Computational Cosmology Initiative: Task Force Report

Scott Dodelson (Co-Chair)
Don Petravick (Co-Chair)
James Amundson
Jim Annis
Nick Gnedin
Don Holmgren
Jim Kowalkowski
Andrey Kravtsov
Andreas Kronfeld

August 13, 2007

Executive Summary

- Cosmological simulations are powerful tools to understand the growth of structure in the universe. This understanding is interesting and important in its own right, but also is essential for extracting precision information about fundamental physics – such as dark energy and dark matter – from upcoming astronomical surveys.

- Groups at the University of Chicago and Fermilab have the tools in the form of state-of-the-art codes to contribute to this important area.

- The local resources for these groups are currently inadequate to run the largest, highest resolution simulations. This inadequacy is representative of a problem on a national scale. In contrast, European groups have pooled resources and currently run the most powerful cosmological simulations.

- Fermilab has the infrastructure to host a cosmological computing initiative on a national scale.

- We propose to develop a national program for computational cosmology stemming from an initial local program which includes the University of Chicago, Fermilab, and possibly Argonne. The path to strengthening the local program is clear. It is less clear how to structure a unified national collaboration, but we offer some suggestions partially based on the experience of the Lattice QCD community.

.

1 Introduction

Observations of the Cosmic Microwave Background (CMB) have pinned down the state of the universe when it was only 400,000 years old. At that time, the universe was very smooth, with different regions sharing virtually identical temperatures and densities. Perturbations around homogeneity were small but important: due to gravitational instability, the slightly overdense regions grew into the massive galaxies and galaxy clusters observed today. Since the initial conditions are known, understanding the growth of structure in our universe is a well-defined problem. It is a difficult problem because gravity is inherently nonlinear, the physics is complicated, and the rich hierarchy of structure resists a simple computational framework, but in principle it is a solvable problem.

In this context, numerical simulations play a dual role. They provide insight into the ways in which the initially smooth universe transformed into one with the majestic structure observed today. They also enable cosmologists to compare cosmological models and theories with ever-more powerful data sets obtained from large telescopes. This second function has taken on increasing importance with the discovery of dark energy. Experiments planned for the near future are challenging the computational cosmology community to produce simulations sufficiently sophisticated and accurate to extract the relevant dark sector (dark matter and dark energy) information from a universe with many layers of complexity.

Fermilab and the University of Chicago have a long history of collaboration in this general area of the large scale structure of the universe. Recently, both institutions added computational expertise to complement their leading theoretical and experimental groups. Numerical cosmology groups at the University of Chicago and Fermilab have jointly developed a suite of powerful computational tools and attracted an active group of students and postdocs. We aim to exploit this expertise to produce state-of-the-art cosmological simulations which will simultaneously push the science forward and support upcoming experiments such as the South Pole Telescope (SPT) and the Dark Energy Survey (DES).

Unfortunately, current computing resources are not sufficient to produce the large simulations that are required. In fact, this local situation is similar to that which presides at other U.S. institutions and is in stark contrast to the unified approach of European cosmologists.

We therefore propose to build machines at Fermilab which will serve two purposes: (i) offer the dedicated time necessary to run world-class cosmological simulations locally and (ii) facilitate the R&D necessary to optimize current codes so that they will run efficiently at National Centers (NERSC, leadership DOE, Teracscale NSF). The initiative will start with a local nucleus but will evolve into a national collaboration capable of producing the best simulations technology allows.

2 Science

Modern cosmological surveys have produced an explosion in the quantity and quality of observational data. The Wilkinson Microwave Anisotropy Probe (WMAP) and the Sloan Digital Sky Survey (SDSS) are nearing completion, but shortly the South Pole Telescope (SPT), Atacama Cosmology Telescope (ACT), the Dark Energy Survey (DES), Panoramic Survey Telescope & Rapid Response System (Pan-STARRS) and the Large Synoptic Survey Telescope (LSST) will begin collecting data. The ability of these surveys to test fundamental physics and to constrain the nature of the dark sector hinges on accurate predictions of the observations given an underlying cosmology. Modern numerical simulations model key observational signatures with varying degrees of confidence, limited mainly by the complex astrophysics of galaxy formation. As surveys grow in size and scale, statistical errors drop, and systematic uncertainties in theoretical modeling become dominant. For some purposes, extracting information from the existing data is already limited by current simulations; this limitation will be the dominant source of systematic uncertainty for upcoming experiments in the near future.

To reduce this systematic uncertainty, computational cosmologists must produce simulations that are large enough to provide the requisite statistics, have enough dynamic range to capture small scale complexity, and include a wide range of physical processes. This last point is necessary because, although there is more dark matter than ordinary matter, the nature of the problem requires baryons as well. Most surveys probe baryons in the form of light-emitting galaxies or quasars; even those dedicated to gravitational lensing probe the dark matter tightly clustered with the baryons. So, the large, high resolution simulations required to extract information about, e.g., dark energy, need include not only gravity but also hydrodynamics, to handle the effects of the baryonic gas. Figure 1 shows the spatial distribution of dark matter and stars in a hydro simulation of a representative volume (~ 100 Mpc) performed by D. Rudd and A. Kravtsov at the University of Chicago. The figure illustrates the range of scales involved in such simulations.

The needs of surveys aimed at understanding cosmological dark matter and dark energy go hand-in-hand with the needs of the broader astrophysical community, which is interested in solving the fundamental astrophysical problems of star and galaxy formation. Since astrophysics is an observational (rather than experimental) science, theoretical modeling and numerical simulations form a crucial part of the process of astrophysical exploration.

Cosmological simulations, therefore, are essential components of both fundamental physics and fundamental astrophysics.

2.1 Computational Cosmology for Fundamental Physics

Cosmological probes of fundamental physics that depend on high quality numerical simulations include:

- **Baryon Acoustic Oscillations (BAO)**. Oscillations of the coupled photon-baryon fluid in the early Universe imprint a "standard ruler" on the matter distribution. The length of this standard ruler can be calibrated using CMB anisotropy measurements. Measurements of the galaxy power spectrum in the transverse and line-of-sight directions can then yield precise measurements of the geometry of the Universe and its expansion rate, both of which are sensitive functions of the mean matter density in the Universe and the properties of dark energy. The use of BAO for precise cosmological constraints hinges upon our understanding of effects of nonlinear clustering of galaxies on the positions and shapes of the acoustic peaks. These effects can be studied and tested reliably only with cosmological simulations.

- **Galaxy cluster counts**. The observed cluster abundance depends on dark energy through the growth factor and the volume element. The primary method of constraining these with a cluster sample is to use an observable proxy for cluster mass to count clusters over a range of redshifts and to constrain dark energy by comparing the counts with counts predicted for different cosmologies. Successful application of such strategies requires 1) accurate theoretical calibration of the cluster mass function, implying large volume simulations and 2) testing mass proxies with simulations which include realistic physics of galaxy formation.

- **Weak gravitational lensing**. The statistics of the tiny distortions of shapes of distant galaxies by the intervening distribution of matter (the cosmic shear) depends on dark energy through the distance-redshift relation, space curvature, and the evolution of the linear growth factor. Although cosmic shear directly probes the mass distribution, which is dominated

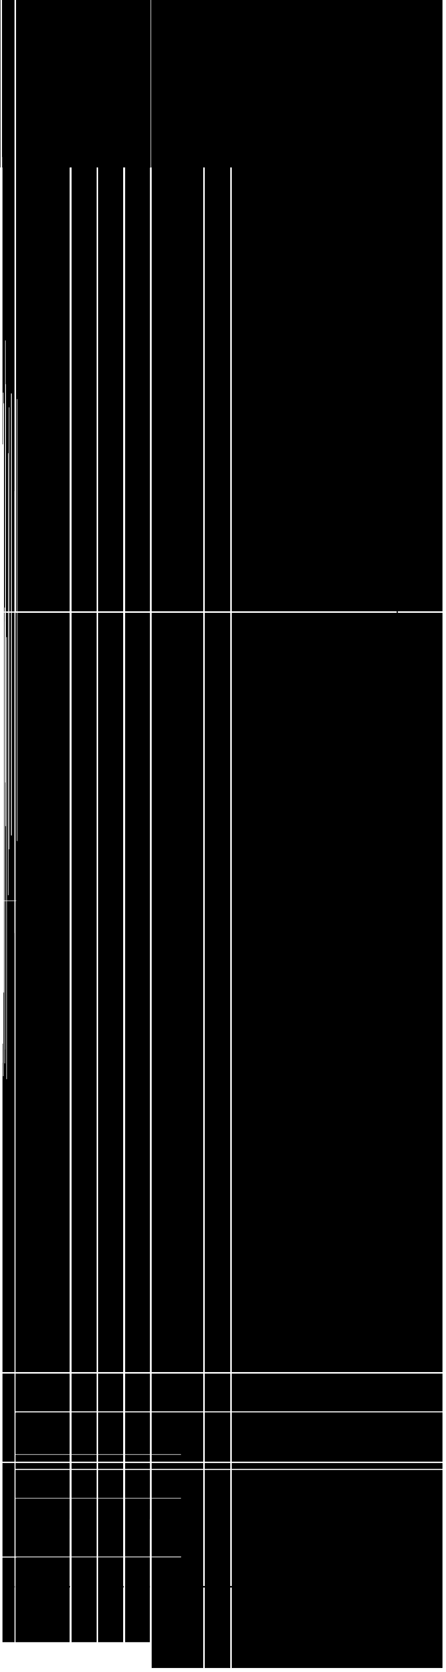

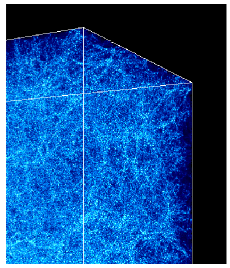

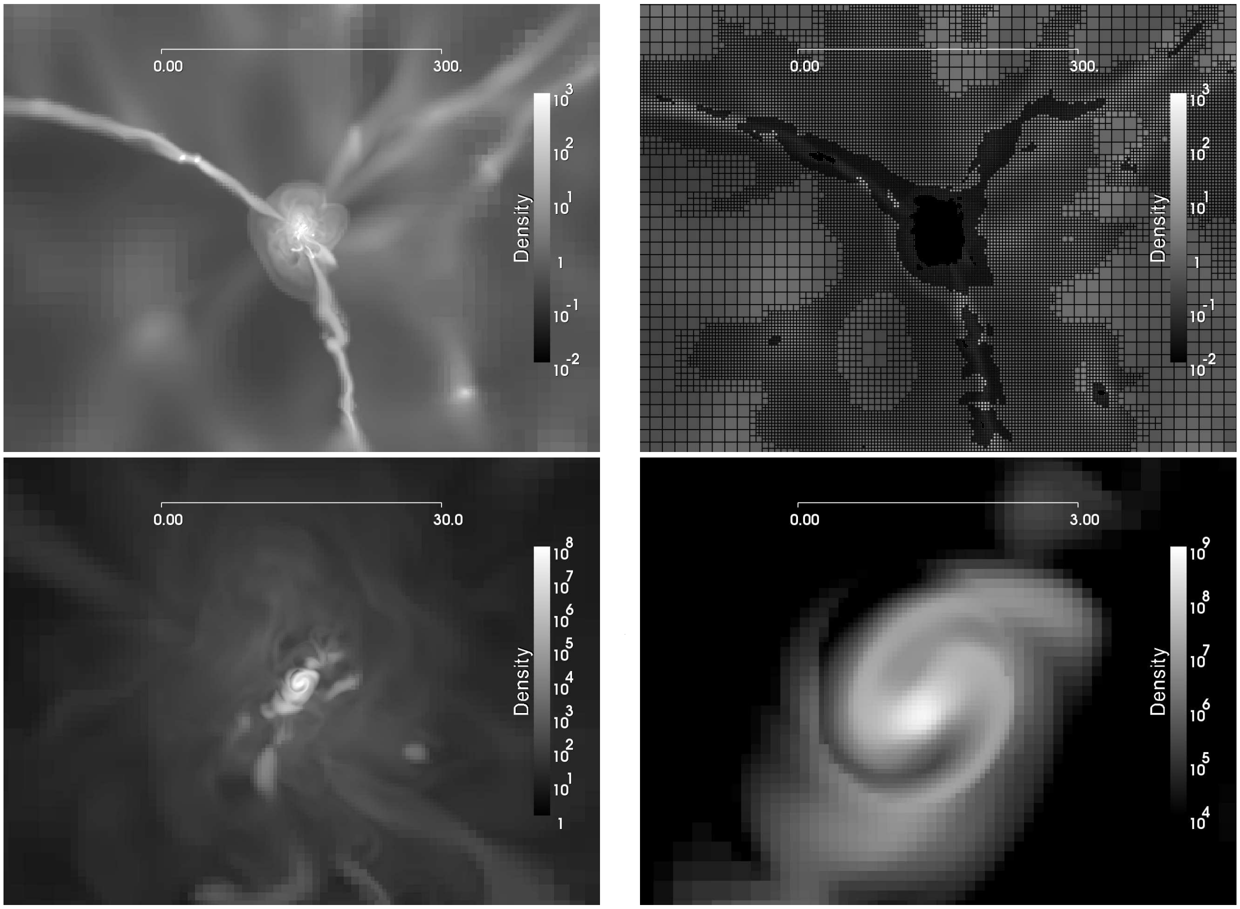

Figure 2: A powers-of-ten zoom into a simulation of galaxy formation from the UC-FNAL group. Panels show the gas density distribution in a thin slice through the center of the object from largest to smallest scale counter-clockwise from the top left panel. The top right panel shows the refinement mesh up to the fourth level overlaid on the density distribution of the top left panel. The ruler on top of each panel shows the scale of 300, 30, and 3 physical kpc. The dynamic range (the ratio of the box size to the size of the smallest cell) of the shown simulation is 32,768.

with a semi-analytic methodology. The DES plans to produce a mock survey from that simulation next year. Note that all three of the base simulations are from Europe and none of them include hydrodynamics, without which galaxies are placed onto dark matter halos using ad-hoc methods. The DES is far from unique in its need of simulations, nor in its need of simulations including the gas dynamics of the baryons.

2.2 Computational Cosmology for Fundamental Astrophysics

Understanding how structure grows in the universe is itself a fascinating scientific topic. Some of the problems which can be attacked with simulations are:

- **Galaxy formation and evolution**: Galaxies are the main building blocks in the universe; they are complex dynamical structures shaped by interactions among dark matter, gas, stars, and radiation. Only by means of ultra-high resolution cosmological simulations can realistic galaxies be modeled and understood. Simulations of isolated galaxies, galaxy mergers, and high resolution simulations of relatively small volumes have always been an important component of numerical cosmology. An example of such simulation and the range of scales involved is illustrated in Figure 2. As the problem of galaxy formation is gradually moving to the forefront of theoretical cosmology, these simulations are becoming more and more sophisticated and include a progressively wider range of new physical effects (from magnetic fields to radiative transfer and molecular hydrogen self-shielding).

- **First stars and reionization of the universe**: This type of simulation focuses on the astrophysical questions of how the first galaxies form and how they affect the rest of the universe. These questions are intricately linked to fundamental physics. For example, two parameters which are nearly degenerate in the CMB are the optical depth to reionization and the shape of the primordial perturbations set down during inflation. A precise determination of the optical depth from simulations would break these degeneracies, thereby adding to our understanding of inflation. An example of a simulation with self-consistent modeling of reionization performed by N. Gnedin at Fermilab is shown in Figure 3. These simulations are also beginning to probe the redshift interval from $z \sim 5$ to $z \sim 6$, captured by the highest redshift quasars discovered by SDSS. Simulations of reionization are also important for understanding the potential of future 21cm observations of the high-redshift Universe.

- **Supermassive black holes at the centers of galaxies**: Studies of these "cosmic monsters" are strewn with major unsolved puzzles. While numerical modeling of supermassive black holes within cosmological volumes is in its infancy, major breakthroughs in our understanding of black hole formation and evolution are expected, as the spatial resolution of modern simulations reaches sub-parsec scales.

These different scientific themes often overlap and merge into each other and the same simulation can often be used to address questions across the theme boundaries. For example, the large-scale simulations are exactly what is needed for calibration of the future dark energy surveys, but the physics of galaxy formation in these large but relatively low resolution simulations is often calibrated from high resolution simulations of small volumes. Thus, *the interests of the fundamental physics and astrophysics communities are rapidly converging*, and this convergence offers a compelling opportunity to advance the field of cosmological simulations to the mutual benefit of both communities.

2.3 Computational Cosmology in Chicagoland

The Theoretical Astrophysics Group at Fermilab has forged a close collaboration with the numerical cosmology group at the University of Chicago and the Kavli Institute of Cosmological Physics (KICP). This collaboration has already produced several state-of-the-art numerical codes for modeling the evolution of cosmic structures on a wide range of scales and at various moments in the history of the universe. The groups' primary simulation tool, the Adaptive Refinement Tree (ART) code, is an implementation of the Adaptive Mesh Refinement (AMR) technique. At present, the ART code is the most comprehensive AMR code in terms of the physical processes incorporated. The Fermilab and the University of Chicago groups have also developed several other specialized codes for special-purpose simulations and analyses.

This collaboration has made a number of important contributions to the fields of large-scale cosmic structures and galaxy formation. Kravtsov and collaborators were one of the two research groups that finally resolved the inner parts of dark matter halos [KKBP00] and solved the long-standing overmerging problems. Their work was instrumental in understanding the properties and

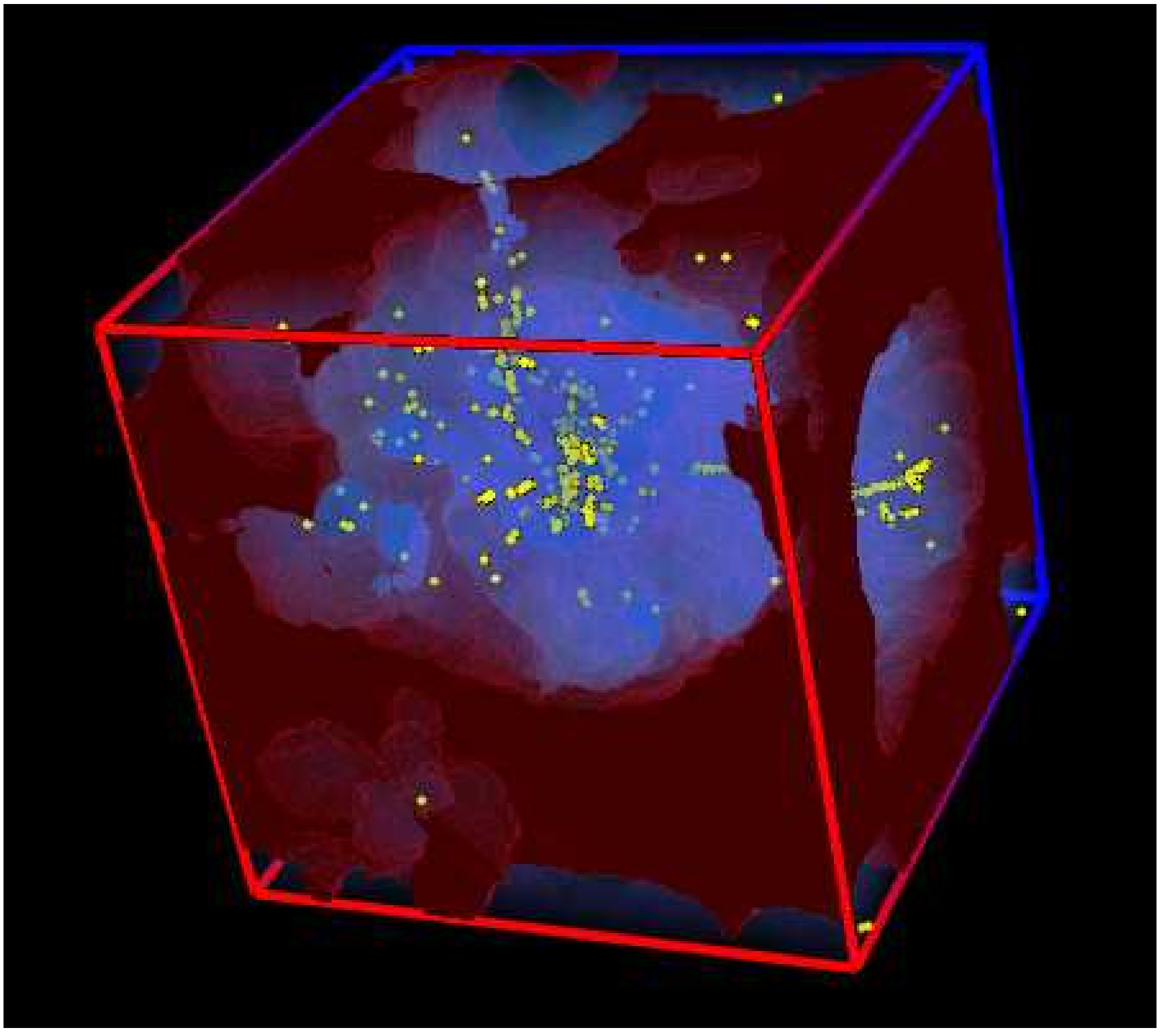

Figure 3: A simulation of cosmic reionization from the UC-FNAL group. Dots show galaxies and dark medium represents neutral gas. Regions around galaxies are ionized by ultraviolet radiation from massive stars and are devoid of neutral gas. (This figure is best viewed in color.)

spatial distribution of satellite halos (often called "sub-halos") and their relationship to observed galaxies (often nicknamed "bias") [WZBK06]. Gnedin has pioneered the numerical modeling of cosmic reionization [GF06] and made important contributions in understanding the thermal evolution of the intergalactic gas and in modeling the physical properties of the first galaxies [GHG06].

In the last five years research efforts of the UC-FNAL computational cosmology group resulted in over 70 research publications, over 100 presentations at international conferences and colloquia at top research universities world wide, and substantial public outreach activities.

3 Applications

3.1 Modern Cosmological Simulations

The scientific objectives outlined in §2 are addressed using two broad classes of numerical simulations mentioned in §2.2. Simulations of the first class model evolution of matter in a large region

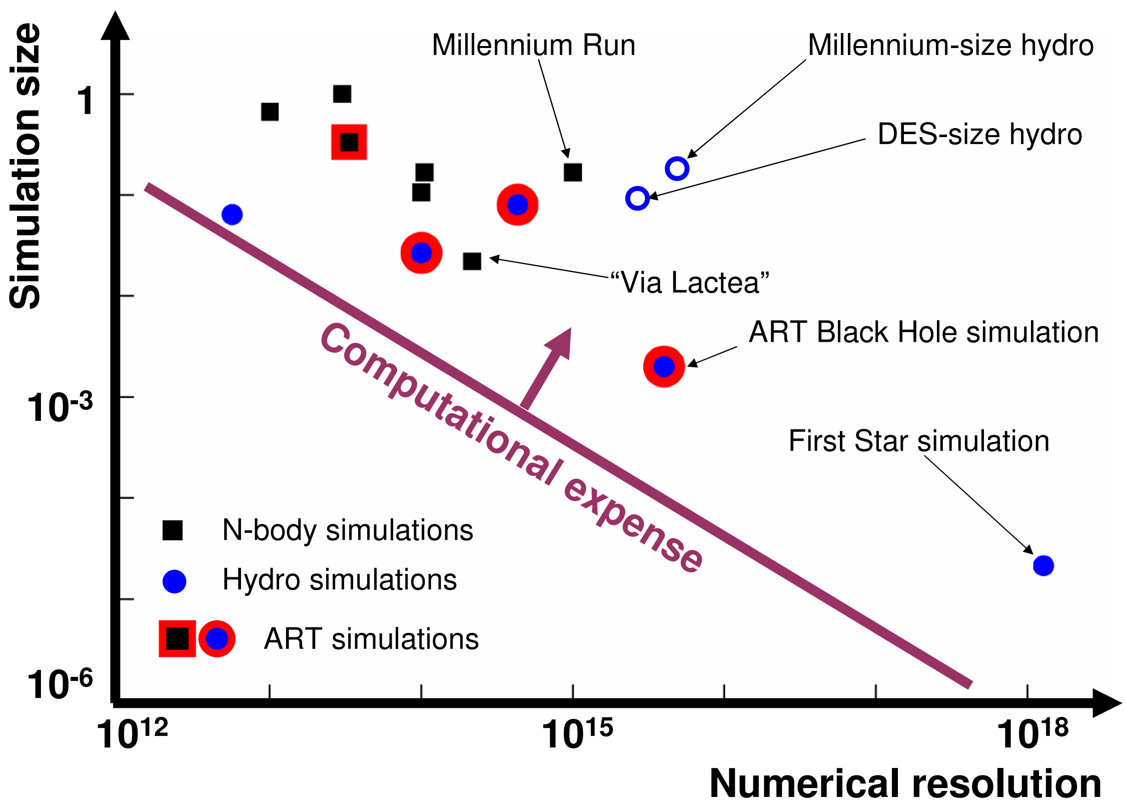

Figure 4: A comparison of resolution and simulation volume of the largest recent simulations. Black squares show N-body only simulations, while blue circles show hydro simulations with cooling and star formation. Simulations produced by UC-Fermilab group are marked by red borders. The simulation volume is measured in units of the Hubble volume, while the numerical resolution is a product of the number of particles and the spatial dynamic range.

of the universe (Fig. 4). A typical cubic volume has a size of $\sim 100 - 1000$ Mpc. Such simulations allow one to perform statistical studies of the matter distribution plus the distribution and properties of cosmological objects, such as galaxies, groups, and clusters (Figs. 1,3). Because of their large volume, these simulations are extremely useful for construction of synthetic mock datasets for large surveys, such as DES. The price is relatively low spatial and mass resolution. In addition, key physics often has to be sacrified lest the simulations become prohibitively expensive. This limits the ability of this class of simulations to address questions that require precision prediction and to explore effects of galaxy formation physics on various statistics.

The second class of simulations focuses all resolution and computational resources on a sub-volume around a single object or at most a handful of objects. Such simulations allow one to achieve very high resolution and include treatment of complex processes of galaxy formation in a relatively small volume (Fig. 2). Such simulations are therefore used to model galaxy formation and associated processes (e.g., formation of supermassive black holes) and their effects. Obvious disadvantages are that only several objects can be simulated and the resolution and the physics are not included uniformly throughout a large, statistically-representative volume.

Despite the significant differences in the nature and setup of these two classes of simulations, the required computational resources are comparable. Recent state of the art simulations in each class (e.g., the Millenium N-body simulation of a 500 Mpc volume and the "Via Lactea" N-body simulation of a single Milky Way-sized halo) cost about 300-500K CPU hours per run. Hydro simulations, which include more sophisticated treatment of baryonic physics, are even more demanding. For example, D. Rudd in Kravtsov's group at the University of Chicago ran a simulation of a volume-limited sample of galaxy clusters in support of the SPT/DES cluster survey in a volume of 340 Mpc. It required more than a million CPU hours and had to be run at the Marenostrum supercomputer in Spain.

These simulations also generate large amounts of data. The cosmological simulations evolve a three-dimensional volume in time; since the simulations are expensive, and are often used for diverse properties, a large number (typically more than 100, sometimes more than 1000) of snapshots, or time moments, need to be stored for the whole computational domain. For example, the Millenium N-body simulation produced over 20 terabytes of raw data, not counting the reduced and analysed data. Thus, data management and storage become crucial components of modern cosmological simulations.

The above examples illustrate that modern and near future state-of-the-art cosmological simulations require millions of CPU hours and tens of terabytes of storage per run. These requirements present a significant challenge to the US computational cosmology community.

3.2 Cosmological Codes

Two classes of cosmological codes produce these simulations: those which ignore baryonic matter and include only the dark matter (often called "N-body") and those that include gas dynamics (referred to as "hydro"). Since the dark matter dominates on cosmologically relevant scales, a hydro code always includes an N-body component responsible for modeling the dark matter.

The reason for the separation into two classes is purely practical: with all other parameters being equal, a hydro simulation typically takes up to 10 times longer than a pure N-body one. As a consequence, at any moment in history, the largest simulations have always been N-body, with hydro simulations being typically 10 times smaller. In addition, modern hydro simulations also include other physical processes, such as gas cooling, star formation and feedback, radiative transfer, etc. Inclusion of these processes makes a simulation even more computationally demanding - sometimes by an order of magnitude!

At present, most of the outstanding problems that can be solved with just N-body simulations have been solved. While this type of simulation will still be used in the future, for example, to improve the accuracy of prior work, it is not expected to result in any major breakthroughs in cosmology. In particular, all the outstanding science questions described above can be solved only with hydro simulations.

While several different methods for modeling the dynamics of cosmic baryons (gas and stars) have been used in the past, only two remain in use at present. The most common method is "Smooth Particle Hydrodynamics" (SPH), that uses particles to represent the gas flow. The main advantages of SPH are its simplicity and versatility, and the existence of a publicly available code named "GADGET". SPH is a Lagrangian method, and is capable of achieving moderately high spatial resolution. Lagrangian methods, however, are inherently incapable of following the fragmentation of cosmic gas. Since stars form by fragmentation of molecular clouds, SPH is of limited use in modeling galaxy formation.

The Adaptive Mesh Refinement technique offers substantial advantages compared to SPH in its

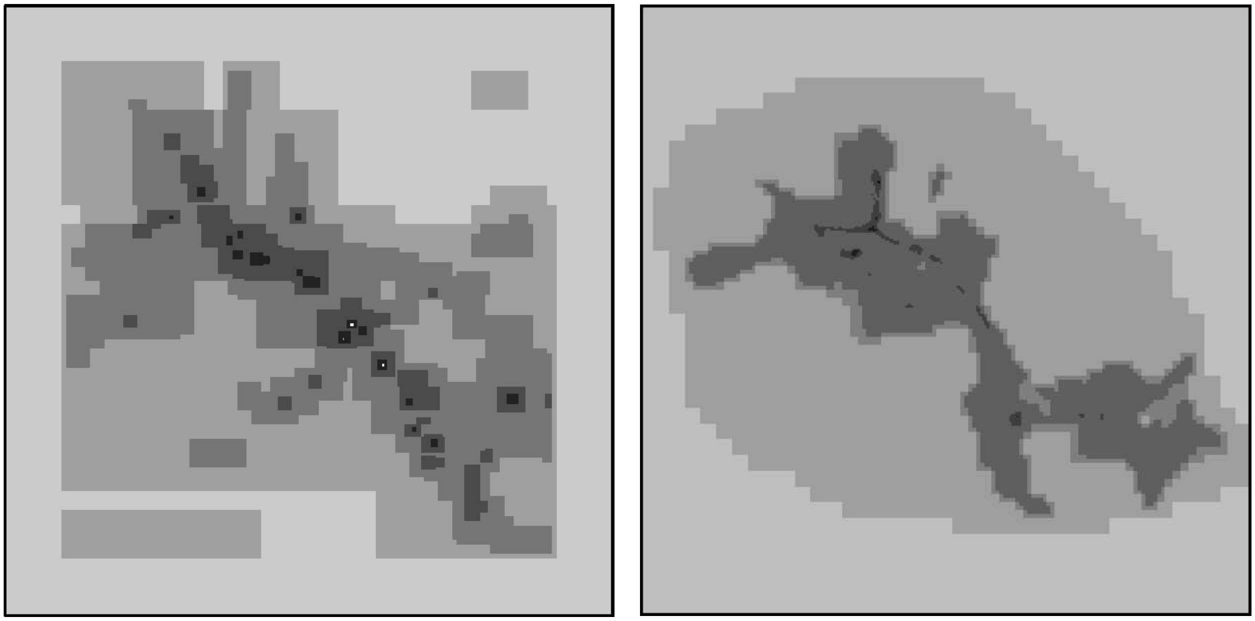

Figure 5: An example of a block-structured AMR (left) and an oct-based AMR (right).

ability to follow the fragmentation of gas down to virtually unlimited small scales. AMR, however, is a much more complicated algorithm than SPH, and is also difficult to parallelize on modern supercomputers.

At present, only two fully working cosmological AMR codes exist in the US, although several groups are actively developing new codes. The first code, named “Enzo”, was developed by G. Bryan, M. Norman, and T. Abel. It is being used by several groups in the US, including the numerical cosmology group at SLAC. The Adaptive Refinement Tree (ART) code was originally developed by A. Kravtsov and A. Klypin, with N. Gnedin joining the ART team in 2003. It is currently Chicagoland’s main simulation tool.

Two different numerical approaches exist which implement Adaptive Mesh Refinement. **Block-structured AMR** achieves high spatial resolution by inserting smaller grids (“blocks”) at places where high resolution is needed. **Oct-based AMR** instead refines on a cell-by-cell basis, subdividing each cell into 8 new cells (an “oct”). While oct-based AMR can be considered a special limit of the block-structured approach, with blocks containing only 8 cells each, in practice the two methods are very different, as they use different data structures and very different methods for distributing the computational load across a large number of processors. For example, running block-structured AMR with 8-cell blocks would incur multi-fold memory and CPU overheads and therefore be an extremely inefficient way of using computational resources. Chicagoland’s ART code is an implementation of the oct-based AMR approach.

In principle, the oct-based AMR is more flexible and more economical than the block-structured approach. In practice, however, the data structures used by oct-based AMR are highly complex, and only modest amount of effort has been invested to date in developing efficient and scalable algorithms. For this reason, the block-structured approach is currently much more popular, and several standard frameworks exist for working with massively parallel block-structured AMR applications.

3.3 Cosmological Code Support

The GADGET code was written by Dr. Volker Springel at the Max Planck Institute for Astrophysics in Germany (MPA) and is currently distributed by MPA. Dr. Springel maintains and provides user support for GADGET. Neither ART nor Enzo currently have such code and release support. Enzo had been supported for some time by the Laboratory for Computational Astrophysics at the University of California, San Diego (UCSD), but that support was terminated due to lack of funds. ART is not currently publicly available due to lack of resources to support a public release, although it is available "as is" to anyone who requests it. Both ART and Enzo are used at about half a dozen research centers around the world; GADGET is used much more widely.

4 Current Situation

The hardware and development resources currently available for computational cosmology research in the United States are considerably inferior to those in Europe or Canada. In addition to powerful, single-instution computational cosmology centers (e.g., the center at the Institute for Computational Cosmology in Durham, and Cosmology and Galaxy Formation Center at the Max-Planck Institute in Garching, Germany; see Table 1), in the last several years the DEISA (Distributed European Infrastructure for Supercomputing Applications) Initiative has provided an avenue for European computational cosmologists to secure large allocations fairly routinely. For example, in the two years since the inception of DEISA, eight large-scale allocations were made for computational cosmology projects at European supercomputer centers with an average allocation size of 1 million hours or more. For comparison, only one computational cosmology project of similar size was approved in the last three years through the DOE INCITE program — a program most similar to DEISA. Allocations at the NSF centers are even smaller. This clearly puts US researchers at a severe disadvantage. It is worth pointing out that even the DEISA's amount of CPU time investment is not sufficient to address the challenging scientific problems faced by computational cosmologists, outlined in § 2.

Although local computing facilities exist in some active U.S. computational cosmology centers, these are not directly accessible by the national computational cosmology community. The computer clusters at the Harvard's Institute for Theory and Computation, Los Alamos National Laboratory, and Princeton (see Table 1) are used almost exclusively for "in-house" calculations. The computing center at the Kavli Insitute for Particle Astrophysics and Cosmology (KIPAC) at SLAC is also used primarily by the members of the KIPAC. It allows the KIPAC scientists to be at the forefront in a wide variety of astrophysical and cosmological disiplines. However, such "in-house" centers do not address the fundamental problem of resources available for computational cosmology nationwide because they do not focus on computational cosmology alone and do not present avenues for nationwide resource allocation or collaboration. Furthermore, they do not provide a vision for a nationwide effort to improve the state of affairs nor plan to take a leadership role in developing such effort.

A computational cosmology center capable of serving the national community would require substantial computational infrastructure support. As computational capabilities increase, so do the demands on data storage and distribution. A national center would require facilities for routinely running million-CPU hour simulations and storing and distributing many terabytes of data. Software infrastructure support will also require dedicated servers, staff capable of directing user queries, relevant local software expertise, and management experience. Fermilab is already supporting such facilities for high energy physics experiment and Lattice QCD applications.

Table 1: Select "in-house" machines at computational cosmology centers around the world. The number of available processor cores is listed for each machine. Note the disparity between resources at US institutions (below the line) and Canadian/European institutions.

Place	Platform	Availability
Canadian Inst. for Theoretical Astrophysics Univ. of Toronto, Canada	2000 CPU dual Pentium 270 CPU Xeon	shared with other theory
Virgo Consortium UK, Germany	816 CPU Power4 670 CPU SparcIII	mostly cosmology
Inst. for Theory & Computation Harvard Univ., USA	316 CPU Opteron 264 CPU Athlon	shared with other theory
Stanford Linear Acc. Center Stanford, USA	360 CPU quad core AMD 128 CPU Xeon 72 CPU SGI Altrix	mostly cosmology
Los Alamos National Lab Los Alamos, USA	294 CPU Pentium4	just for cosmology
Department of Astrophysical Sciences Princeton Univ., USA	188 CPU Xeon	mostly cosmology
Department of Astronomy Univ. of Washington, USA	64 CPU Pentium	shared with other theory

5 Proposal

A concise overview of the proposed effort is shown in in Figure 6. Tables 2 and 3 complement it with the break-down of the required resources in their natural units and in the approximate cost estimates. The resources listed under FY07 already exist (or are currently being procured) and are reflected in the present Fermilab budget; any excess over FY07 is an additional, presently unavailable, resource.

5.1 Science and simulations

By the end of 2007, the proposal is to obtain local resources capable of running, storing, and processing simulations of $500-1000$ Mpc^3 volumes. These will have the dynamic range sufficient to generate mock galaxy catalogs for the Dark Energy Survey. A single simulation will be useful to the DES collaboration, but multiple runs are needed to investigate parameter space, both cosmological and astrophysical. For example, to calibrate the weak lensing convergence power spectrum, a suite of simulations will have to be run to explore dependence on cosmology and galaxy formation

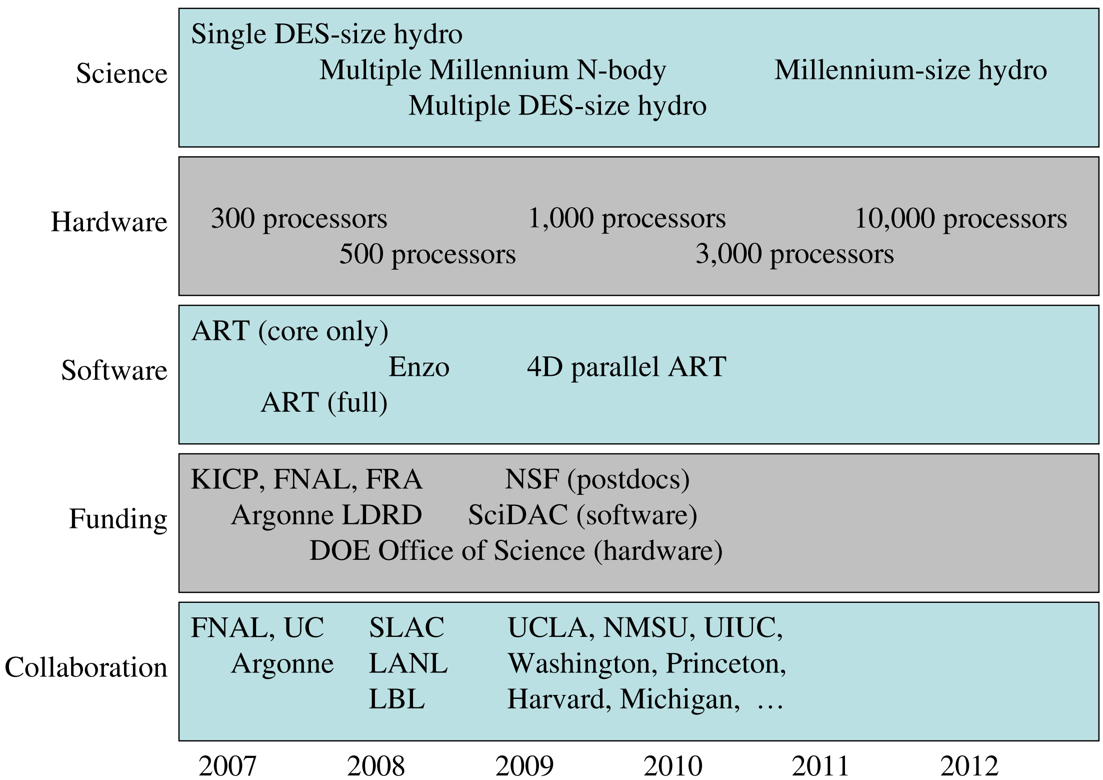

Figure 6: Timeline for cosmological computing initiative. "DES-size" is 340 Mpc, not capturing the full DES volume, but rather a simulation useful for developing a variety of tools that will be used in DES. "Hardware" refers to the local cluster; as the collaboration expands, other hardware will become available.

Table 2: Required Resources by Fiscal Year

Category	FY07	FY08	FY09	FY10	FY11	FY12
Hardware: processors (cores)	300	500	1000	3000	5000	10000
Hardware: disk storage (TBytes)	36	60	120	360	600	1200
Hardware: tape storage (PBytes)	0	0.6	1.2	3.6	6.0	12.0
Hardware: floor space (sq ft)	32	53	107	160	267	533
Hardware: power (KWatts)	29	48	97	145	242	483
Hardware: cooling (tons)	8.1	13.5	27.0	40.5	67.5	135.0
System Administration (FTE)	0.1	0.2	1.0	2.0	2.5	4.0
Software Development (FTE)	0.0	1.0	2.0	3.0	4.0	5.0
Project Management (FTE)	0.0	0.5	0.5	0.5	0.5	0.5
Postdoctoral Fellows (FTE)[a]	4.0	5.0	8.0	8.0	8.0	8.0

Table 3: Estimated Cost in K$ by Fiscal Year

Category	FY07	FY08	FY09	FY10	FY11	FY12	TOTAL
Hardware: processors[b]	$200	$233	$433	$867	$867	$2,167	$4,767
Hardware: disk storage[c]	$26	$13	$24	$73	$55	$103	$294
Hardware: tape storage[d]	$0	$75	$75	$225	$188	$375	$938
Hardware: power[e]	$15	$25	$51	$78	$134	$276	$579
Hardware: cooling[f]	$10	$17	$35	$55	$94	$193	$404
System Administration[g]	$12	$25	$130	$270	$351	$584	$1,372
Software Development	$0	$166	$346	$540	$749	$973	$2,774
Project Management	$0	$83	$87	$90	$94	$97	$451
Postdoctoral Fellows	$300	$390	$649	$675	$702	$730	$3,446
Total	$563	$1,027	$1,830	$2,873	$3,234	$5,498	$15,025

[a] This number includes 3 KICP and 1 Fermilab postdoc in FY07; a postdoc at Argonne is included in FY08; three more postdocs are added when the collaboration goes national in FY09.

[b] Computer costs are based on the cost for FY07 purchases, with additional cost for high performance network fabrics for all acquisitions in FY08-FY12. Per core costs decrease by 50% starting in FY10 to reflect predicted growth in number of cores per processor socket.

[c] Disk storage costs are based on the cost for FY07 purchases, with a cost reduction factor of 0.75 per year per terabyte of storage.

[d] Tape storage costs are based on the cost for LTO-3 media in FY06 ($50 per 400 GB), with a factor of two reduction in $/TByte every other year to reflect technology improvements.

[e] Power requirements are based on power consumption of existing computers. Costs are based on Fermilab power costs of $0.059/KW-Hr in FY07, $0.058/KW-Hr in FY08, and 3% annual inflation thereafter.

[f] Cooling costs assume that power consumption for cooling will be 70% of power consumption for computing hardware.

[g] Salaries are escalated 4% annually.

physics. In galaxy formation modeling, multiple simulations will be needed to sample a number of representative evolutionary histories and to study effects of different assumptions about star formation and feedback.

A simple implication then is that to address computational cosmology challenges in the next several years will require multi-million CPU time allocations *per project*. Local resources will be sufficient to run multiple large (Millennium-size) N-body simulations by the end of 2008. The next step, a year or so later, is to run multiple DES-size hydro simulations. The aim ultimately is to run hydro simulations on the scales which today are probed by N-body simulations.

5.2 Collaboration

The aim of this initiative is to start local and end up national. The first step is relatively straightforward, and the pathway to create a national collaboration is one of our goals in the near term.

The local collaboration begins with the Fermilab Center for Particle Astrophysics (which houses Gnedin's group) and KICP at the University of Chicago (Kravtsov's group). Argonne, also under the University of Chicago umbrella, is interested in joining the collaboration, leveraging their computational resources with a Laboratory Directed Research and Development (LDRD) proposal. The details of this are currently under negotiation, but one route would be to hire joint postdocs who can connect the computational resources at Argonne with the science and code activity at Fermilab and KICP. The next natural partners in the collaboration are the other interested national labs, Los Alamos, LBL, and SLAC. Due to our close collaboration with personnel there (Salman Habib at LANL, Tom Abel and Risa Wechsler at SLAC, Martin White and Eric Linder at LBL), we expect these additions to proceed very smoothly. Preliminary meetings with representatives from these labs are anticipated in Summer/Fall 2007.

The eventual goal of the proposed effort will be an establishment of a national collaboration or consortium in computational cosmology. There exist two possible models for such a collaboration.

In a first model a national consortium is built, in which most of the US computational cosmology centers become formal participants, in a process similar to the formation of the USQCD Collaboration for Lattice QCD. The consortium then distributes the available computational resources internally, applies coherently to the national centers for allocations, and presents unified proposals to funding agencies.

In a second model, Fermilab becomes a national center for computational cosmology, providing the resources and supporting a limited but comprehensive set of highly developed and optimized software tools. The computational resources are then distributed to different US cosmology groups on a competitive basis. This model differs in important functions from the existing national centers in that it is limited to computational cosmology and to large (million CPU hours and above) projects; medium and small size projects would still be applied for at the national centers by individual PIs.

There are pros and cons for both models. The first model follows a successful example of the USQCD Collaboration. Computational cosmology, however, is not LQCD. All LQCD computations are based on a common set of lattice configurations that can be computed by the consortium as a whole and become joint resources. However, in numerical cosmology simulations are very diverse and are performed in different physical regimes and with different algorithms. This is due to the multi-faceted and diverse nature of scientific problems in this field. This often results in little overlap between different groups. For example, the Grand Challenge Cosmology Consortium (GC^3) existed from 1993 to 1998 as one of several NSF funded Grand Challenge projects (Gnedin was a member of the GC^3 during his postdoctoral position at MIT). The GC^3 Consortium was envisioned as a close collaboration of independent research groups motivated by common broad scientific theme.

However, the difficulty of coordinating work of different groups with diverse interests, which use different tools and methods, has shown that a close collaboration on common projects is rather restrictive and is not optimal.

Thus, the second model envisions a loose collaboration. While retaining the advantange of presenting most of US computational cosmology as a coherent group to the funding agencies, it places less restriction on individual PIs. Another advantage of this model is that, by restricting the focus of the "national machine" to large projects only, it does not supplant or interfere with the national centers, which will still be used for projects that require fewer than 1,000,000 CPU hours. Of course, this model also allows tighter collaborations between groups interested in closely related problems.

At this stage it does not appear possible to select one model or another. Such a selection will require a consensus of the US computational cosmology and a negotiation process. We, therefore, envision a two-stage process of creating a national collaboration. In the first stage, the local Chicagoland resources at Fermilab are improved to be on par with the computational resources of comparable groups. Without this step, Fermilab cannot be considered an equal partner in a national collaboration.

In the second stage, a national collaboration is built either gradually or all at once. This stage would need to be preceeded by discussions and negotiations within a community. We plan to start such negotiations later in 2007, with the goal of having a national working group set up in early 2008. This group can then coordinate the negotiation process among the different US computational cosmology centers.

5.3 Hardware

Fermilab currently operates a pair of quad-socket, dual-core Opteron computers dedicated to computational cosmology simulations. During the summer of 2007, four quad-socket, dual-core Opteron systems, and 67 dual-socket, dual-core Opteron systems, will be added to the cluster, bringing the total available processor count to 308 and providing over 2.5 million processor hours per year for simulations. Further, 36 TBytes of disk storage will be added to supplement the existing five TBytes. This new cluster will bring the Fermilab computing capability to a level comparable to the dedicated resources at the other US sites listed in Table 1. Astrophysicists from the University of Chicago, Fermilab, and other collaborating institutions will use this cluster for simulations, for algorithm development, and as a prototype for the design of larger systems.

Expansion of the computational facility over the next several years to one thousand processor cores, and ultimately to a total of ∼10,000 by 2012, would raise the computing capability of the collaboration to be on par with other large international sites (see Table 1) as they continue to grow. The one thousand processor facility would provide computing capacity in excess of eight million processor hours per year for simulations. As discussed below, Fermilab is in a strong position to provide the key infrastructure for such a facility, including space, power, cooling, mass storage, data movement on- and off-site, as well as considerable expertise in all aspects of high performance computing. A large computing facility would not only give very significant throughput for large simulations, but it would also provide a platform of sufficient scale to develop algorithms and codes that would run effectively on the largest leadership-class computing facilities in the US.

In order to design the most cost effective system, the performance of codes on the cluster built in 2007 will be understood in detail, including the sensitivity to the performance of the network fabric interconnecting the individual machines. Such studies will determine, for example, whether more costly high performance network fabrics, such as Infiniband, Myrinet, or 10 Gbit Ethernet, are required to maximize the effectiveness of the system. Also of concern is the impact of the on-going industry trend towards increasing the number of processor cores per processor socket, in contrast to the prior trend of increasing the clock speed of processors. This dedicated computing facility would be able to leverage the results and experience of the other high performance computing efforts at Fermilab, such as the lattice QCD and accelerator modeling projects.

5.4 Software

The current version of ART consists of a "core" that includes N-body, hydrodynamics, equilibruim cooling, and star formation. At present, the core of ART is fully operational and scales to hundreds of processors on modern supercomputers. The full version of ART, which includes additional components such as non-equilibrium cooling, chemistry, and radiative transfer, does not yet scale to the same degree the core does on large parallel machines. The work of restructuring the full version of ART to allow for more efficient use of resources on parallel machines is expected to be complete by the end of 2007.

After a full version of ART is ready, a program for helpdesk support, code maintanence, feature enhancements, testing and validation, and release management and distribution systems will need to be established. Such a program can be developed and run by groups within Fermilab. A major reason for the wide spread use and popularity of GADGET is the uninterrupted support that the project gets from the Max Planck Institute (MPA). This aspect has also given the MPA a competitive advantage, because scientists within MPA make use of not-yet-released versions of the code. As other institutions join, the support role can be expanded to be a collaborative effort. A distributed, collaborative support model can be developed to bring together the various applications. As an example, aiding in the development and support of Enzo will be beneficial to the entire community. As the national collaboration develops and SLAC and UCSD join the partnership, the Enzo code can be added to the software resource. The futher software development work will ensure that both codes contain the same physical components and are capable of addressing a similar range of problems.

The support and maintance of ART and Enzo as public resources is a crucial component in this proposed plan. By maintaining and supporting these two codes we will serve the US cosmological simulation community with two modern, sufficiently different but comparably efficient simulation packages.

During and after this transition to a supported product, the performance characteristics of ART and Enzo can be analyzed with a goal to improve their efficency (speed and accuracy) when running on large clusters. Automated build and test environments will also be evaluated and added.

The techniques used to implement AMR are not yet conductive for petascale computing. We envision the scaling of ART to increase by an order of magnitude and possibly reaching petascale levels by implementing a novel method for parallelizing oct-based AMR codes using a new technique, which is nicknamed "4D parallelization". This algorithm is highly complex and will take a substantial effort to be implemented and validated. This upgrade presents opportunities for collaboration with other numerical computation research groups.

Locally this effort will require up to five professionals from the Computing Division at Fermilab and/or from Argonne.

5.5 Funding

Several sources of funding are potentially available, locally and within DOE. During the first phase of the timeline discussed above, the astrophysics groups at Fermilab and the University of Chicago plan to establish themselves as a local computing center for astrophysics. From the perspective of future funding, an aim of this phase is to establish the local collaboration as a force that can competently construct, operate, and exploit a computer facility designed for astrophysics. During this phase it is sensible and appropriate for local resources—from Fermilab, the University of Chicago, the Fermi Research Alliance, and, possibly, Argonne National Laboratory—to fund the effort.

Once the local effort in computational cosmology has been established, we plan to evolve toward a national collaboration. Based on the experience of similar efforts (particularly that of the national effort in USQCD Collaboration), one should distinguish between support for software development, hardware construction, and (long-term) postdoctoral positions. Because of existing funding patterns, the natural source of funding is the Office of High-Energy Physics in the DOE's Office of Science. During this period, it may make sense to begin the national collaboration with a software effort funded by the program for Scientific Discovery through Advanced Computing (SciDAC) in the Office of Science. (The first SciDAC program funded prototype hardware for lattice QCD, but not for others; the current SciDAC program funds no hardware.)

As the national effort grows in scope, it may become necessary to fund more postdoctoral researchers. In particular, a serious software effort could be more successful if five-year postdoctoral positions were made available. This is the kind of support that NSF has provided for lattice QCD, and we plan to approach them for computational cosmology.

Based on the costs of similar computing hardware, such as the tightly-coupled clusters built at Fermilab for LQCD, the approximate investment for computers, high performance networking, and storage outlined in §5.3 during 2007-2012 is $6M. This figure does not include steady-state operating costs, such as power and cooling, nor the cost of labor. The labor required for operating such a facility, including system administration, user support, procurement and deployment of hardware, project management, and supervision, would grow to at least four computer professional full time equivalents (FTE) by 2012. Note that this labor is in addition to the five software development professionals discussed in §5.4.

6 Why Fermilab?

6.1 Science

Fermilab and the University of Chicago are involved in several experiments (SDSS, DES, SNAP) which aim to measure properties of the dark sector. Proper interpretation of these experiments demands rigorous, state-of-the-art simulations, so the computing initiative outlined here is an essential component of the experimental astrophysics program in the Chicagoland area. Indeed, part of the motivation for hiring Gnedin and Kravstov was for the combined UC/Fermilab groups to become competetive in this increasingly important area of astrophysics. Now that they are here and have successfully accreted a core of students and postdocs, it is imperative to support this effort with the required resources.

6.2 Code Development

Fermilab has the experience necessary to create and maintain production quality software. This includes the collaborative code development environments for building and managing releases and large scale deplyoment of runtime environments on multiple platforms. Developers, researchers, and support staff contribute directly to coding and management of software currently running at many large experiments including CDF, D0, and CMS. The software infrastructure running at these experiments demonstrates Fermilab's expertise in large, complex software frameworks and numerical algorithms using programming languages such as C++ and C. Fermilab staff also has the knowledge necessary to diagnose diffcult performance program in scientific code including speed and accurancy.

The Fermilab Computational Physics for Accelerators group (see §A.2) is involved with development, optimization, and deployment of parallel programs, including collective operation analysis and performance analysis across different architectures. The techniques they use for the analysis of collective effects in accelerators are very similar – portions of one of the accelerator codes used by the accelerator group (IMPACT) are used in a cosmological structure formation code (MC^2). They also have a working relationship with computer science and applied mathematics groups.

The Fermilab Lattice QCD Facilities Department is involved in many aspects of code development for lattice QCD (LQCD) simulations (A.1). As part of the SciDAC-I "LQCD Computing Infrastructure Project", department personnel implemented optimized math kernels for two processor architectures, and helped design, implement, and test the community's parallel I/O application library. Work on the current SciDAC-II "LQCD Computing Infrastructure Project" includes the maintainance and extensions of the math kernels developed under SciDAC-I, optimizations for multicore processors, the design and implementation of a system to increase the reliability of tightly coupled clusters through monitoring, diagnosis, and where possible automated repair or reconfiguration, and the implementation of an automated workflow system tailored to LQCD calculations. Department personnel have considerable expertise in instrumenting parallel codes for the analysis and optimization of performance.

6.3 Data Handling

Fermilab has extensive experience specifying, integrating, and delivering both High Performance Computing (HPC) and capacity computing clusters, and has delivered successful and useable Infiniband and Myrinet systems for lattice QCD and accelerator modeling. Its HPC skills include developing application-specific, performance-predictive benchmarks and the ability to relate hardware features to application performance. The laboratory has relationships with CPU, motherboard, and interconnect vendors, which provide for early access to engineering units for benchmarks. This allows sequencing procurements to steps in performance. The laboratory normally uses "white box" vendors for large procurements at the best possible price.

Fermilab hosts archives for the Sloan Digital Sky Survey; data sets for its Run II and other experiments; and the largest Tier 1 data center for the LHC CMS experiment. Currently, the Fermilab storage systems hold over 5 petabytes of data, have ingested over 500 TB of archival data in a month, and have sourced over a petabyte in a month to offsite users and data centers. A disk and tape based storage system, with both archival and modern data serving capabilities, is available to the center.

Two independent optical fibers connect Fermilab to network exchange points in downtown Chicago, where the laboratory connects with ESnet, Internet2, international and national research networks. The fibers are controlled jointly by Fermilab and Argonne National Laboratory, and currently provide 80 gigabits per second of connectivity dedicated to Fermilab. The capacity is expandable at modest equipment cost. With current technology, the fibers represent over a terabit/sec of potential throughput.

Fermilab has extensive experience helping scientific collaborations turn networking bandwidth into application level network utility. Fermilab hosts the network integration manager for the CMS LHC experiment and routinely helps LHC sites improve their access to remotely hosted data sets. Complementary to network access and equally essentially, Fermilab provides interoperating computer security, allowing access by the US and international scientific communities to its data and systems, and is a leader in the US Open Science Grid. These skills will ensure that facilities at Fermilab are accessible and useable by the cosmological computing community.

6.4 Space and Power Capabilities

The Fermilab Computing Division makes a 8-year forecast of the laboratory's computing facility needs. The envisioned computational cosmology center is incorporated into the Fermilab Facility plan. Compared to the laboratory's total needs, the computational cosmology center is of modest size.

The art of providing computing facilities is not static. While computational power increases exponentially for a constant budget, there has been no exponential increase in the efficiency of providing computing floor space. A facility requirement is to deal with computing equipment that has unprecedented power density. Fermilab continuously upgrades its computing rooms, seeking to match its physical facilities to the changing requirements of computing systems.

6.5 Support for Machines

Fermilab operates substantial computing facilities that support experimental and theoretical programs. The support for these machines includes housing, power, and cooling; operating system support (Fermilab personnel have joint responsibility with partners at CERN for the Scientific Linux distribution); computer security, including all necessary planning following NIST guidelines as mandated by Federal policy, as well as active vulnerability scanning, incident detection, incident management, and automated software patching; data management and protection, including backup facilities; hardware and software problem tracking and resolution ("helpdesk" facility).

Fermilab currently operates:

- 1250 systems (over 3000 processor cores) tightly coupled via Myrinet and Infiniband for Lattice Quantum ChromoDynamics simulations
- 1000 systems (4280 processor cores), including purchases to be completed in 2007, for the US CMS Tier 1 and analysis facilities
- 2420 systems (11960 processor cores), including purchases to be completed in 2007, on the FermiGrid, a facility used for offline reconstruction and other processing pipelines for the Run II experiments (CDF, D0), the Sloan Digital Sky Survey, the Pierre Auger Observatory, and a number of "fixed target" particle physics experiments
- Over 300 fileservers, including purchases to be completed in 2007, providing in excess of 4.5 petabytes of online disk storage

6.6 Project Management and Execution Experience

Fermilab personnel have a long history of providing the Department of Energy with project management and project execution expertise on projects both local and remote to the laboratory. Local multiyear projects include the construction, operation, and upgrades of the CDF and D0 detectors. Fermilab provided the project management, manufacturing, testing, and delivery of key superconducting magnets for the Large Hadron Collider (LHC) accelerator. In astrophysics, Fermilab personnel serve as project managers for the Sloan Digital Sky Survey, the Pierre Auger Observatory, the Cryogenic Dark Matter Search, and the Dark Energy Survey. In high performance computing, Fermilab personnel serve as project managers for the US Lattice QCD Computing Facilities project and the SciDAC-2 Accelerator Modeling Project.

References

[DES05] J. Annis et al. (DES Collaboration), *Dark Energy Studies: Challenges to Computational Cosmology*, (2005), astro-ph/0510194.

[GF06] N. Y. Gnedin and X.-H. Fan, *Cosmic Reionization Redux*, Astrophys. J. **648**, 1 (2006), astro-ph/0603794.

[GHG06] A. Gayler Harford and N. Y. Gnedin, *Star Formation in a Cosmological Simulation of Reionization*, (2006), astro-ph/0610057.

[KKBP00] A. Klypin, A. V. Kravtsov, J. Bullock and J. Primack, *Resolving the Structure of Cold Dark Matter Halos II*, (2000), astro-ph/0006343.

[S+05] V. Springel et al., *Simulating the joint evolution of quasars, galaxies and their large-scale distribution*, Nature **435**, 629–636 (2005), astro-ph/0504097.

[WZBK06] R. H. Wechsler, A. R. Zentner, J. S. Bullock and A. V. Kravtsov, *The Dependence of Halo Clustering on Halo Formation History, Concentration, and Occupation*, Astrophys. J. **652**, 71–84 (2006), astro-ph/0512416.

A Projects with Relevant Expertise and Goals

A.1 Lattice Quantum ChromoDynamics

Lattice Quantum Chromodynamics (LQCD), the numerical simulation of Quantum Chromodynamics (QCD) requires very large computing resources with properties similar to those required for computational cosmology. Physicists run LQCD simulations on large commercial supercomputers, purpose built supercomputers, and tightly coupled clusters. Two classes of computing are done on these machines. In the first class, a simulation of the QCD vacuum is carried out, and a time series of configurations, representative samples of the vacuum, are generated and archived. Ensembles with varying lattice spacing and quark masses are generated. This class of computing requires machines capable of sustaining on the order of a TFlop/s for days or weeks at a time. The second class, the analysis phase, uses hundreds of archived configurations from such ensembles to calculate quantities of physical interest. These analysis computations also require large floating point capabilities; however, the calculations performed on individual configurations are independent of each other. While configuration ensemble generation requires single machines of as large computing capability as practical, analysis computing can rely on multiple machines each capable of sustaining at least on the order of tenths of TFlop/s for days at a time. Both classes of computations are memory bandwidth and communications intensive.

The LQCD community in the US is organized into a national collaboration, the USQCD Collaboration. Members of USQCD compete for time on the various national supercomputing centers, as well as use various dedicated machines that have been funded by the Department of Energy that are located at several national laboratories. These dedicated machines include the QCDOC supercomputer at Brookhaven National Lab (BNL) and several clusters at Fermilab and the Thomas Jefferson National Accelerator Facility (JLab). The total capacity of these dedicated machines as of May 2007 was approximately 11 TFlop/s sustained, roughly equivalent to 30 TFlop/s on the HPL Linpack ("Top500") benchmark or to roughly 38 TFlop/s peak performance.

The Department of Energy has also funded the development of LQCD software via the SciDAC (Scientific Discovery through Advanced Computing) LQCD Computing Infrastructure Projects.

The first of these, running from 2001-2006, resulted in the creation of substantial software infrasture, including libraries and applications, that enable LQCD simulations to run on essentially any supercomputing hardware available worldwide. The libraries abstract hardware details such as communications networks, optimize the performance of mathematical routines on a number of specific microprocessor architectures, and provide a data parallel computational framework that facilitates the development of new algorithms and physics applications. Significant performance improvements of legacy codes have resulted from the application of the software infrastructure developed during the first SciDAC grant. This grant also funded the development of prototype tightly coupled clusters that were designed to minimize the ratio of price to performance on LQCD codes. The SciDAC-2 LQCD Project, running from 2006-2011, will extend the software infrastructure to include optimizations supporting emerging multi-core microprocessors, automated workflow frameworks, visualization support, and toolkits to support algorithm development and analysis.

The USQCD Collaboration was formed from a very diverse theoretical community in high energy (particle) physics and nuclear physics, as well as physicists working on areas beyond the standard model (such as supersymmetry). Fermilab physicists currently hold major roles on the governing committees.

A.2 Fermilab Computational Accelerator Physics

Modern computational accelerator physics includes modeling of interactions of beam particles with accelerator components (single-particle optics) as well as interactions of beam particles with themselves and other particles (collective effects). Collective effects of interest include the interaction of beam particles with themselves (space charge), with residual electrons in a beam pipe (electron cloud) and with the other beam in a collider (beam-beam). All of these effects require solving field equations in a manner very similar to that used in computational cosmology.

While single-particle optics modeling can be carried out on desktop machines, the computation demands for collective effects are far greater. Simulating beams with over 10^{12} particles utilizing particle-in-cell techniques requires the computational power of supercomputers and/or tightly-coupled clusters. The Fermilab Computational Accelerator Physics group is experienced in developing on and working with such large, parallel machines. The group's expertise includes not only accelerator physics, but also parallel algorithm design and optimization and infrastructure management.

The accelerator group is funded by a DOE SciDAC2 grant to the COMPASS collaboration, which includes computational accelerator physicists from every relevant national laboratory as well as some universities. The collaboration PI, Panagiotis Spentzouris, is also the head of the Fermilab group. COMPASS also includes collaborators from computer science and applied mathematics COMPASS evolved from a similar, smaller, group funded from 2001-2006 under the first SciDAC program.

A.3 Potential Collaborators from Computer Science and Applied Mathematics

A national effort for computational cosmology should involve computer scientists and applied mathematicians in addition to physicists. This type of interdisciplinary collaboration has proven effective in the Lattice QCD and Computational Accelerator Physics projects. We list here a few potential collaborators from computer science and mathematics representing the general sort of collaborations we envisage.

The PETSc group is at Argonne. It is part of the SciDAC TOPS project. PETSc is a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations. It employs the MPI standard for all message-passing communication.

Chombo comes from the Applied Numerical Algorithms Group from LBNL. The Chombo package provides a set of tools for implementing finite difference methods for the solution of partial differential equations on block-structured adaptively refined rectangular grids. Both elliptic and time-dependent modules are included. Support for parallel platforms and standardized self-describing file formats are included.

Chombo provides a distributed infrastructure for parallel calculations over block-structured, adaptively refined grids. Chombo's design is uniquely flexible and accessible. Any collaborator will be able to develop parallel applications to solve the partial differential equations in which she is interested with far shorter development times than would be possible without the infrastructure. Very careful design and documentation allows said collaborator to enter the software at many levels. She will be able to use Chombo to investigate deep technical issues of adaptive mesh refinement algorithms or to simply adapt the example applications to solve different scientific problems.

The Argonne Modeling, Simulation and Visualization Group is a likely candidate for collaboration on advanced visualization. Argonne MS&V Group designs and develops multidisciplinary, integrated modeling and simulation software. They work closely with research sponsors and domain scientists to bridge the gap between concrete problems and modeling, simulation and visualization solutions.