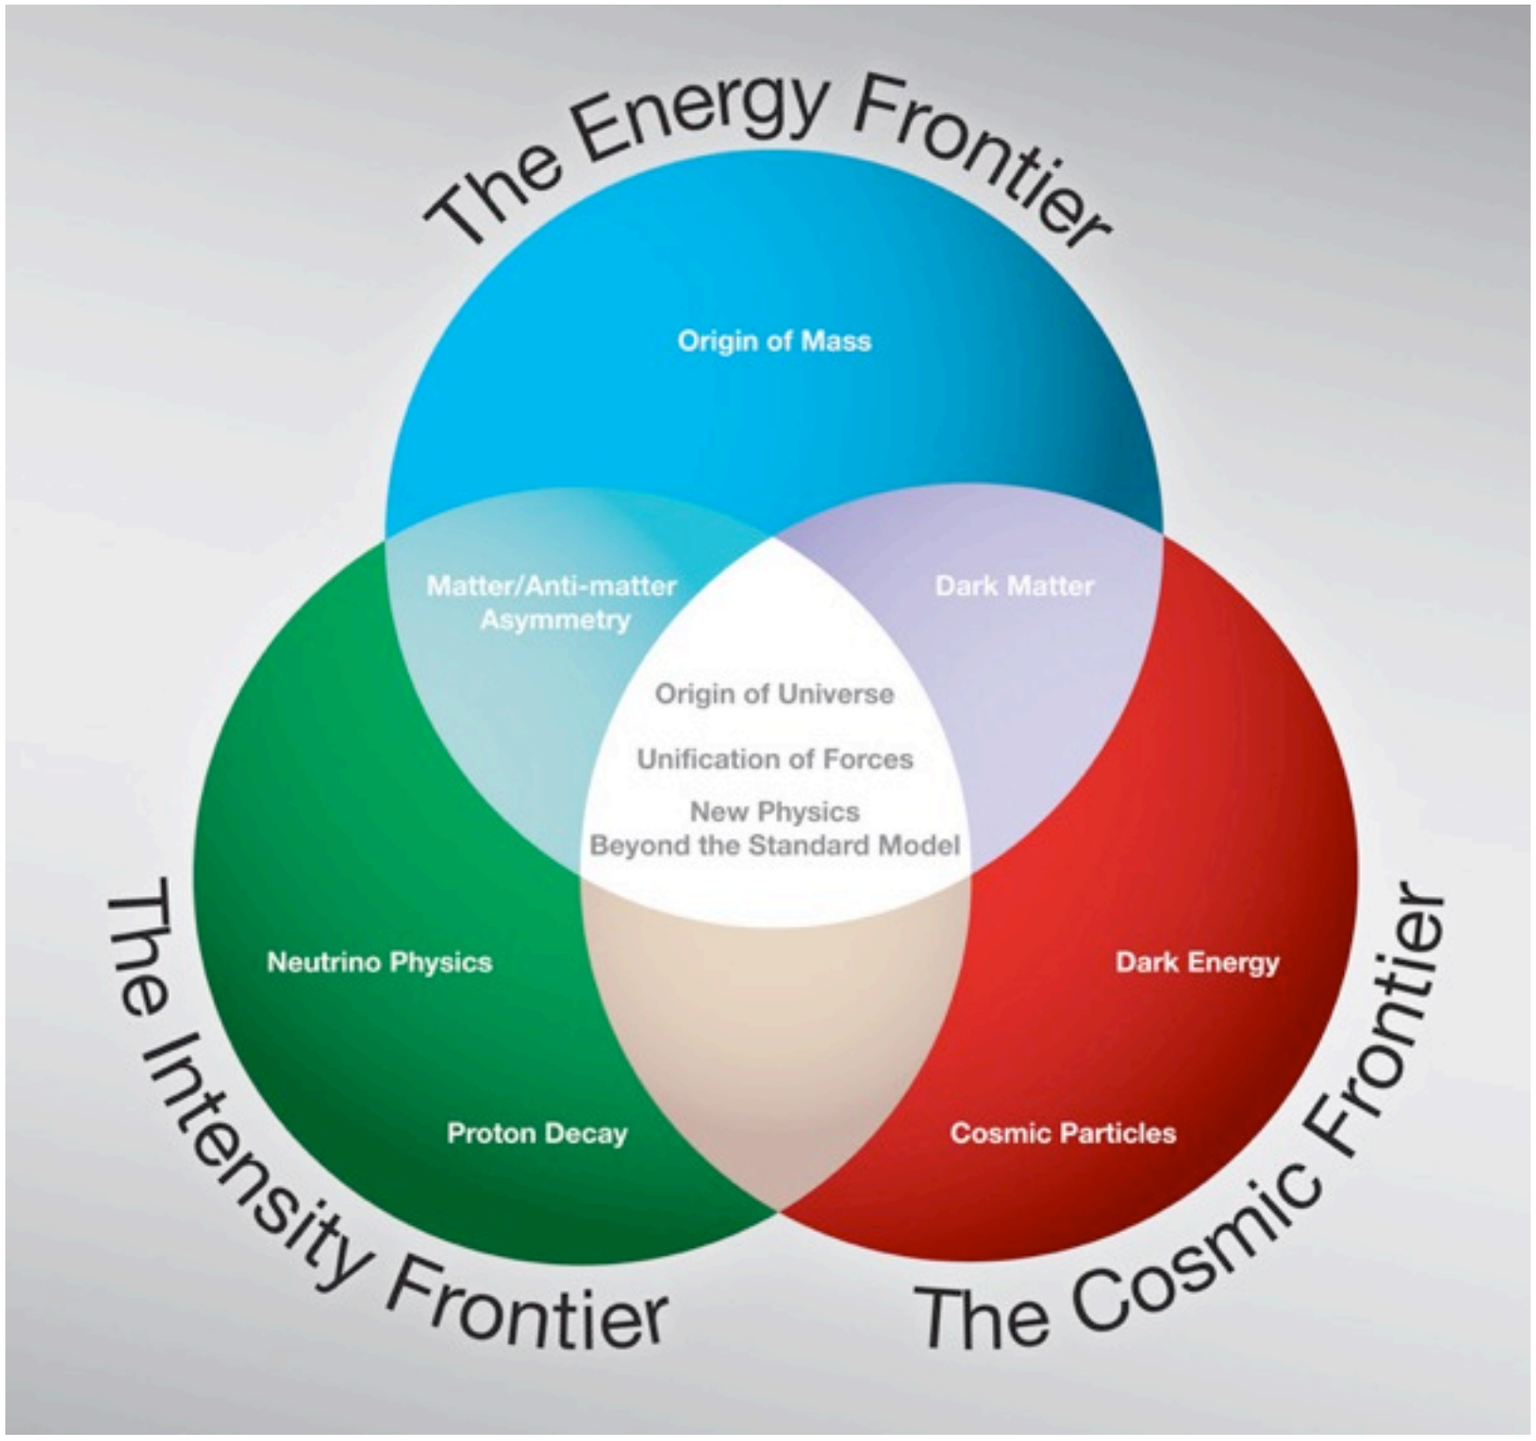

Computing Frontier

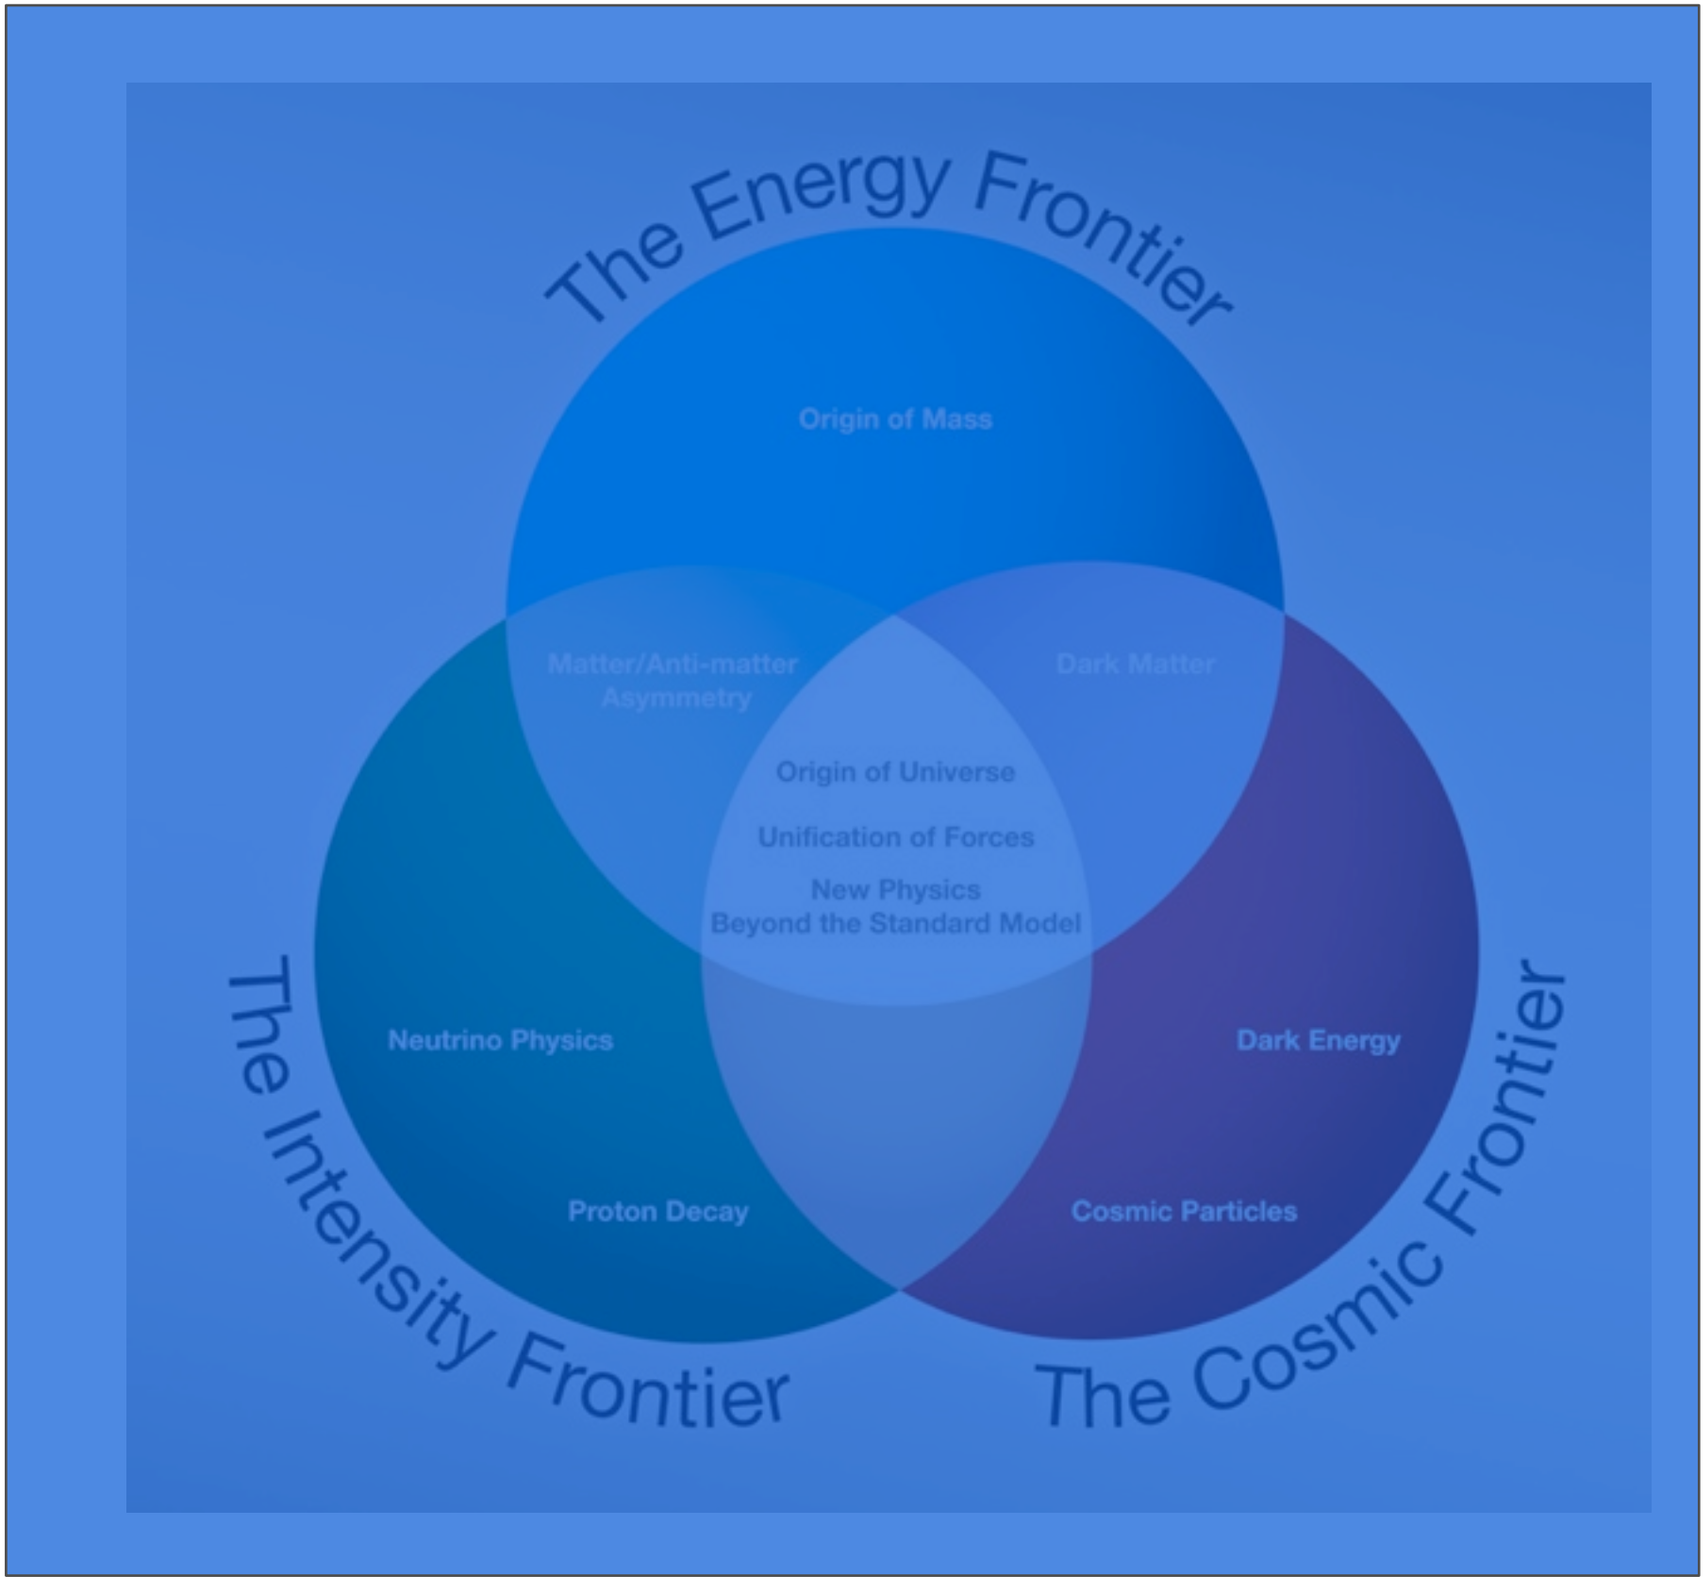

Computing Frontier Charge

- 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?

Organization I

- We have subgroups for User Needs and for Infrastructure.
- User Needs subgroups:
  - CpF E1: Cosmic Frontier
  - CpF E2: Energy Frontier
  - CpF E3: Intensity Frontier
  - CpF T1: Accelerator Science
  - CpF T2: Astrophysics and Cosmology
  - CpF T3: Lattice Field Theory
  - CpF T4: Perturbative QCD
- Each subgroup with interact with the other frontiers to assess the computing needs to advance the science.

Organization II

- Infrastructure subgroups:
  - CpF I1: Computing, including special purpose hardware
  - CpF I2: Distributed Computing and Facility Infrastructures
  - CpF I3: Networking
  - CpF I4: Software Development, Personnel, Training
  - CpF I5: Data Management and Storage
- The infrastructure groups are supposed to project computing capabilities into the future and see how the user needs map onto the trends.
- If the trends indicate that research is needed to meet some computing needs, we will point that out to the funding agencies.

Experimental Tasks

✦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 T1: Accelerator Science

- Estelle Cormier (Tech-X), Panagiotis Spentzouris (Fermilab); Chan Joshi (UCLA)
  - In the last decade, numerical modeling and simulation have played a major role in the theory, design and development of accelerators and associated technologies. This group will identify the research and development necessary to maximize the impact of advanced computation on the design of future particle accelerators and the development of new accelerator techniques and technologies for the Energy and Intensity Frontiers.
  - In the context of this workshop, our working group will address the following issues:
  - Document the current computational and computing needs for supporting the optimization and design of accelerators, new accelerator concepts, and accelerator technologies. What are the available tools, their respective modeling capabilities, performance and computing and data resource needs
  - What are the new capabilities (physics models, numerical algorithms, computational infrastructure) necessary to perform the necessary simulations for future accelerators and accelerator technologies?
  - What is the role of new computing technologies for parallel computing? How do existing models and algorithms evolve in the context of the evolution of the computing hardware (multi-core, multi-threaded, GPU, etc)?
  - Is there a need of common infrastructure such as analysis tools, data model, data portals, etc?

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)
  - The US Lattice QCD Collaboration recently wrote several white papers in preparation for the workshop "Fundamental Physics at the Intensity Frontier" in 2011. The lattice-QCD working group of the 2012 Project X Physics Study also wrote a summary of the working-group activities. These documents focus on the areas where lattice-QCD will play a key role in interpreting the results of future measurements at the intensity frontier: heavy flavors, muon anomalous magnetic moment, and nucleon matrix elements. They describe the current status of lattice-QCD calculations and prospects for the next decade.
  - Heavy flavors
  - Muon g-2
  - Nucleons
  - The US Lattice QCD Collaboration also recently organized a workshop "Lattice Meets Experiment 2011: Beyond the Standard Model". This workshop focused on lattice gauge theory calculations of theories other than QCD, such as Supersymmetry or Technicolor, and the interplay between lattice gauge theory and measurements at the LHC. The talks given can be found on the workshop website.

CpF T4: Perturbative QCD

✦ Stefan Hoeche (SLAC), Laura Reina (FSU); Markus Wobisch (Louisiana Tech)

- ✦ In the last decade the field of perturbative QCD has reached unprecedented results in providing higher-order calculations for processes with large multiplicity and/or involving several massive particles. Major developments have been the implementation of new techniques for one- and two-loop calculations, the consistent interface of NLO calculations with parton-shower Monte Carlo programs, the realization of sophisticated jet algorithms. They have directly contributed to the amazing results of the first year of run of the LHC and will be at the core of LHC physics for the years to come. To guarantee continued success, a reliable computing infrastructure must be in place.

In the context of this workshop we would like to:

1. Provide a compact summary of current computing needs for perturbative QCD.
    1. What are the available tools and their respective computing and storage requirements?
    2. What are the prospects for exploiting these tools beyond their original scope? Do increased computing and storage capacity imply increased potential? Can we facilitate (semi-)automatic reproduction of results for different input parameters (e.g. PDFs)?
2. Describe the elements of advanced infrastructure needed to perform perturbative QCD calculations required for precision SM and BSM measurements in the future
    1. What is the role of parallel computing models? What are the benefits of inter-thread communication vs. inter-process communication? What are the prospects for using GPUs on a large scale?
    2. What can be gained from consolidating resources? Will QCD theorists have access to parallel computing that is widely supported in Grid environments? Can storage needs be met? What dependences on local resources remain that might be difficult to meet?
    3. What role does the software environment play? What is the role of proprietary software? Are there limitations in programming languages, threading and inter-process communication models that should be addressed? Are new common libraries needed?

We plan to coordinate our work with other working groups, in particular the Energy Frontier working groups HE5, QCD, and HE3, Top Quark.

Infrastructure Subgroups I

CpF I1: Computing, including special purpose hardware

CpF I2: Distributed Computing and Facility Infrastructures

Ken Bloom (U.Nebraska/Lincoln), Sudip Dosanjh (LBL)

(charge in draft)

CpF I3: Networking

Gregory Bell (LBNL); Michael Ernst (BNL)

They have a long and lovely charge, which is on the next page.

- The role of high-energy physics is fundamental to the study of the most complex, subtle, and elusive natural phenomena. Such inquiry is wholly dependent on worldwide collaborations involving thousands of scientists, with workflows that include many geographically-dispersed resources, including detectors and instruments, storage assets, compute nodes, and networks.
- Emerging HEP applications, driven by the data-analysis needs of specific discovery processes, will increasingly be “network-aware,” user-focused, dynamic, distributed, and may leverage cloud and virtualization technologies. Current-generation network services and capabilities will need to be transformed, so that previously-opaque infrastructure can become visible, manageable and deeply integrated into distributed data and computing architectures.
- Given the ever-increasing computing and storage requirements of present and future HEP experiments, enhanced network capabilities may offer something more than an opportunity to boost analysis performance: they are also likely to be a critical factor in our community’s effort to provide computing at affordable costs.
- In the context of the Community Planning Meeting 2012, our working group will address the following four goals :
- We will assemble information pertaining to the process of science, describing architectures and quantifying performance requirements associated with vastly-distributed computing facilities to support data analysis of exascale datasets by thousands of physicists.
- We will explore how HEP applications need to evolve to benefit from enhanced network capabilities (i.e., bandwidth, functionality).
- We will ask how network infrastructures, capabilities, and service models should be modified, so that networks can become key components of the next technology innovation cycle.
- We will articulate HEP’s request to network designers, in order to inspire and enable enhanced application system innovations.
- The results of our analysis will be documented in a comprehensive report that discusses Findings and Recommendations for the following:
- Science and ultra-high-performance networking scenarios
  - Characteristics of HEP use cases that motivate the need for enhanced network capabilities.
  - Attributes and features of such enhanced capabilities.
  - A vision of the process of the science over the next decade, with particular attention to the opportunities for accelerated discovery made possible by the development of enhanced network capabilities.
- A provisional strategy
  - Objectives of the strategy.
  - Process and timeline for enacting the strategy.
  - Evaluation of the ability of the strategy to address the vision of extreme-scale scientific collaborations in high energy physics, with requirements for applications and network technology at the level of exascale computing.

Infrastructure Subgroups II

- CpF I4: Software Development, Personnel, Training
  - David Brown (LBL), Peter Elmer (Princeton); Ruth Pordes (FNAL)
- CpF I5: Data Management and Storage
  - Michelle Butler (NCSA), Richard Mount (SLAC); Mike Hildreth (Notre Dame)

Infrastructure Subgroups III

- The infrastructure subgroups will be predicting how the user requirements will match onto the projected computing technology at the time the computation needs to be done.
- If they feel that the natural evolution of computer technology will not meet the user needs, they will let the funding agencies know that research is necessary to accelerate meeting user needs.
- This also applies to software development, personnel and training.
- Group charges will also reflect specific group issues.

Final Remarks

- ✦ We will need extensive input from the rest of the community to identify user needs.
- ✦ Please feel free to send us advice or suggestions
  - computingfrontier@denali.physics.indiana.edu
    - for Bauerdick & Gottlieb
  - allcomputingfrontier@denali.physics.indiana.edu
    - for all subgroup conveners

- ✦ Thanks for your attention!