I. Introduction

HPC and storage resources are requested to support two new research grants to the PI in the area of early cosmological evolution: NSF AST-0708960, "*Toward Petascale Simulations of the Early Cosmic Evolution*," and NASA ATFP07-0087 "*Molecular Cloud-Resolving Simulations of First Light Objects: Predictions for JWST*", as well as ongoing research in other areas of cosmological structure formation. These includes simulations of the intergalactic medium under a pre-existing NSF grant AST-0507717 "*Precision Measurements of the Intergalactic Medium: Simulations and Observations*" (D. Tytler PI, M. Norman co-I). The PI has discontinued his use of DOE machines; therefore this proposal constitutes the only request for supercomputer resources by the PI this year.

II. Project Descriptions and SU Request

This year we will put into production our new implicitly coupled cosmological radiative transfer and ionization kinetics solver for the Enzo code (Reynolds et al. 2008). This will be used for several projects described below, including the first self-consistent radiation hydrodynamic simulation of the epoch of Helium II reionization. The large computational grid required (2048^3), and the large cost per timestep for the implicit solves translates into a very large SU request on TACC Ranger. Because of this and other time-consuming simulations, we are making our first 2-year request. Also new this year, we are requesting resources for a project that will be carried out using the LCA's newest parallel code ZEUS-MP2 (Hayes et al. 2006) which has been extended with multifrequency radiative transfer and primordial gas chemistry (Whalen and Norman, 2007). Finally, in view of the increasing dataset sizes we are requesting a large allocation on SDSC DataStar *strictly for the purpose of data analysis*. No other NSF center has the combination of archival, SAN disk, and shared memory resources at present to store, stage, and support data analysis of 2048^3 datasets and larger. A summary of the projects and the resource request is given in Table 1, together with the name of the individual who will be doing the simulations. Enzo and ZEUS-MP2 performance and scaling on NSF platforms is given in Sec. III. Fifteen publications resulting from last year's allocation are given in Sec. IV.

Project	SUs	Machine	Student/Collaborator
1. Early Cosmic Evolution: The Rise and Fall of Pop III	10,000,000	TACC Ranger	Paschos (PD) Reynolds (PD)
2. Chemical Enrichment and Triggered Star Formation by the First Stars	400,000	SDSC DataStar p655 →NICS Kraken	Whalen (PD)
3. Lyman α Forest Signatures of He II Reionization	5,400,000	TACC Ranger	Paschos (PD) Reynolds (PD)
4. Fisher Matrix for the Lyman α Forest	430,000	TACC Ranger or NCSA Abe	Paschos (PD)
5. Lightcone Data Analysis: All Sky Maps of CMB Anisotropies	400,000	SDSC DataStar p655 & p690	Wagner (GR) Skory (GR)
6. ENZO Petascale Deveopment	70,000 200,000 160,000	NCSA Cobalt TACC Ranger NICS Kraken	Harkness (staff)

Table 1. Summary of projects and SU request. PD=postdoc, GR=grad, F=faculty.

Project 1: Early Cosmic Evolution: The Rise and Fall of Pop III

Dark matter mini-halos in the mass range 10^5-10^6 solar mass virializing at z~20 collect enough primordial gas within them to host the formation of the first generation of stars (Population III). Mediated by molecular hydrogen cooling, gas in the centers of these halos condenses and becomes unstable to gravitational fragmentation with a typical mass scale of a few hundred solar masses (Abel, Bryan & Norman 2002; Bromm, Coppi & Larson 2002; Norman 2008). Because of their large mass and primordial composition, Pop III stars are intense UV emitters, and they begin the process of cosmic reionization (Barkana & Loeb 2001). The contribution of Pop III stars to the eventual complete reionization of the intergalactic medium (IGM) is poorly known both observationally, and theoretically. One of the key theoretical uncertainties is the duration of the Pop III star-forming era. There is good agreement among simulators that Pop III stars begin forming in abundance in the redshift interval $30 > z > 20$ ("the rise of Pop III"; Yoshida et al. 2003). Two effects reduce their abundance at lower redshifts: chemical feedback from supernovae, which pollutes the pristine gas with heavy elements, and radiative feedback, which creates a UV radiation background which photo-dissociates the H_2 molecules which are needed to cool and collapse the primordial clouds. Previously, it was believed that Pop III star formation became self-limiting at z~25 due to radiative feedbacks. However, recently it was found that Pop III stars can form in rather high UV backgrounds (O'Shea & Norman 2007, Wise & Abel 2007a). Furthermore, it was found that Pop III stars can form in the relic HII regions produced by earlier Pop III stars (O'Shea et al. 2005, Mesinger et al. 2006, Yoshida et al. 2007). Given these findings, it is not clear when the Pop III era wanes in the absence of chemical feedback.

We propose to examine the rise and fall of Pop III by simulating a statistical volume (4 Mpc) over the redshift interval 100 > z > 10 taking the feedback of ionizing and photo-dissociating radiation into account. We will use Enzo in its "AMR everywhere" mode with a root grid of 1024^3 cells and 10 levels of refinement—for an effective dynamic range of 10^6. We will simulate the evolution of a large population of Pop III star forming halos with sufficient resolution to determine that gravitational collapse is occurring in their cores. Once collapse conditions have been met, we will introduce Pop IIII star particles using the procedure of Wise & Abel (2007b). Photodissociating radiation in the Lyman-Werner (LW) band (11.2-13.6 eV) will be treated as an evolving background radiation field as described in Ricotti, Gnedin & Shull (2002). Ionizing radiation will be transported using our new implicitly-coupled radiative transfer and ionization kinetics algorithm. As this is our first attempt, we will use a single group approximation for the radiation spectrum above the hydrogen Lyman edge.

Resource Justification In year 1, we will do the cosmological hydro calculation using AMR, but transport the ionizing radiation on the root grid only, as we do not yet have our algorithm coupled to AMR. This is justifiable because it is found that in most cases, the halos that host Pop III stars are completely overrun by their HII regions in a few 100 kyr (Whalen, Abel & Norman 2004). By the beginning of year 2, we expect to have AMR radiative transfer working, and we will repeat the calculation more self-consistently. We emphasize the year 1 calculation will be fully self-consistent as regards the negative feedback due to the evolving LW background. The resources required for the Year 1 simulation can be broken down into the cost of the 1024^3 AMR cosmology simulation and the cost of the 1024^3 unigrid radiation transport simulation. We estimate the former by scaling up the cost of a 512^3 7 level simulation shown in Fig. 1carried out on DataStar and Cobalt last year. It required 350,000 SUs on 256 cpus. We multiply this by 12 (8 for the ratio of the root grid sizes, and 1.5 for the additional levels), we estimate this portion of the calculation to cost 4.2 Million SUs. Performance tests of our implicit radiation algorithm on DataStar yield a speed of 1.7×10^4 cell/cpu-sec on a 256^3 test. The algorithm is implemented using the optimally scalable multigrid solvers in the HYPRE package developed at LLNL. Assuming perfect scaling one timestep on the 1024^3 grid will cost 6.3×10^4 cpu-sec. Assuming 10^4 root grid time steps, we calculate the cost of the radiation calculation at 1.75×10^5 cpu-hrs. Making some allowance for non-ideal scaling, false starts, and analysis, we estimate the total cost of each run at 5 Million SUs. As Enzo currently runs at roughly the same speed and scalability on Ranger as on SDSC DataStar (Fig. 1), we do not adjust these numbers. We are requesting the same amount for the fully coupled rad-hydro AMR run as we expect substantial speedups in many areas as we implement the changes described in Project 6, and we learn more about the Ranger architecture.
Resource Request: 10 Million SUs, TACC Ranger, 2048-4096 cores (128-256 nodes), 4-8 TB RAM

Project 2: Chemical Enrichment and Triggered Star Formation by the First Stars

Ambitious numerical attempts to model the assembly of primitive galaxies do so one primordial star at a time, with each new star forming in the relic H II regions of its predecessors (Yoshida, et al 2007). A crucial effect completely excluded in current studies is the contamination of these primeval structures with metals from supernovae (SN), which can radically alter their evolution. When mixed with heavy elements, pristine gas cools far more efficiently, fragmenting into stars much less massive than their progenitors. Expanding SN remnants are also prone to violent instabilities that form clumps that may collapse directly into a second, prompt generation of stars. Accurate simulations capturing these effects are key to predicting chemical abundance patterns in early galaxies and the ionization of the universe at high redshifts, both of which will soon be probed by the *James Webb Space Telescope*. The models require both the H II region of the progenitor, which also exhibits strong shocks and clumping due to dynamical instabilities, and the SN blast, which must be evolved on spatial scales from 0.001 pc to 1000 pc for self-consistency.

We propose twelve unigrid 1024^3 simulations of primordial SN in cosmological halos with the ZEUS-MP2 astrophysical radiation hydrodynamics code (Hayes et al. 2006; Whalen & Norman 2006, 2007a). The numerical suite will address Type II SN, hypernovae, and pair-instability SN in four dark matter halos from 6.9×10^5 to 1.2×10^7 solar masses, those in which Population III stars are expected to form via H_2 cooling. Our recent 1D work indicates that the SN remnant evolves according to whether or not the halo is photo-ionized by the original star (Whalen & Norman 2008). If so, the remnant exits the halo and chemically enriches the IGM, possibly with prompt 2^{nd} star formation. If not, the SN ejecta instead decelerates violently in the halo, enriching its interior with metals but failing to escape. The remnant eventually falls back to the center of the halo, either fueling a central black hole or fracturing into a swarm of low mass stars confined to the dark matter potential of the halo. Twelve models are the minimum set required to determine in what circumstances primordial supernovae chemically enriched the IGM, formed globular clusters, or created the high redshift black holes that may be the origin of the supermassive black holes at the centers of most massive galaxies today.

Resource Justification

Each simulation proceeds in two steps: first we ionize the halo with multifrequency radiative transfer turned on; second, after the star explodes, we propagate the supernova remnant without radiative transfer. Standard multifrequency calculations of ionization front (I-front) transport in $1000 \times 250 \times 250$ boxes with 120 energy groups cost 1000 SU on 8 P655 nodes at SDSC DataStar. Since ZEUS-MP UV transport scales nearly linearly in the 2- and 3-coordinates, we estimate that each 1024^3 simulation will cost 16,000 SU, and the 12 H II region calculations will require 192,000 SU. Preliminary timing profiles of 1D SN explosions in primordial halos on a 1000-zone expanding grid indicate that one 3D blast model, which takes an H II region model as its initial conditions, will consume 16,000 SU. We therefore request an allocation of 400,000 SU on DataStar (until 10/1/08) and Kraken (starting 6/1/08) for this simulation campaign.

Project 3: Lyman alpha Forest Signatures of Helium II Reionization

The standard simulation of the Lyman alpha forest (LAF) assumes that the IGM is fully ionized at z~6 and calculates its subsequent evolution assuming an evolving UV radiation background treated in the optically thin limit (e.g., Jena et al. 2005). There is now ample observational evidence that the hydrogen and and by inference helium in the IGM was singly ionized at redshifts z~6 by stellar sources, while helium only became doubly ionized at substantially lower redshifts by quasars (Sokasian, Abel & Hernquist 2003; Paschos, Norman & Bordner 2007). In Paschos, Norman & Bordner (2007) we simulated the inhomogeneous reionization of He II to He III in the IGM by a time-evolving quasar population, which were treated as discrete point sources. It was found that the mean temperature of the IGM is raised by 1.7x relative to the standard calculation**.** This is a huge effect, and must be taken into account when comparing simulations with precision observations of the LAF. This late reheating result is due to opacity effects in the gas, and was obtained by post-processing density fields from an optically thin Enzo simulation of the Lyman alpha forest with a standalone radiative transfer code. By splitting the calculation in this way, we treated the ionization balance, thermal evolution, and hydrodynamics only approximately. We propose to repeat this calculation fully self-consistently in order to get a more accurate description of the thermal evolution of the IGM, and to do this at a high enough spatial resolution so that we can calculate the expected signature on the Lyman alpha forest. Motivation for this calculation comes from a recent paper from Faucher-Giguere et al. (2007), who have measured the redshift evolution of the mean opacity of the Lyman alpha forest over the redshift range $2 < z < 4.2$. They find a departure from a smooth powerlaw evolution at z~3.2 which they suggest might be related to the end of Helium II reionization.

To explore this we will simulate a cosmological volume large enough to contain a representative sample of quasars (~100 Mpc box) with a spatial resolution high enough to resolve the LAF absorbers (~50 kpc). This will be done with a uniform grid of 2048^3 cells, and Enzo's implicit radiation transfer-ionization kinetics solver. We will generate synthetic absorption spectra for the H I and He II forests and analyze their statistical properties as he have done in Paschos et al. (2007). We will compare the results of this experiment with a standard optically thin simulation of the same size.

Resource Justification

Just as above, we can estimate the resource requirements by calculating the cost of a standard LAF simulation of that size, and add to that the cost of the radiation transfer calculation. For the first, we have the cost of the 2048^3 LAF simulation done on NERSC Seaborg last year, which required 1.2 Million cpu-hrs on 4096 Power3 processors. Multiplying this by 3 to account for the factor 3 smaller cell size, and

dividing by 2 to account for the differential in speed of the machines at scale (see Fig. 1), we arrive at 2 Million SU for the hydrodynamic cosmology portion of the calculation. As for the radiation transfer calculation, since it is implicit there is no Courant condition on cell size. In fact, we find that the number of timesteps to completion on a simple Stromgren sphere test to be independent of resolution. Therefore we simply multiply the estimate made above of 175,000 SUs for 104 timesteps on a 1024^3 grid by a factor of 8 to account for the larger number of grid points. We thus estimate 8x175,000=1.4 Million SUs for the radiation portion. Our total estimate for the self-consistent simulation and the optically thin comparison run is therefore 2 x 2 + 1.4 = 5.4 Million SUs.
Resource Request: 5.4 Million SUs, TACC Ranger, 4096 cores (256 nodes), 8 TB RAM

Project 4: Computing the best unbiased estimators (BUE) from the Fisher matrix of cosmological parameters using the Lyman alpha forest

We propose the use the Lyman alpha absorption spectra in the intermediate to high redshift range (z = 2-5) in order to obtain constraints on a set of key cosmological parameters. We intend to run a series of cosmological numerical simulations covering the following parameter space: [Ω_Λ, Ω_b, σ_8, h, L, ΔL, Γ, Ξ]. The first four parameters have their usual cosmological meaning, L and ΔL are the box size and mesh resolution of the simulation, Γ is the amplitude of the hydrogen ionizing UV background, and Ξ is a parametrization of the slope of the UVB spectrum at the He II ionization edge. We will construct synthetic Lyα absorption spectra through the simulated volume to compute a list of the physical properties of the forest that have been also measured in observed quasar spectra. These include the mean, variance and the pdf of the transmitted flux, the flux power spectrum, the autocorrelation and cross correlation functions and the line broadening width and column density distributions. Therefore our objective is to construct the covariance matrix of our parameter space otherwise known as the "Fisher Matrix". Sophisticated algorithms exist to optimally sample the 8-dimensional parameter space required to compute the unbiased estimators of the individual parameters, including Latin Hypercube and Centroid Vornoi Tesselation (CVT). As a rule of thumb, for N independent parameters, ~N^2 points in parameter space are required. For our problem we have estimated that we need ~64 512^3 simulations. Given the relatively modest size of these simulations and the capacity of TACC Ranger, this should not be a problem.

Based on the measured time to solution on SDSC DataStar, each 512^3Lya simulation requires about 6700 SUs on 256 processors. At present, Ranger and DataStar run at the same speed (Fig. 1), therefore we do not scale this estimate. For 64 simulations we therefore request 430,000 SUs on TACC Ranger. NCSA Abe would be an acceptable alternate platform for this parameter survey.
Resource Request: 430,000 SUs, TACC Ranger, 256 cores (16 nodes), 512 GB RAM

Project 5: All-sky maps of gravitationally-induced secondary CMB anisotropies (PhD thesis research) The anisotropies, or variations, of the temperature of the cosmic microwave background (CMB) provide a direct measure of several cosmological parameters, including the curvature and matter densities in the universe [1]. The primary source of these anisotropies are the fluctuations in the density field at the surface of last scattering. Secondary sources of anisotropies include red-shifting of CMB photons due to evolving gravitational potentials, the integrated Sachs-Wolfe (ISW) effect [2], Compton scattering of CMB photons, the Sunyaev-Zel'dovich (SZ) effect [3], and small-angle deflection of CMB photons by potential wells, or weak gravitational lensing [4]. Quantitative understanding of each of these signals will be needed for the next series of CMB experiments [5]. This project will focus on the gravitationally induced signals, the ISW effect and weak gravitational lensing. We will create high-resolution (~0.5 arcminute per pixel) all-sky maps of lensing potential, deflection angle, and temperature variations along lines of sight from the present epoch to a redshift of 10. These maps will be used to measure the impact of these secondary effects on the angular power spectrum of the CMB in the range of $500 < l < 4000$ and the gravitationally induced B-mode polarization.

The process for creating the maps is based on one used on the Millennium Simulation data [7]. In this method, the simulation volume is stacked in randomly oriented and translated spherical shells around the origin of the rays at a redshift of zero. The rays are then cast radially through the simulation, going backwards in time. A ray passes through several datasets (time-based outputs) as it travels through an individual shell. At the end, the desired quantities (deflection, potential, etc.) are integrated along the rays assuming the Born approximation. We improve upon the current work by having higher resolution data that includes baryonic physics and a larger realization of the universe. The maps will be generated from existing data--a "light cone" consisting of 8 statistically independent simulations of varying physical box sizes designed to provide a large comoving volume (over 1 Gpc/h per side) at a fixed angular resolution of < 30 arcseconds. These simulations were produced with the AMR cosmology code Enzo running on LLNL Thunder, and include both dark matter particles and baryons. Over 200 TB of data has been transferred to SDSC for analysis. Measurements of the Sunyaev-Zel'dovich effect from a portion of the light cone have already been published [6]. We will use SDSC DataStar for the analysis, as this is where the data is located.

Resource Justification

At least five maps will be made: two preliminary maps at a lower resolution of approximately 2 arminutes to verify the process and measure the scaling of the analysis software; two high-resolution maps at approximately half-arcminute per pixel resolution; and a high-resolution map comparing the results of the Born approximation to a ray-tracing method that deflects the ray along the path. Creating the maps will use two steps, post-processing and ray-casting. The post-processing step is a single pass through all of the datasets (~50) to be used in the light cone to compute and store the gravitational potential, its gradient and time derivative. Then, for each map, a set of orientations and translations is chosen for each shell, and the rays cast through the datasets. The post-processing step requires the same number of MPI tasks as the original simulation (256 to 1372), and based on the simulation performance, each dataset will take 2 wallclock hours to process. Further, due to the memory footprint of the AMR hierarchy, this step needs a minimum of 2 GB of memory per task. This could lead idling processors on some multi-processor nodes, with a subsequent SU consumption penalty. The time needed to complete this step is 100,000 SUs (50 datasets X 2 wallclock hours X 1000 tasks). The ray-casting step will be done using 1024 processors, and will have the same minimum memory requirements. We estimate that each dataset will take 1 wallclock hour, using the performance from creating high-resolution (8196^2 pixel) SZ maps as guide. This adds another 250,000 SUs to the request (50 datasets X 1 wallclock hour X 1024 tasks). Finally, analysis of the maps themselves will be compute intensive, and we would like to request an additional 50,000 SUs for data analysis, bring the total to 400,000 SUs (100,000 + 250,000 + 50,000).

First choice - SDSC Power4+ p655 DataStar
Second choice - NCSA Linux Cluster Abe

Project 6: ENZO Petascale Development

During the last allocation period our scaling experiments were limited by the number of processors actually available on NSF platforms. With the availability of TACC Ranger and NICS Kraken we propose to continue exploring the limits of scalability with representative full-scale problems of interest.

At present, ENZO uses MPI domain-decomposition only. While weak scaling has been successful up to 8000 cores, this approach is limited by contention amongst cores sharing interconnect access and bandwidth. By using N OpenMP threads per N-core processor within a single MPI task per (N-core) processor we can minimize contention and make some computational use of the many cores on each processor. In this allocation period we expect to be able to make effective use of quad-core processors with 4 OpenMP threads per MPI task.

(a) Unigrid development

In the present cycle we used NCSA Cobalt to generate initial conditions for the scaled Lyman alpha forest benchmark sequence shown in Fig. 1, including top grid mesh dimensions of 2304^3, 2560^3 and 2816^3. The task decomposition is shown int Table 2. In this cycle we would like to generate 3072^3 initial conditions for the benchmark sequence and also for a potential unigrid model to be conducted on TACC Ranger or NICS Kraken if resources allow.

Table 2. ENZO Unigrid Peta-scaling tests

Top grid	Decomp.	Tasks	Systems
2048^3	16^3	4096	Ranger, Kraken
2304^3	18^3	5832	Ranger, Kraken
2560^3	20^3	8000	Ranger, Kraken
2816^3	22^3	10648	Ranger, Kraken
3072^3	24^3	13824	Ranger, Kraken

Accurate determinations of unigrid performance can be made with just a few time steps, plus the time required to read initial data and initialize the simulation and to write a few restart dumps. On the NERSC Cray XT4 even a 2560^3 test can be accomplished in less than 1 wall clock hour on 8000 nodes with 8000 cores active, i.e. < 16,000 SU.

We also plan to use the test sequence to develop and validate the hybrid MPI/OpenMP version of ENZO. In the unigrid case the cost of running ENZO is dominated by the hydrodynamics and the physics package and both of these are easily threaded with approximately equal load balancing amongst the threads. Thus with the unigrid case we can reliably test the hybrid code and also have a reasonable expectation of a significant gain in efficiency by utilizing cores that would otherwise be idle due to limitations on memory per node and/or the intrinsic limits of ENZO scalability in the present version.

(b) AMR development
Once the hybrid code is operational we can also expect significant gains for full-AMR cases although the overall efficiency may be lower than obtained in unigrid due to the currently non-threaded hierarchy rebuild which is more costly for deep AMR. The main challenge with AMR applications of ENZO is to increase the root grid dimension from the present record-breaking 512^3 7-level AMR case up to 1024^3 7-level AMR which should be possible with TACC Ranger. We are working on methods to distribute the AMR hierarchy and to reduce the storage requirements. With TACC Ranger it should be possible to accomplish the 1024^3 root grid with 7-levels of refinement using ~4096 MPI tasks with 4 OpenMP threads per MPI task. Several limited test runs will be required to establish that the entire run is feasible.

Table 3. ENZO AMR Peta-scaling Tests

Top grid	Decomp.	Tasks	Comment
512^3	8x8x4	256	L7 run on Cobalt & DataStar completed
1024^3	8^3	512	L7 run already in progress on TACC Ranger
1024^3	16^3	4096	Target for TACC Ranger

It is very difficult to estimate AMR performance in advance of an actual run. Depending on the details of the physical model, the onset of rapid subgrid generation will occur at varying redshift and obviously one cannot generate a representative case at, say, z=3 without actually performing the computation to reach this point. We will therefore use the restart dumps from the existing 512^3 7-level restart dumps for tests at various redshift points and utilize the restart dumps from the 1024^3 7-level WHIM model as it progresses.

<u>Resource Request</u>
We request the following resources for this project.
<u>NCSA Cobalt</u>: At present, the ENZO initial conditions generator is a shared-memory OpenMP-parallel code which performs effectively on 32 processors (or more) of the NCSA Altix. Only the Altix can accommodate the necessary shared memory so this platform is critical until we develop a fully distributed memory version. Since the capability is only rarely used we request access to Cobalt to maintain progress in the near term. 7 individual runs of the initial conditions code are required to generate density, gas velocities and particle positions and velocities.

Table 4. Scaling tests resource requirements

Top grid	Memory	SU (hrs x procs x runs)
2304^3	683 GB	4 x 32 x 7
2560^3	935 GB	6 x 32 x 7
2816^3	1242 GB	7 x 32 x 7
3072^3 (est)	1619 GB	10 x 32 x 7

If the largest scale run requires dedicated access the total cost would be 10 x 512 x 7 = 35,000 SU. For two cases we request 70,000 SU on Cobalt.
<u>TACC Ranger</u>: We wish to perform tests of hybrid ENZO scaling & performance analysis on TACC Ranger using 1, 2 or 4 OpenMP threads per MPI task with unigrid cases up to 3072^3. Cost is mainly determined by the largest test case and increases with the number of cores we can utilize effectively. Since we hope to reach > 20,000 cores and the small benchmark tests require ~1hr we estimate that 10 development runs at this scale would cost 200,000 SU on Ranger.

NICS Kraken: Similarly, we wish to perform hybrid ENZO scaling & performance analysis on NICS Kraken using 1, 2 or 4 OpenMP threads per MPI task with unigrid cases up to 2560^3. The largest test case would employ 16,000 cores. 10 development runs at this scale would cost 160,000 SU on Kraken.

III. Methodology: ENZO and ZEUS-MP2 Scaling Results

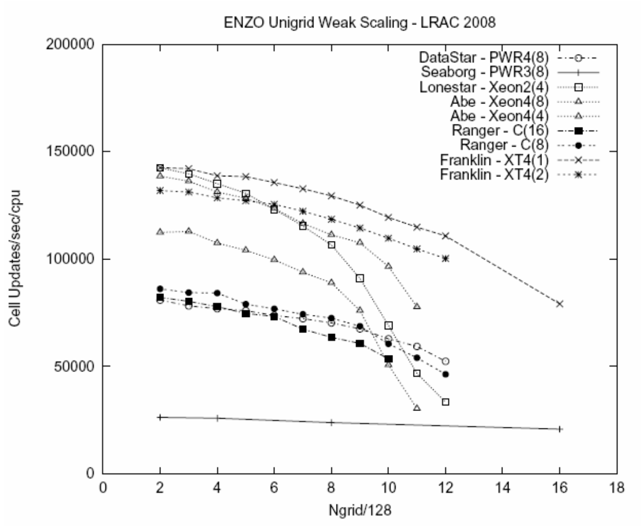

Fig. 1. Cell update rates per cpu for *Enzo* unigrid multispecies hydrodynamic cosmology simulations on a variety of NSF and DOE platforms for different grid sizes (weak scaling).

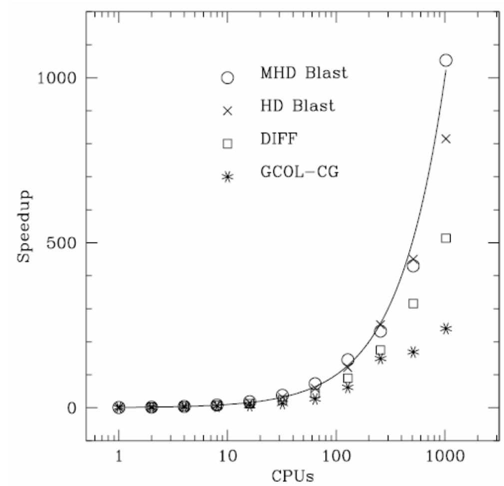

Fig. 2. Parallel speedup for ZEUS-MP2's hydro, MHD, radiation diffusion, and self-gravity algorithms for 2563 strong scaling tests. Solid line is ideal (from Hayes et al. 2006).

The standard ENZO benchmark sequence consists of series of Lyman Alpha Forest models where the physical box size is increased in proportion to the number of top grid mesh points in one dimension. The number of MPI tasks is set so that the work per MPI task consists of a 128^3 root grid "tile" and such that the Courant condition is constant across the entire sequence. Ideal scaling performance would be indicated by constant grid cell updates/sec as a function of the top grid dimension.

Table 5: ENZO weak scaling test description

X-dimension	Box size (Mpc/h)	Decomposition	MPI tasks
512	27.264	4x4x4	64
1024	54.528	8x8x8	512
1536	81.792	12x12x12	1728
2048	109.056 Mpc/h	16x16x16	4096
2560	136.320	20x20x20	8000
3072	163.584	24x24x24	13824

The results are shown in Figure 1. The Cray XT4 exhibits the highest peak performance and the best scaling behavior overall. The NERSC Cray XT4 "Franklin" was the only system available capable of running the 2304^3 and 2560^3 benchmarks. The Cray XT4 performance and scalability is closely matched by the Cray XT3 systems BigBen (at PSC) and Jaguar (at ORNL) and we confidently expect similar results on UT/OR Kraken when it enters service.

TACC Ranger "early user" performance and scaling closely match results from SDSC DataStar. It is important to note that the Ranger results are very preliminary and we expect the absolute performance and scalability to improve as the system is tuned and we gain more experience with ENZO on Ranger.

The Dell Infiniband clusters both show high performance at smaller scales but relatively poor scaling at larger scales. Consequently, we expect to use NCSA Abe for medium-scale applications and for ensembles of small-scale models.

We note that since these measurements were made, we have identified a significant optimization which dramatically improves scalability at > 1024 MPI tasks. The scaling limitation was identified using CrayPAT on the NERSC Cray XT4 and the optimization verified on the same system. The algorithmic change applies equally to all systems and should allow for reasonably efficient execution with > 10,000 cores, or equivalently an improvement in *strong scaling* at a fixed problem size.

IV. Publications Supported by Previous LRAC Award (cosmology only)

1. **Norman, M. L.**, et al. 2008. "Simulating Cosmological Evolution with Enzo", in Petascale Computing: Algorithms and Applications, Ed. D. Bader, Chapman & Hall/CRC Computational Science Series
2. **Norman, M. L.** 2008. "Population III Star Formation and IMF", to appear in "First Stars III", eds. B. O'Shea, A. Heger & T. Abel, AIP Press
3. Xu, H., Collins, D., **Norman, M. L.**, Li, S., & Li, H. 2008. "A Cosmological AMR MHD Module for Enzo", ApJ submitted
4. Xu, H., Collins, D., **Norman, M. L.**, Li, S., & Li, H. 2008. "A Cosmological AMR MHD Module for Enzo", to appear in "First Stars III", eds. B. O'Shea, A. Heger & T. Abel, AIP Press
5. Whalen, D., Van Veelen, R., O'Shea, B. & **Norman, M. L.** 2008. "The Destruction of Cosmological Minihalos By Supernova Explosions", ApJ, submitted
6. Whalen, D. & **Norman, M. L.** 2007. "Three-Dimensional Dynamical Instabilities in Galactic Ionization Fronts", Ap. J., in press, arXiv e-print (arXiv:astro-ph/0703463)
7. Whalen, D. & **Norman, M. L.** 2007. "Ionization Front Instabilities in Primordial H II Regions", Ap. J., in press, arXiv e-print (arXiv:0708.2444)
8. Whalen, D.; O'Shea, B. W.; Smidt, J.; **Norman, M. L.,** 2007. "Photoevaporation of Satellite Halos by the First Stars", Ap. J., in press, arXiv e-print (arXiv:0708.1603)
9. Whalen, D.; O'Shea, B. W.; Smidt, J.; **Norman, M. L.,** 2007. "Photoionization of Clustered Halos by the First Stars", to appear in "First Stars III", eds. B. O'Shea, A. Heger & T. Abel, AIP Press, arXiv e-print (arXiv:0708.3466)
10. Paschos, P., **Norman, M. L.** and Bordner, J. O. 2007. "Late Reheating of the IGM by Quasars: Radiation Hydrodynamic Simulations of Inhomogeneous He II Reionization", Ap. J., submitted arXiv e-print (arXiv:0711.1904)
11. Tytler, D.; Paschos, P.; Kirkman, D.; **Norman, M. L.**; Jena, T., 2007. "How Simulated Spectra of the Lya forest Change with the Size of the Box of Numerical Simulations", Ap. J., submitted, arXiv e-print (arXiv:0711.2529)
12. O'Shea, B. W., & **Norman, M. L. 2007**. "Population III Star Formation in a Lambda CDM Universe, II: Effects of a Photodissociating Background", Ap. J., in press, arXiv e-print (arXiv:0706.4416)
13. Hallman, E. J.; O'Shea, B. W.; Burns, J. O.; **Norman, M. L.**; Harkness, R.; Wagner, R. 2007. "The Santa Fe Light Cone Simulation Project. I. Confusion and the Warm-Hot Intergalactic Medium in Upcoming Sunyaev-Zel'dovich Effect Surveys", ApJ, 671, 27
14. Burns, J. O.; Hallman, E. J.; Gantner, B.; Motl, P. M.; **Norman, M. L.,** 2007. "Why Do Only Some Galaxy Clusters Have Cool Cores?", Ap. J., submitted arXiv e-print (arXiv:0708.1954)
15. Hallman, E. J.; Burns, J. O.; Motl, P. M.; **Norman, M. L**., 2007. "The β-Model Problem: The Incompatibility of X-Ray and Sunyaev-Zeldovich Effect Model Fitting for Galaxy Clusters", ApJ, 665, 911

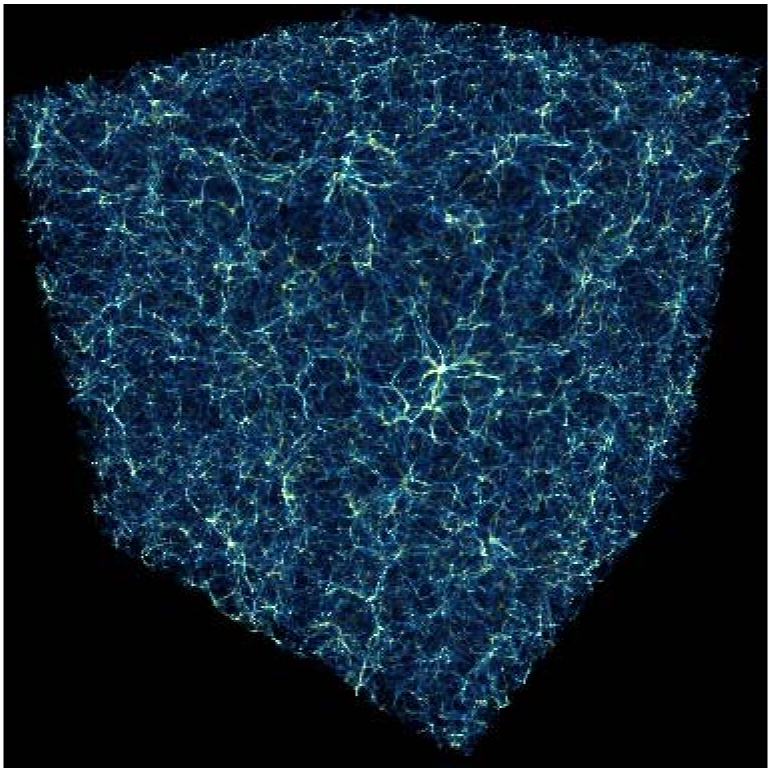

Fig. 1: A globally refined AMR simulation of cosmic structure simulated with Enzo on SDSC DataStar and NCSA Cobalt. The calculation generated >350,000 subgrids at 7 levels of refinement around >10,000 collapsed objects. From Norman et al. (2008).

References

[1] Spergel, et. al., ApJS, 170, 377 (2007), http://arxiv.org/abs/astro-ph/0603449

[2] R. K. Sachs and A. M. Wolfe, ApJ., 147, 73 (1967)

[3] Birkinshaw, Phys.Rept. 310 (1999) 97-195, http://arxiv.org/abs/astro-ph/9808050v1

[4] Lewis and Challinor, Phys.Rept. 429 (2006) 1-65, http://arxiv.org/abs/astro-ph/0601594

[5] Legacy Archive for Microwave Background Data Analysis, http://lambda.gsfc.nasa.gov/

[6] Hallman, et. al., ApJ, 671, 27, http://arxiv.org/abs/0704.2607

[7] Carbone, et. al., MNRAS Submitted, http://arxiv.org/abs/0711.2655v1

Abel, T., Bryan, G. & Norman, M. L., 2002, Science, 295, 93

Bromm, V., Coppi, P., & Larson, R., 2002, ApJ, 564,23

Barkana, R., & Loeb, A., 2001, ARA\&A, 39, 19

Faucher-Guiere, P., Prochaska, J., Lidz, A., Hernquist, L. & Zaldarriaga, M. 2007, arXiv e-print (arXiv:0709.2382)

Hayes et al., 2006, ApJ Supp., 165, 188

Jena et al. 2005, MNRAS, 361, 70-96

McDonald, P 2003, ApJ, 585, 34-51

Mesinger, A., Bryan, G., & Haiman, Z., 2006, ApJ, 648, 835

Norman, M. L., et al. 2008, in Petascale Computing: Algorithms and Applications, Ed. D. Bader, Chapman & Hall/CRC Computational Science Series
O'Shea B., Abel. T., Whalen, D., & Norman M. L., ApJL, 628, L5
O'Shea B., Norman M. L., 2007, astro-ph/0706.4416
Paschos, P., Norman, M. L. and Bordner, J. O. 2007, Ap. J., submitted, arXiv e-print (arXiv:0711.1904)
Reynolds et al. 2008. *in prep.*
Ricotti M., Gnedin N. Y., Shull J. M., 2002, ApJ, 575, 49
Seo, H-T; Eisenstein, D J. ApJ 2007 665, 14-24A

Sokasian, A., Abel, T. & Hernquist, L. 2003, MNRAS, 332, 601
Tegmark, Taylor & Heavens ApJ, 1997, 480, 22-35
Whalen D., Abel T., Norman M. L., 2004, ApJ, 610, 14WN07a
Whalen, D., & Norman, M. L., 2006, Ap. J. Supp., 162, 281
Whalen, D., & Norman, M. L., 2007a, Ap. J., in press, arXiv e-print (arXiv:0708.2444)
Wise, J. & Abel, T. 2007a, ApJ, 671, 1559
Wise, J. & Abel, T. 2007b, arXiv e-print (arXiv:0710.3160)
Yoshida N., Abel T., Hernquist L., Sugiyama N., 2003, ApJ, 592, 645
Yoshida N., Oh S. P., Kitayama T., Hernquist L., 2007, ApJ, 663, 687