SEPTEMBER 13-14, 2011 • WASHINGTON DC

Computing at the Cosmic Frontier

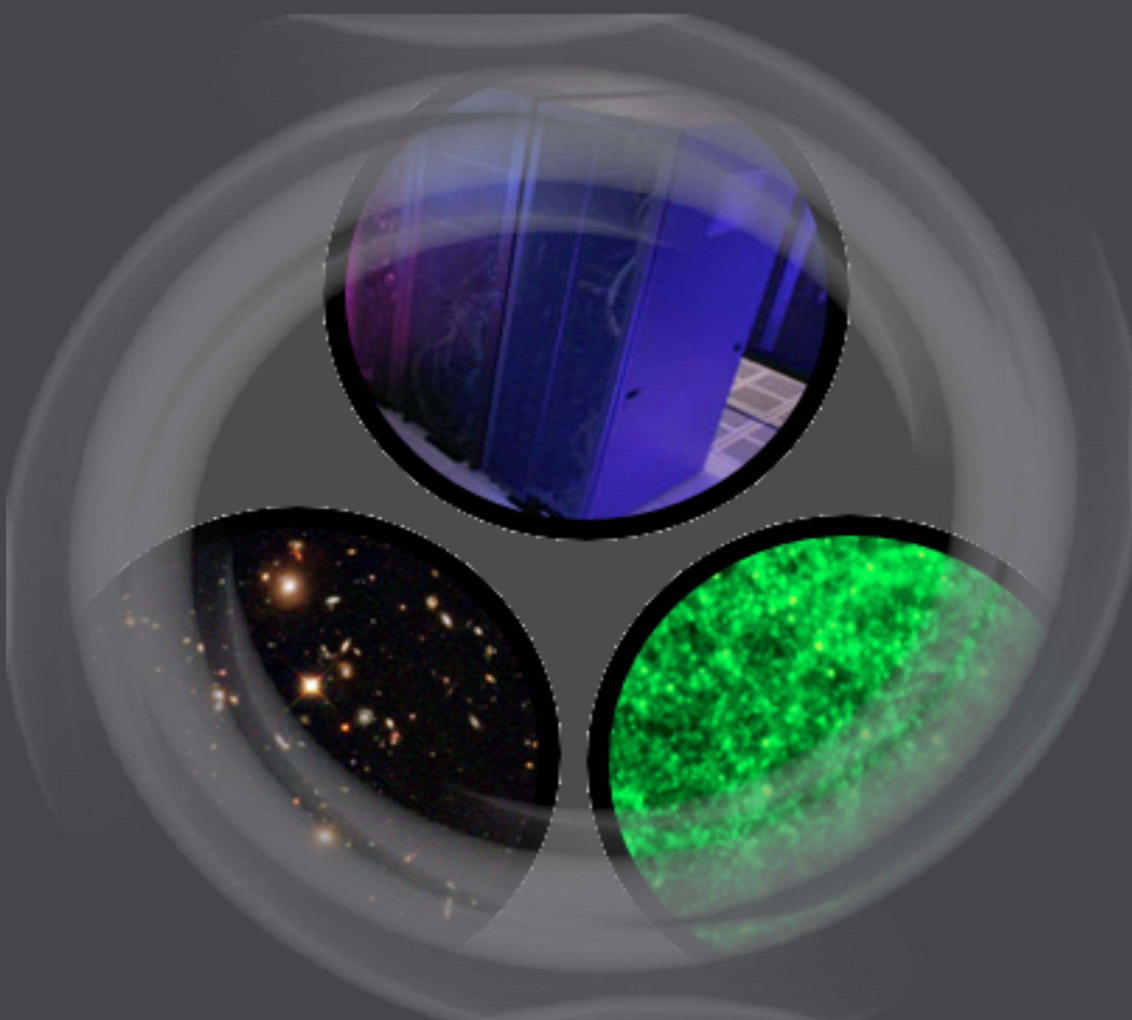

Computational Cosmology Workshop Report

Sponsored by the Department of Energy, Office of High Energy Physics

Executive Summary

The current science thrusts of DOE HEP's Cosmic Frontier program are dark matter and dark energy, high energy cosmic and gamma rays, and studies of the cosmic microwave background (CMB). The award of the 2011 Nobel Prize in physics shared by Saul Perlmutter, Brian Schmidt, and Adam Riess, for the "discovery of the accelerating expansion of the Universe through observations of distant supernovae", marks the latest high point in Cosmic Frontier research. Cosmology – an essential component of the Cosmic Frontier research portfolio – is undergoing a period characterized by remarkable discoveries and explosive growth in the power and number of interconnected cosmic probes. The research community is pursuing an exciting and challenging multi-agency-sponsored experimental program to investigate the mysterious foundational pillars of the highly successful cosmological standard model. An essential aspect of these experiments is that the Universe is part of the apparatus, blurring the line between theory and experiment.

The current accuracy with which cosmological parameters are determined is already impressive, yet in the near future, this accuracy will be enhanced very significantly as a new generation of instruments turns on to probe the sky and unearth a wealth of new information. To carry out this next generation of observations, to interpret the results, and to increase the science reach of the surveys, a large-scale computational cosmology program operating at a number of different levels is required. Because of the large scale and significant complexity of the effort, national laboratories are expected to play a central role. Driven by this imperative, the DOE HEP Laboratories have joined together in a collaborative effort to nucleate and develop a next-generation computational cosmology program.

The program would reside as a long-term core capability providing an in-house theory, modeling and simulation capability. It would be a one-point contact for scientists, projects, and programs, contributing to Cosmic Frontier 'discovery space' via new concepts and theoretical predictions, helping to catalyze the development of concepts into projects, and playing a key role in project optimization. The capabilities and tools built by members of the collaboration would be an essential component of the 'data to science' step in projects. Finally, the Lab effort would function as a major community resource enabling efficient collaborations and continuous development paths, providing a repository of best practices and 'lessons learnt' and connections across multiple projects.

The planning process has been sponsored by the Laboratories and also by DOE HEP primarily through the support of a planning meeting in September 2011 held in Washington DC. This report, a direct result of previous work, as well as the planning meeting, lays out the basic case for the computational cosmology program. It provides some background material, presents the rationale for the scientific organization and prioritization, describes current capabilities, and discusses the key aims and challenges. The report also outlines the scope of the effort, its connectivity to HEP Cosmic Frontier projects, and possible synergies with other DOE HEP activities, such as the SciDAC program, as well as linkages to the wider community. An appendix lists the key points made by the invited presenters and collaboration members in their presentations at the planning meeting. A list of acronyms is provided at the end of the report.

Table of Contents

1. Introduction

This document is a report from the planning meeting, "Computing at the Cosmic Frontier", held in Washington DC, September 13-14, 2011. This report also includes input from a collaboration workshop held at LBNL, August 18-19, 2011, as well as individual contributions. The goal of these meetings and discussions was to converge on, and to lay out a plan for a DOE HEP Laboratory-led effort in computational cosmology with the general aim of pushing the frontiers of the field, and in particular, to target the challenges most pertinent to DOE HEP's current and future experimental program. The final plan would include creation, coordination, and consolidation of resources as appropriate.

The purpose of this report is to present the conclusions in outline form, leaving out most of the technical, topic-specific details. The organization of the report is as follows. Section 2 introduces DOE HEP's Cosmic Frontier program and the key role of cosmology-related investigations within it. Section 3 discusses an example set of cosmic probes and outlines the computational requirements for each, as well as more general concerns (e.g., archiving and analyzing large datasets). Section 4 presents a short introduction to simulations of cosmological structure formation, a primary aspect of the computational cosmology program. Section 5 discusses the role of the Lab collaboration in the contexts of connecting to HEP projects and to the wider community. Section 6 describes the connectivity and relevance to the HEP Cosmic Frontier experimental program, adding further details to the material of Section 3. Section 7 gives a short overview of the current capabilities of the Lab collaboration and Section 8 lists elements of an initial collaborative program. Section 9 lists synergies with other computationally intensive research programs conducted within DOE HEP. Finally, Section 10 presents the current organization of the collaboration. The Appendix lists, in bullet form, the key points made by the presenters at the September meeting.

2. The Cosmic Frontier and Cosmology

The overall DOE HEP research program is organized on the basis of three interrelated scientific frontiers: The Energy Frontier, the Intensity Frontier, and the Cosmic Frontier [1]. The Energy Frontier concentrates on the use of accelerators at the highest energies to explore the ultimate nature of matter and spacetime; the Intensity Frontier follows a complementary path by targeting large particle number probes and high-sensitivity detectors to investigate rare events; finally, the Cosmic Frontier focuses on the detection and mapping of galactic and extra-galactic sources of radiation and particles to help reveal the fundamental nature of the Universe.

The current science thrusts of DOE HEP's Cosmic Frontier program are dark matter and dark energy, high energy cosmic and gamma rays, and studies of the cosmic microwave background (CMB). The award of the 2011 Nobel Prize in physics shared by Saul Perlmutter, Brian Schmidt, and Adam Riess, for the "discovery of the accelerating expansion of the Universe through observations of distant supernovae", marks the latest high point in Cosmic Frontier research.

As a core aspect of the Cosmic Frontier, modern cosmology is one of the most exciting areas in all of physical science. Progress over the last two decades has resulted in cementing a 'Consensus Cosmology' that, defined by only half a dozen parameters, is in excellent agreement with a host of cross-validated observations. Although the fact that a simple model can be so successful is already remarkable, three of its key ingredients – dark energy, dark matter, and inflation – point to significant future breakthroughs and discoveries in fundamental physics, as all require ingredients beyond the Standard Model of particle physics.

Perhaps the most pressing of the problems posed by cosmology is the understanding of 'dark energy', technical shorthand for the cause of the observed late-time acceleration of the Universe. Additionally, the fact that the mass content in the Universe is dominated by nonbaryonic 'dark matter' requires a compelling fundamental explanation. Inflation, the current paradigm for describing the early Universe and the source of primordial perturbations, also requires going beyond the Standard Model. Finally, cosmological observations can place strong constraints on the physics of the neutrino sector, a key area of current interest in particle physics and one of the focal points within the Intensity Frontier.

Priorities for the path forward in these areas have been laid down in various reports [1–3] and constitute a major component of DOE HEP's long-term Cosmic Frontier vision. In the case of dark energy, the first task is to test the null hypothesis (tcosmological constant) and the consistency of general relativity at the percent level of accuracy, as well as to investigate clues such as a variation in the equation of state parameter that characterizes dark energy. Measuring the geometry of the Universe and the growth of cosmic structure provides complementary insights into this problem. For dark matter, the target is to identify its particle nature via experiments at the Large Hadron Collider (LHC) and by direct and indirect detection efforts. The distribution of dark matter in the solar system, in the Galaxy, and in extragalactic space, can also be used to probe its properties and interactions. The inflation paradigm faces precision testing of its predictions (flatness, perturbation spectrum, Gaussianity) as well as searches for tensor modes (primordial gravitational waves). Connections to a more fundamental theory need to be established.

As a driver of progress along the broad front of the Cosmic Frontier research program, computing plays a central role [3]. Computational cosmology functions in three key roles: **1)** providing the direct means for cosmological discoveries that require a strong connection between theory and observations ('precision cosmology'); **2)** as an essential 'tool of discovery' in dealing with large datasets generated by complex instruments, and, **3)** as a source of high-fidelity simulations that are necessary to understand and control systematics, especially astrophysical systematics.

3. Cosmic Probes: The Role of Computation

Taking the measure of the Universe requires combining results from a variety of cosmic probes. These range from standard candle-based distance measurements (e.g., Sn Ia surveys) to statistical analyses of the galaxy distribution across a large fraction of the observable sky. The central role of simulations in cosmology is to model the formation and evolution of cosmic structures over a wide range of scales, opening up new discovery channels and providing accurate predictions for cosmological measurements.

The large-scale structure of the Universe can be measured in a variety of ways using observations across multiple wavebands. The two-point statistics of the matter distribution – specified by the correlation function or its Fourier analog, the power spectrum – has so far provided the most robust characterization of the nontrivial clustering properties of matter, arising from the dynamics of the gravitational instability in an expanding Universe. To list some of the important techniques (Cf. Fig. 1), **1)** precision CMB datasets probe the largest scales and early times, anchoring the basic paradigm, **2)** neutral hydrogen seen by 21-cm emission provides a tracer of structure (at redshifts $z > 0.5$), as does, **3)** the neutral hydrogen distribution seen as absorption features in the spectra of distant objects (Ly-α forest, $2 < z < 3$), **4)** gravitational lensing measures density inhomogeneities via the associated light deflection from background sources ($0 < z \leq 1$), **5)** galaxies themselves are biased tracers of the density field ($0 < z \leq 2$), and **6)** the abundance of massive objects (galaxy clusters) probes the tail of the density distribution ($0 < z \leq 2$).

The different cosmic probes are each associated with a major computational campaign, of varying specificity [3]. Below, we list a subset of important probes and associated projects that provide a flavor of the challenges ahead.

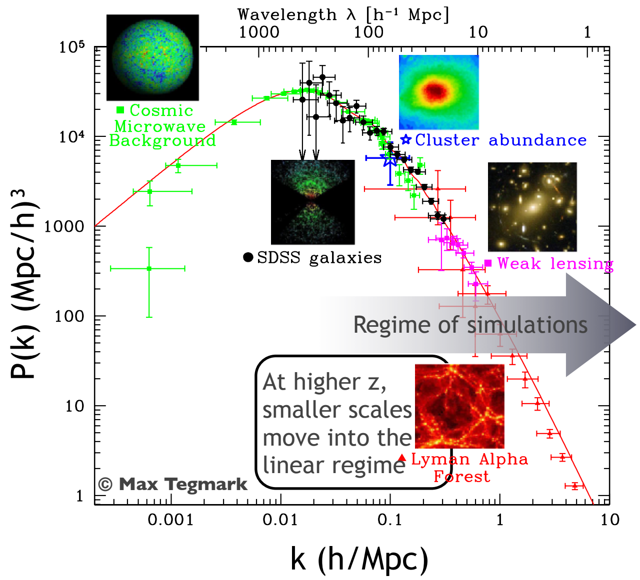

Fig. 1 Measuring the Power Spectrum

Baryon acoustic oscillations (BAO) accessed from galaxy surveys [Baryon Oscillation Spectroscopic Survey (BOSS), Dark Energy Survey (DES), Big Baryon Oscillation Spectroscopic Survey (BigBOSS), Large Synoptic Survey Telescope (LSST)] together provide a precision measurement of the geometry of the Universe at $z < 1.6$. The challenge here is the ability to run large-volume N-body simulations that can precisely determine the BAO signature in the power spectrum or the corresponding peak structure in the correlation function. At $z > 2$, the BAO signature can be extracted from the spatial statistics of the quasar Ly-α forest (BOSS, BigBOSS) – a probe of the intervening intergalactic medium (IGM). This requires running large hydrodynamics simulations to model the distribution of neutral hydrogen.

Cluster counts (DES, LSST) provide measurements of both geometry and structure growth. Here, large-volume N-body simulations are required to provide sufficient statistics, and hydrodynamic simulations are necessary to characterize observable-mass relations.

Weak gravitational lensing (DES, LSST) has multiple uses – measurements of geometry, structure growth, and cluster masses. All of these need large N-body and hydrodynamic simulations to accurately predict the mass distribution responsible for the lensing signal.

Redshift-space distortions (BOSS, BigBOSS, DES, LSST) measure the growth of structure and can test theories of modified gravity; these require large-volume N-body simulations to determine and characterize individual galaxy velocities.

Ly-α forest measurements of the matter power spectrum (BOSS, BigBOSS) are sensitive to small length scales and hence to probing the neutrino mass and thermal weakly interacting massive particle (WIMP) mass limits. Properly exploiting this probe demands hydrodynamics simulations with radiative transfer to interpret quasar spectra.

Sn Ia light curve-based techniques of distance measurement constitute one of the cornerstones of cosmology, providing essential knowledge regarding the geometry of the Universe. Simulations of Sn Ia light curves are valuable in providing an understanding of the empirically known correlation between peak magnitude and light curve duration. Although such simulations are currently not in a predictive regime (given the accuracy requirements for next generation probes of dark energy) and systematic uncertainties are dominated by instrumental calibration, progress is possible in important areas such as understanding the host/Sn-brightness correlation and the description of the intrinsic scatter in the Sn Ia Hubble diagram.

A key aspect of modern cosmology is the extraction of information from cross-correlations between different sets of observations in order to reduce parametric uncertainties and to control systematics. Consequently, several of the simulation campaigns mentioned above will have to be conducted so as to cut across the domains of the individual probes.

A significant computational task relates to the analysis of large datasets sourced by sky surveys and by simulations. The observational datasets are expected to be ~1 PB for DES and ~100 PB for LSST, while simulation data generation is constrained only by storage and I/O bandwidth and can potentially produce much larger datasets. Traditional high performance computing platforms are quite unsuited to data-centric computations, and new approaches to scalable data-intensive computing are needed. It is also apparent that managing a complex workflow with very large datasets will be a significant component of computing at the cosmic frontier. This provides an excellent area for collaboration with HEP experimentalists who have a long history of facing and overcoming data-related challenges.

Aside from the intrinsic difficulties in theoretical modeling of the individual and collective science cases and dealing with large observational datasets, there is a major added complication: Data analysis in cosmology is in fact a high-dimensional problem of statistical inference where one solves for cosmological and modeling parameters, requiring many solutions of the forward model (predictions for the observations) within a Markov chain Monte Carlo (MCMC) framework. A large number of simulation runs are also needed to determine error covariances. These requirements motivate the development of a new set of fast statistical techniques that at the same time can provide results with small, controlled errors. Going beyond this, there is the question of how results from multiple observations should be combined, given that there are nontrivial correlations between different cosmic probes. All of these problems make computational cosmology a prime candidate for exploiting recent advances in uncertainty quantification (UQ), some of which have in fact arisen from collaborations between cosmologists and statisticians [4].

4. Computational Cosmology: Modeling Cosmic Structure

By its very nature, cosmology is an observational science, where the Universe also functions as part of the apparatus, blurring the line between theory and experiment. The current flood of data from sky surveys across multiple wavebands has dramatically reduced statistical uncertainties in cosmological measurements and this trend will accelerate into the future. As a result, large-scale theoretical modeling and data analysis are required to design new cosmological probes and to interpret results from observations. The role of computation in what is now termed ‘precision cosmology’ is pervasive, complex, and crucial to the success of the entire enterprise.

Computational cosmology is primarily concerned with structure formation-based probes of the Cosmic Frontier. All of these probes are essentially measures of the dark matter-dominated density field, or quantities related to it, as demonstrated by the examples listed in the previous Section. The success of the overall approach rests on a solid first principles understanding of the basis of structure formation: initial fluctuations laid down by inflation, to be later amplified by the gravitational instability giving rise to the complex structures observed today – the growth rate of structure formation being a competition between the attraction of gravity and the expansion of space. In a standard cosmological analysis, this process is fully described by general relativity and atomic physics.

The earlier epochs of structure formation can be treated using linear perturbation theory, as applied with spectacular success to the cosmic microwave background. As the process of structure formation continues to evolve and highly overdense regions form, perturbative methods are no longer adequate. Consequently, the latter half of the history of the Universe falls in the nonlinear domain of structure formation – significantly more complex and impossible to treat without large-scale computing.

Simulations play multiple roles in cosmology. First, **1)** they enable fundamental advances in basic science by providing the most realistic physical descriptions of the known Universe. Simulations have established the current detailed picture of hierarchical structure formation where dark matter-dominated halos and sub-halos are distributed within a filamentary superstructure called the 'cosmic web' (see Fig. 2). Second, **2)** simulations are often the only 'theory' with which to compare to observational data. Matching such simulations to observations validates physical models and constitutes the key scientific return from observations; the use of simulations in a Monte-Carlo fashion is required to assess statistical significance and to determine error bars. Thus simulations are essential to extracting scientific inference from observational data. Finally, **3)** simulations play a key role in mission optimization. They are crucial for trade-off studies and in determining the impact of systematic errors; simulations can be used to design, refine, and test data processing and mission operations.

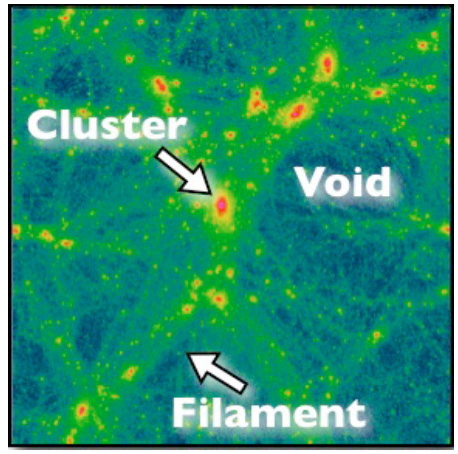

Fig. 2 The Cosmic Web

Cosmological simulations can be usefully classified into two types: gravity-only N-body simulations, and 'hydrodynamic' simulations that also incorporate gasdynamics, sub-grid modeling, and feedback effects. Because gravity dominates on large scales, and dark matter outweighs baryons by roughly a factor of five, gravity-only N-body simulations provide the bedrock on which all other techniques rest. These simulations can accurately describe matter clustering well out into the nonlinear regime, possess a wide dynamic range (Gpc to kpc, allowing coverage of survey-size volumes), have no free parameters, and can reach sub-percent accuracies. Reasonable run times can be achieved given efficient numerical implementations, which can be purely particle-based or particle/grid hybrids. Several post-processing strategies exist to incorporate additional physics on top of the basic N-body simulation, e.g., halo occupation distribution (HOD), or semi-analytic modeling (SAM) for adding galaxies to the simulations. The key shortcoming is that much of the physics of the baryonic sector cannot be treated directly.

Whenever the dynamics of baryons is important, gasdynamic, thermal, and radiative processes – among others – must be incorporated along with sub-grid modeling of processes such as star formation and local feedback mechanisms. Such simulations are substantially more complex and difficult to carry out. 'Gastrophysics' is added to gravity-only N-body simulations via either grid-based adaptive mesh refinement (AMR) solvers or via particle-based methods such as smoothed-particle hydrodynamics (SPH). Because these so-called 'hydro' simulations are computationally very expensive and have a large number of parameters, there is a need to encapsulate their results in terms of inputs into simpler modeling paradigms, either semi-analytic, or as post-processing modules added to gravity-only N-body simulations. Characterizing and controlling the robustness of these approaches will be a significant activity in the near future.

5. Role of the Lab Collaboration

From the discussion above, and as noted in the individual presentations at the planning meeting, the complexity and difficulty of the outlined tasks makes a compelling case for a computational cosmology program supported by DOE HEP and residing at the DOE HEP Labs, with appropriate university partnerships. This program would reside as a long-term core capability providing an in-house theory, modeling and simulation capability (Fig. 3), with strong connections to HEP computing, especially in the domain of high-throughput computing and large datasets. It would be a one-point contact for scientists, projects, and programs, contributing to Cosmic Frontier 'discovery space' via new concepts and theoretical predictions, helping to catalyze the development of concepts into projects, and playing a key role in project optimization. The capabilities and tools built by members of the collaboration would be an essential component of the 'data to science' step in projects. Finally, the Lab effort would function as a major community resource enabling efficient collaborations and continuous development paths, providing a repository of best practices and 'lessons learnt' and connections across multiple projects.

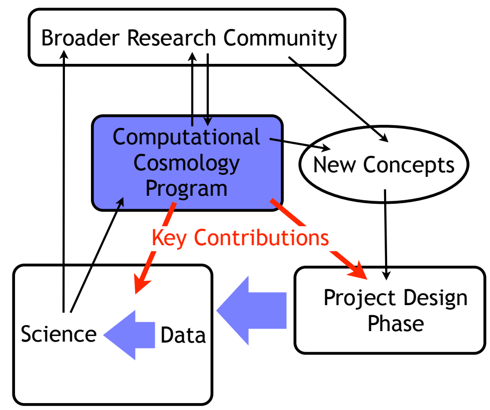

Fig. 3 Computational cosmology program research loop

The role of theory and modeling in Cosmic Frontier research is crucial at both ends of the project timeline. The initial phase requires conceptual design and optimizations followed by ever more realistic data challenges of the pipelines using input catalogs produced from simulations. Once data are in hand, extraction of cosmological information (e.g., parameter estimation) also requires a large simulation effort. It is important, however, to distinguish activities that are very project specific from those that are of a more general nature.

An example of an optical survey is shown in Fig. 4. The construction of sky catalogs proceeds from cosmological simulations augmented by HODs or SAMs. Computational work relevant to the middle stage ('sky to data') is largely project-specific, but once the data are available, cosmological simulations and large-scale data analysis tools become essential to interpret the results. Many of these simulations and tools can be relatively generic although others would have to be tuned to the specific requirements of the projects.

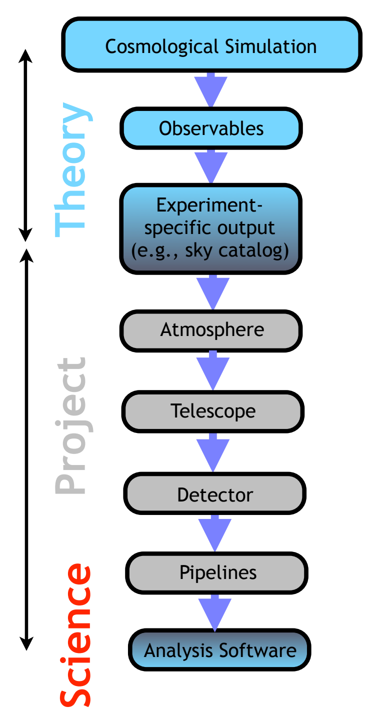

Fig. 4 Notional 'theory' and 'project' task division

The aim of the collaboration would be to develop and maintain capabilities that are relatively general in application, while being able to turn them into more specialized tools, if so desired by the projects. This being the case, the collaboration provides a fertile environment to enable consolidation of computational techniques and elements that may be common across experiments, adding efficiency to utilization of DOE resources.

To be more specific, 'first-level' data pipeline, curation, and archiving work (Level 1), followed by the production of science-ready data (Level 2) is a project responsibility. The analysis of Level 2 data is where Cosmic Frontier science results are obtained. At the same time, large-scale simulations in themselves create a need for large data storage and analysis capability in order to carry out the required extraction of science from Level 2 data. Thus there is a clear need to bring the observational and simulation data together at a Level 2 "data analysis center". Such a center would be the appropriate place for the computational cosmology collaboration to handshake with projects. Potentially, it could be set up so that a single entity handles multiple projects, reducing overheads and duplication of effort. (Such a center could be mirrored appropriately – i.e., fully or partially – at multiple sites, so it is only "single" in the restricted sense of providing a central access hub.)

Another important point to keep in mind is that as our understanding of structure formation improves, the analysis of observational data can be made more sophisticated. Thus theoretical progress can add a significant 'value-added' component to the original aims of a survey, thereby increasing its science reach. Specific examples of this include extending the range of weak lensing to smaller length scales for DES and extracting more information from the Ly-α forest for BOSS and BigBOSS.

Currently running DOE HEP-supported survey projects include BOSS, DES, and the South Pole Telescope (SPT), with BigBOSS and LSST slated for the near future. Collaboration team members have strong connectivity to the associated science, and some already play key roles within the projects. Suggestions for specific interactions with projects will be given below, with the expectation that the boundary conditions be flexible and be allowed to change with time.

6. Connections to HEP Cosmic Frontier Experiments

There is a strong history of Lab theorists collaborating with experimentalists and observers, both within and across individual institutions. Therefore, the situation in cosmology differs considerably from that in particle physics where theorists are rarely directly involved in experiments. As mentioned in Sec. 4, the collaboration expects to interact with experimental efforts at two levels (although other interactions are possible too), **1)** as a source of simulation data input, and **2)** as a resource for data analysis. Currently, there are three cosmology experiments that are funded by DOE HEP: BOSS, DES, and SPT. In all three cases, good examples of the general interactions between computational theory and projects can be found, following the discussion of task divisions in Sec. 4.

In the case of BOSS, large-scale structure simulations were run at LBNL to make mock catalogs, as well as simulations for the Ly-α forest (also in collaboration with the ANL group, then at LANL). The mock data generated from these simulations have been made public and are in use by multiple groups (e.g., by the BNL group). SPT analyses are currently being carried out at NERSC and future simulations for cross-correlation studies are being planned by the ANL group. Sets of simulations and mock catalogs for DES have been produced by the SLAC group (in collaboration with Fermilab). These are playing a valuable role in testing the data pipelines and getting the science teams ready to analyze data, which is expected to arrive on the timescale of roughly a year. Many simulations are being planned in support of science group activities within DES, and will likely be run – for the most part – using non-project computing resources. A very similar situation is likely to arise for LSST. Therefore, the division between project activities and those to be undertaken by the Lab computational cosmology collaboration require further clarification, as well as a process for resource allocation.

In order to make many of the simulation results and data products available to project science teams and the broader research community, independent resources are needed for data archiving, management, and serving. Logically, these resources should be primarily situated where the data is generated, e.g., at ALCF, NERSC, and OLCF (with secondary resources at other sites). An argument can also be made that simulation data should be mirrored at project sites that possess project-specific computing resources (e.g., a Fermilab mirror for DES simulation data). The amount of useful data associated with ‘synthetic observations’ that can eventually be produced is very large, much larger than the data volumes from the actual observations. Since, in principle, these datasets are reproducible (by re-running the simulations), strategies for optimizing use of limited storage and communication bandwidth can and will be implemented.

In the short term, we are looking at serving data in the 100s of TB, moving quickly to the PB scale as the systems mature. Resource planning in the realm of large data is therefore a significant issue. The temporary storage of data at the computing centers is probably not an insurmountable problem, but it is unlikely that they will provide independent resources for data storage, analysis, and serving without additional investments.

7. Collaboration Capabilities

The key institutions making up the collaboration are the major research facilities for HEP: Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (Fermilab), Lawrence Berkeley National Laboratory (LBNL) and SLAC National Accelerator Laboratory (SLAC). The broad science goals at each of the HEP Laboratories are well aligned with Cosmic Frontier priorities, i.e., targeting physics beyond the Standard model using cosmological observations. The theoretical efforts at the Laboratories all have a strong component in structure formation-based probes of cosmology, relating to studies of cosmic probes and the understanding of their astrophysical systematics. The balance between theory and observation (and participation in particular projects) varies from Lab to Lab, but there are significant cross-cuts for most major projects (BOSS, DES, BigBOSS, LSST).

Simulation capabilities at the Labs cover the wide range needed for a major program. At ANL, the main focus is on the Hardware/Hybrid Accelerated Cosmology Code (HACC) framework for large-scale structure simulations, now running on several petascale supercomputer architectures. HACC uses several algorithms [Particle Mesh (PM), Particle-Particle Particle Mesh (P^3M), Tree/PM] and is being extended in new directions [Particle-In-Cell (PIC) hydrodynamics, neutrinos]. The HACC framework is one of 15 Early Science projects on ANL’s new 10 PFlops BG/Q system, Mira, now being installed; ports to BG/Q prototypes are at an advanced stage of readiness. The public code Gadget is also used by the ANL group. Central to HACC is its fast ‘on the fly’ analysis toolkit which uses the framework’s data structure to auto-parallelize serial algorithms; some of these tools have been folded into ParaView, an open source parallel analysis and visualization engine. A key development direction is the link to observational programs (e.g., DES weak lensing, CMB/large-scale structure cross-correlations with SPT and DES).

BNL has Gadget running on their BG/P system as part of an external collaboration; this effort has been targeted to weak lensing studies for large-area surveys, including the generation of a large simulation suite. In addition, general and special purpose analysis codes running on local clusters are used for analyzing simulation results. A dedicated simulation analysis capability will be built for dealing with IGM simulation outputs for Ly-α forest studies. (The parent simulations will be conducted by the LBNL group.) BNL will also take early delivery of a BG/Q system; this will be a natural point of collaboration with the ANL effort.

Fermilab researchers have contributed to the development of the Adaptive Refinement Tree (ART) code, an Oct-based adaptive mesh refinement (AMR), hybrid (OpenMP+MPI) cosmological hydrodynamics code that treats N-body forces, gas dynamics, cooling, chemistry, star formation, and astrophysical feedback mechanisms. Unique to ART is a full 3D 'on the fly' radiative transfer capability, H2-based star formation, and a new treatment of the DC mode for finite simulation volumes. Planned for the future are radiation field-dependent cooling, cosmic rays, and new feedback models. ART possesses a rich analysis tool library and an initial conditions code (GIC) for generating arbitrary-shape multi-resolution initial data. Because of major local involvement in DES, and previously in the Sloan Digital Sky Survey (SDSS), there is strong connectivity to observations, especially with the Experimental Astrophysics Group ('end to end' simulations capability).

LBNL has two major codes, Nyx and TreePM. Nyx is a hybrid (OpenMP+MPI) AMR hydro and N-body code currently running on petascale systems. The TreePM code uses a parallel hybrid Tree/PM algorithm for N-body simulations and has an established history of successful large-scale simulations. In the coming years, Nyx will be the focal point of development, with the major goal of adding radiative transfer for realistic IGM studies, constraints on inflation, warm dark matter, neutrino studies, etc. Work is in progress to combine VisIt and yt and produce a state-of-the-art parallel analysis code (several other codes may be folded into this suite). With a strong history in observational programs (SDSS, BOSS, DES), the LBNL group has a broad range of codes for measuring statistics on simulations, creating mock catalogs, and weak lensing maps.

SLAC researchers have made major contributions to the development of the public AMR N-body/hydro code Enzo, including magnetohydrodynamics (MHD), radiative transfer, and many cooling, chemistry, and feedback routines. They are regular users of the public code Gadget including its application to high-resolution re-simulations. SLAC has developed the MUSIC code for generating multiscale initial conditions. The yt analysis code (now public) was incubated at SLAC; the SLAC group has also developed a number of structure analysis routines and many methods for generating mock survey catalogs. They have developed tools for testing mock surveys against survey data. Further development of these tools is planned for the future, including their integration to improve modularity. There is strong interconnectivity with DES; the SLAC group has provided DES simulations and built mock galaxy catalogs for the survey.

In terms of institutional facilities, ANL hosts a DOE Leadership Computing Facility (LCF) with a large BG/P (Intrepid, 160K cores) and a much larger BG/Q (Mira, 750K cores) arriving by mid-2012. In addition there are other machines such as CPU/GPU clusters, large analysis clusters, and visualization facilities. Substantial local computer science and applied mathematics expertise is available to the collaboration. The BNL group has access to BG/P and BG/L racks locally and will have access to 16K cores of the new BNL BG/Q. A smaller dedicated cluster is sufficient for simulation analysis. The Fermilab group has access to a local 1200-core cluster; significant expertise is available in workflow development. LBNL hosts the National Energy Research Scientific Computing Center (NERSC), the main provider of capacity supercomputing to DOE researchers. NERSC has a stable of large systems, currently led by Hopper (153K cores). In addition, there is significant computer science and applied mathematics support, including data analytics. An initial NERSC allocation has already been provided to the multi-Lab collaboration and a follow-on proposal has garnered a significantly larger allocation in 2012. SLAC researchers have access to a local 768-core system as well as a visualization center. There is substantial local expertise in the Scientific Computing group relevant to large data management and analysis.

The future computing landscape appears to be somewhat complex. Aside from the leadership-class machines available at ANL and ORNL, and also at NERSC, a variety of other high-performance architectures will also be available (e.g., heterogeneous systems based on GPUs or the Intel MIC chip, new low power many-core CPUs, 'science clouds') optimized for a variety of uses, from traditional high performance computing to data-intensive applications. It is likely that at least some portion of the current code suite would require additional development in order to take advantage of these systems.

8. Initial Program

Several collaborative science areas have been identified; these cover key aspects of the Cosmic Frontier program, including topics such as weak lensing, the Ly-α forest, clusters, galaxy surveys, neutrino constraints, inflation and non-Gaussianity, and baryonic effects. Topics such as the CMB and "ancillary" astrophysics were considered to be relevant but lower on the priority list.

To provide a specific example, weak lensing is an area of interest to all the Labs. As a technique, it is potentially very powerful – an essential component of DES and LSST – but the underlying systematics, especially the control of baryonic physics, are a serious concern. A large simulation program is essential; such a program requires running many precision N-body simulations to build fast emulators for observables and understand covariances, as well as a major effort in hydrodynamics simulations to model baryonic effects, and potentially control the associated systematic errors, especially in the context of using observational data to determine modeling parameters ('self-calibration'). A similar case can be made for the Ly-α forest, a potentially unique source of cosmological information (BAO, dark matter properties, neutrinos), but which, once again, requires a major hydrodynamics simulation effort, including radiative transfer.

To facilitate collaborative efforts, a basic scientific organization has been proposed around three overlapping themes, with corresponding institutional representatives: **1)** Cosmological Probes, **2)** Simulations, and **3)** Middleware/Tools. It is expected that inter-institutional teaming will naturally tend to follow this rough partitioning. This group will also explore connections to university groups and their research programs.

In the areas relevant to simulation and data, there was a general consensus that fruitful collaboration could initially commence on several concretely defined fronts. These include, **1)** Generation and sharing of initial conditions with the aim of producing a single code framework; **2)** Improving interoperability, e.g., ability to plug in simulation/analysis modules across codes, resolve data format issues (interchange formats for cross-code data sharing), accelerate the trend towards use of more general analysis and visualization tools (non-code specific, etc.); **3)** Develop "CosmoPack", a multi-algorithmic set of computational kernels to test and validate new computational architectures; **4)** Implement data serving of simulation outputs and products, with remote analysis capabilities to follow; **5)** Exchange expertise on code development and analysis (scaling, new physics, algorithms, new architectures, statistical methodologies, etc.); **6)** Build the "Sky Shop" (database that lists which simulation results are available where); **7)** Investigate network connectivity across the Labs and help optimize data transfer rates.

In terms of utilizing external expertise in scientific computing there are two obvious directions: **1)** ASCR-supported activities, primarily in high performance computing (with associated computer science and applied mathematics), and **2)** management of complex scientific workflows and management and analysis of large datasets, both of which fall within the domain of HEP computing expertise. Because ANL and LBNL have a significant presence in ASCR-sponsored research, these Labs have the primary responsibility to bring this leverage to bear across the entire collaboration. All of the Labs possess significant resources in HEP computing that can be brought to bear on issues of common interest. There has been a suggestion that a joint workshop be organized to help explore this direction and to determine concrete plans of action.

9. Other Synergies

The computational cosmology effort has several synergies – some potentially significant – with other DOE HEP-sponsored activities. As already mentioned in Section 6, the most obvious synergy is in the realm of large data management, analysis, and serving, where significant HEP computing expertise already exists. In terms of specific activities, one can point to the Open Science Grid (OSG), the computing efforts for the ATLAS and CMS detectors at the LHC, and the Data Preservation and Long-term Analysis (DPLTA) efforts within the HEP experimental community. In terms of computing, there are some shared algorithmic interests with the accelerator physics community (e.g., parallel Poisson solvers) as well as shared interests in areas such as visualization and parallel I/O for particle codes. The architectural complexity of the future supercomputing landscape is a key driver for sharing ideas and experiences with other HEP supercomputing communities (accelerators, lattice QCD).

10. Organization

Each institution has designated science and management leads. The current membership is as follows – ANL: Salman Habib, Harry Weerts; BNL: Anze Slosar, Morgan May; Fermilab: Scott Dodelson, Craig Hogan; LBNL: Peter Nugent, Bob Cahn; SLAC: Risa Wechsler, Roger Blandford. Salman Habib is the current spokesperson for the collaboration.

References

[1] See the HEPAP-commissioned reports at http://science.energy.gov/hep/hepap/reports/
[2] Astro2010 National Academies report, http://www.nap.edu/catalog.php?record_id=12951
[3] See the DOE ASCR and DOE HEP (jointly-sponsored) Extreme Scale Computing Report, "Challenges for Understanding the Quantum Universe and the Role of Computing at the Extreme Scale", at http://extremecomputing.labworks.org/highenergyphysics/index.stm
[4] See the section on Uncertainty Quantification and Error Analysis in the DOE ASCR and NNSA (jointly-sponsored) Extreme Scale Computing Report, "Scientific Grand Challenges in National Security: the Role of Computing at the Extreme Scale", available at http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/

Appendix: Presentation Summaries

The appendix consists of short summaries (in bullet form) of the talks given by the invited participants and collaboration members at the planning meeting, 'Computing at the Cosmic Frontier' (Washington DC, September 13-14, 2011). The individual presentations can be found at the planning meeting website:
https://twindico.hep.anl.gov/indico/conferenceDisplay.py?confId=647

Michael Turner (U Chicago): The Cosmic Frontier

1. Current status: A highly successful cosmological model accommodates all observational data and points to big discoveries ahead about the Universe and the fundamental laws that govern it.
2. Modern cosmology is characterized by being data-driven and living in the era of large instruments, both ground and space-based ('big science').
3. Today's 'consensus cosmology' is based on a set of precision measurements of cosmological parameters.
4. The three key pillars of modern precision cosmology are all mysterious: dark energy, dark matter, and inflation; all implicate new fundamental physics.
5. Cosmological computing is essential to progress – (i) in connecting theory and observations, (ii) in dealing with large datasets generated by complex instruments, and (iii) in providing simulations to understand and control systematics, especially astrophysical systematics.

Alex Szalay (Johns Hopkins): Data-Intensive Computing

1. Data is growing exponentially in all of science, changing the balance between hypothesis-driven and data-driven discovery, with more and more emphasis on data analysis aspects.
2. New data-intensive scalable architectures and new algorithms are both needed – this requires bringing together several different communities.
3. Complex cosmological simulations require an associated data-intensive analysis framework. (Simulations are a key source of 'large data'.)
4. Scalable and balanced (computation vs. communication) architectures are needed to address the data-intensive challenge.
5. The Johns Hopkins 'Data-Scope' is one approach to meeting the design goals of a data-intensive scalable computer.

Paul Messina (Argonne): Computing Future(s)

1. Evolution of HPC systems and building blocks: Technology limits are being reached, driving major changes in system architectures. Future trends include much increased concurrency, decrease in flops to memory ratio, increasingly complex memory hierarchies, increased code complexity, etc.
2. Coping with this evolution requires (among other strategies) new computational and programming models and algorithms. The key issues are accessing, moving, and storing data.
3. Guiding the evolution is important for HPC-based science. This motivates co-design as a way of matching applications to architecture using a cross-cutting design strategy across hardware and software. The aim is to avoid the twin extremes of application-driven design and technology-driven design.
4. Discussions with code teams have brought out several issues. Portability is considered highly desirable because of the large current code base; shared worries include fault management, data sizes, debugging, how to express (multi-level) parallelism.
5. One of the key functions of leadership computing facilities is to serve as test-beds for next-generation systems.

Daniel Eisenstein (Harvard): BOSS

1. BOSS is making the largest 3-dimensional map of the Universe; the primary target is the expansion history of the Universe and investigation of dark energy. BOSS will have 1.5 million galaxies and 160,000 quasars at $z > 2.1$ (resulting in a 3-dimensional map of the IGM with the Ly-α forest).
2. Cosmological opportunities include BAO, structure growth, non-Gaussianity, broad-band power spectrum, neutrino masses, –.
3. Achieving percent-level precision requires detailed modeling. Cosmological simulations are needed to interpret clustering in the quasi-linear regime and to model the covariance matrix.
4. BOSS galaxy simulations requirements include halo-level modeling and follow-on simulations to test validity of these models.
5. Simulations are needed with up to two orders of magnitude better mass resolution at two orders of magnitude larger simulation volume – this is a major challenge.

John Ruhl (Case Western): CMB

1. A new aspect of CMB observations is detection of polarization, in both E and B modes.
2. The <TT>, <EE>, and <TE> fluctuation power spectra are sourced by density perturbations at $z = 1000$. The <BB> signal is sourced by gravitational waves at $z = 1000$ at low multipoles; at high multipoles <EE> is converted into <BB> by intervening large-scale structure.
3. Current <TT> results hint at possible interesting physics (number of relativistic neutrino species > 3, etc.). The <TT> constraint on the inflationary tensor to scalar ratio is, $r < 0.17$, the best limit from <BB> is ~0.7, the eventual hope is to push this down to $r = 0.01$ with <BB>.
4. CMB lensing is sensitive to the large scale mass distribution at redshifts $0.5 < z < 4$ (relatively long tail to high z), therefore it can be sensitive to high-z dark energy. It can help break degeneracies.
5. SPTpol aims to get a neutrino mass (sum) limit of 0.15 eV and to reach $r = 0.023$ (2-σ).
6. CMB data analysis challenge needs a serious HPC effort.

Wick Haxton (UC Berkeley): Neutrinos

1. Outline of open questions in neutrino physics: New neutrino properties are discoveries that go beyond the Standard Model – an exciting interface exists between laboratory and cosmological neutrino probes.
2. Major aspects of the laboratory/cosmology neutrino program include (i) the absolute mass scale, (ii) the lepton number and mass mechanism, (iii) the neutrino hierarchy, and (iv) CP violation and measurements of θ_13.
3. Systematics in neutrino measurements are difficult to assess – this underscores the need for a unified computational team that develops and maintains a standard cosmological model and allows for consensus to emerge when discrepancies across datasets show up.
4. There are two areas where there is tension in current data: (i) 7Li abundance in BBN is puzzling (too low as measured by a factor of two), (ii) N_eff for relativistic species from CMB appears to be too high (but one need's to wait for Planck results).

Tony Tyson (UC Davis): LSST

1. LSST will carry out a wide-deep-fast survey, with precision measurements of all dark energy signatures in a single dataset (weak lensing, BAO, cluster counts, supernovae, dark energy anisotropy). Joint analyses are very important to break degeneracies and control systematics.

2. Many large dataset problems are posed by the LSST survey, e.g., there will be ~1000 deep images, each containing 10 billion galaxies.
3. Realistic image simulations are necessary to validate system performance, data processing pipelines, etc.
4. Drivers for computational challenges include cross-matching, photometric redshifts, spatial correlations, outlier detection, statistical vs. systematics errors, galaxy shapes, comparing observations to models, repeated end-to-end simulations, scalability, and so on.
5. There are important 'data to science' issues: large data, high-dimensional correlations, automated quality assessment, novel science database, –
6. Much breakthrough science occurs at the limits of surveys; computation is required to make this process more robust.

Josh Frieman (Fermilab): DES
1. DES goes after dark energy science with 4 complementary techniques (clusters, weak lensing, large-scale structure, supernovae).
2. Simulations are needed for all of these programs – (i) for clusters, focus on issues such as halo abundance, halo mass-observable relations, testing and refining methods for optical cluster finding, quantify uncertainties (ii); for weak lensing, need to quantify baryonic effects and develop methods for self-calibration.
3. The domain of DES computation covers (i) the data management system, (ii) science analysis codes, (iii) computational cosmology.
4. The roles of cosmological simulations within DES include (i) testing of data analysis and reduction software, (ii) data interpretation, (iii) controlling theory uncertainties to be less than observational systematics. Some aspects are survey-specific, while others can benefit from a broader community approach.

Rick Kessler (U Chicago): Sn Ia
1. Sn Ia have a key role in cosmology as 'standardizable' candles (brightness variation of ~1 mag can be empirically corrected to ~0.1 mag).
2. While systematic uncertainties are dominated by instrumental calibration, nevertheless improvement is possible in three other areas using supernova simulations – (i) Sn modeling, (ii) Sn evolution, and (iii) the host/Sn-brightness correlation.
3. Qualitative guidance from explosion model simulations could clarify the description of the intrinsic scatter in the Sn Hubble diagram. Sn brightness depends (empirically) on host galaxy mass – Sn simulations may be able to reveal the rationale for this.
4. Overview of Sn Ia simulations and computational demands (many CPU-hrs, large simulation datasets); complex radiation transfer problem. Comparison of predicted light curves to data is qualitatively good.

Salman Habib (Argonne): Computational Cosmology
1. The basic theory of structure formation is well understood and forms a solid basis for large-scale simulations. The current paradigm of nonlinear structure formation is based almost entirely on simulation results; quantitative probes of the 'Dark Universe' rely essentially on sophisticated simulations.
2. Examples of current results include large-volume simulation suites, with smaller-box simulations used for understanding smaller scale issues (e.g., galaxy clustering at small scales, baryonic effects on weak lensing, cluster physics, single galaxy halo simulations motivated by indirect dark matter detection).

3. There is a clear need to move on from intermediate phenomenology (simulations leading to simple fits) acceptable at the 10% accuracy level, to direct numerical approaches in order to reach the 1% level.
4. Computational cosmology is closely tied to the HPC state of the art. Simulations generate very large amounts of data, thus the ultimate vision is to combine simulation results and observational data in one analysis platform.

Martin White (UC Berkeley/LBNL): Structure Formation Probes
1. Review of structure formation probes of fundamental physics: The nonlinear domain of structure formation is of key importance. In overviewing the role of simulations, stress is laid on the importance of a sustained effort and the existence of a community of simulators and users.
2. The observational workhorse is the galaxy survey. Interpreting surveys requires precision simulations. Probes discussed included galaxy distribution statistics, weak lensing, galaxy clusters, and the IGM.
3. Galaxy power spectrum: Information can be extracted from sharp features (BAO), broad-band power (slope of initial fluctuations, non-Gaussianity, dark matter thermal properties, –), and angular dependence (tests of gravity theories).
4. Weak lensing probes distribution of luminous and dark matter in the cosmic web; it measures both geometry and structure growth. Cluster counts over time depend on geometry, growth of structure, and initial density perturbations.
5. The role of the IGM as a cosmic probe: Density fluctuations between the observer and source quasars imprint themselves as absorption features in optical spectra. These features can be used to measure neutrino mass, running of the spectral index, and also BAO (new); they are only accessible to spectroscopic surveys.
6. The computing challenges relate to high dynamic range in space and mass and to modeling complex physics. The analysis is characterized by complex workflows and large datasets. Experiments need to be designed to be robust to uncertain physics; phenomenological models are needed where ab initio modeling is insufficiently accurate.

Katrin Heitmann (Argonne): N-Body Simulations
1. The central role of simulations lies in their being the source of all of our detailed knowledge of structure formation.
2. Basics of the N-body technique were presented with code examples. Advantages (accurate, fast) and disadvantages (needs more physics at small scales) were listed. In order to match simulations to surveys, large volume/large particle number runs are needed.
3. Discovery examples from N-body simulations include: the ‘universal’ halo profile, ‘universal’ halo mass function, non-Gaussianity and bias. Some precision cosmology exploration examples are: parametric sensitivity of cosmological probes, theoretical estimates of signal to noise (BAO, weak lensing), early dark energy, etc.
4. The analysis of results includes applications to clusters, weak lensing, and small scale clustering. A new technique is the development of simulation-enabled emulators for MCMC in the nonlinear regime (‘LSSFast’).
5. Collaboration codes include the hybrid mesh/particle HACC (Hardware/Hybrid Accelerated Cosmology Code) framework from ANL and TreePM from LBNL. Adaptive mesh refinement codes are targeted primarily to hydrodynamics applications. Analysis/viz codes include ParaView and yt (both open source with Lab inputs).
6. Collaboration science highlights from N-body simulations were listed, as were simulation challenges and opportunities for the future.

Nick Gnedin (Fermilab): Hydro Simulations

1. The important role of baryons and baryonic physics was stressed: Applications include clusters, weak lensing, the Ly-α forest, improvement of CMB constraints, dark matter detection, etc.
2. The case for hydro simulations for clusters, weak lensing, Ly-α forest, CMB, dark matter detection was laid out; a view of the future was presented.
3. Cosmic gas dynamics with adaptive mesh refinement (AMR) was presented as the basis for cosmological hydrodynamical simulations.
4. The main collaboration codes are ART (UChicago/FNAL), Enzo (public, SLAC co-development), and Nyx (LBNL). All codes use adaptive mesh refinement: ART is oct-based, while Enzo and Nyx are block-structured.
5. Highlights from hydro codes were listed: Ly-α forest as a cosmological probe, physics of clusters (mass-observable scaling relations), first star masses, discovery of 'cold mode' accretion of gas onto galaxies, modeling of cosmic reionization.
6. Code discussions: Parallel performance, path to exascale, possible future benchmark problems.

Risa Wechsler (Stanford): Data Issues

1. The importance of connecting simulations to observations ('middleware') was stressed as a major challenge in its own right; it is an absolute requirement for carrying out survey science.
2. Simulations are reaching levels of dynamic range where populating them with galaxies is feasible via a statistical approach (halo occupancy distribution, abundance matching, semi-analytic galaxy modeling). Despite recent progress, more detailed and robust techniques are needed.
3. The above technique forms the basis for mock catalogs, essential for testing survey data and analysis pipelines and for helping with control of systematic errors.
4. Dark matter-dominated halos play a basic role – more work is needed in characterizing halos. Astrophysically correct halo modeling is useful to interpret dark matter experiments.

Acronym Index

ALCF ……………Argonne Leadership Computing Facility
AMR …………....Adaptive Mesh Refinement
ART ……………Adaptive Refinement Tree
BAO ……………Baryon Acoustic Oscillations
BOSS ……………Baryon Oscillation Spectroscopic Survey
BigBOSS…………...Big Baryon Oscillation Spectroscopic Survey
CMB ……………Cosmic Microwave Background
CPU ……………Central Processing Unit
DES …………….Dark Energy Survey
DPLTA …………….Data Preservation and Long-term Analysis
GPU ……………Graphics Processing Unit
HACC ……………Hardware/Hybrid Accelerated Cosmology Code
HOD ……………Halo Occupation Distribution
HPC ……………High-Performance Computing
IGM ……………Intergalactic Medium
I/O ……........…Input/Output
LHC ……………Large Hadron Collider
LSS ……………Large-Scale Structure
LSST ……...……Large Synoptic Survey Telescope
MCMC ……………Markov Chain Monte Carlo
MHD ………...…Magnetohydrodynamics
NERSC ………...…National Energy Research Scientific Computing Center
OLCF ………...…..Oak Ridge Leadership Computing Facility
OSG ………...…..Open Science Grid
PB ……….......Petabyte
PIC ……….......Particle-In-Cell
PM ……...…......Particle-Mesh
QCD …………....Quantum Chromodynamics
SAM ………...…..Semi-Analytic Modeling
SDSS …………....Sloan Digital Sky Survey
Sn Ia …………....Type Ia supernovae
SPH …………....Smoothed Particle Hydrodynamics
SPT ………...….South Pole Telescope
WIMP ………...…Weakly Interacting Massive Particle