Creating the Observable Universe on a Supercomputer

Richard Bower, Tom Theuns, Adrian Jenkins (Durham University), Jason Beech-Brandt (Cray CoE), Robert Crain (Swinburne University)

Scientific Aims

Overview.

A great challenge for Computational Cosmology is to produce a realistic model of the observable universe. Although on large-scales matter interacts primarily through gravity, the visible parts of the Universe are dominated by gas which has condensed into stars and which are created through a complex interplay of gravity, atomic and molecular cooling, and energetic outflows due to stars themselves. Fortunately, recent advances have allowed simulators to approach the dynamic range required to tackle this problem head-on. Our proposed simulations will track the formation galaxies from their cosmological initial conditions in sufficiently detail to resolve the condensation of individual star forming clouds. This proposal aims to facilitate this advance by unlocking the world-leading Gadget code so that it can efficiently run across 100s to 1000s of processors.

The main matter component in today's Universe appears to be a yet-undiscovered elementary particle whose contribution to the cosmic density is more than 5 times that of ordinary baryonic matter. This dark matter particle interacts extremely weakly with regular atoms and photons, so that gravity alone has affected its distribution since very early times. Recent observations have established a standard paradigm in which dark matter emerged from the early Universe with negligible thermal velocities and a Gaussian and scale-free distribution of density fluctuations. This ``Cold Dark Matter" (CDM) hypothesis determines the statistical properties of dark matter structures at early epochs when the universe was almost uniform.

The non-linear growth of structure in the dominant dark matter component can be followed accurately with N-body simulations, since the initial conditions and the evolution equations are known. Over the past two decades, numerical simulations of dark matter played a primary role in establishing the viability of the LCDM paradigm (where 'L' stands for an additional dark energy field), and the scientists on this proposal provided several of the most important key contributions in achieving this progress. Today, the increasingly sophisticated simulation algorithms and the expansion of computational capabilities have led to solid and increasingly detailed theoretical predictions that now make the CDM paradigm falsifiable by observation. So far, the paradigm has successfully passed tests based on a direct comparison of the distribution of matter on large scales (from ~1 Mpc to the size of the observable Universe) with a wide array of observations. With more than 10^{10} simulation particles it is the largest cosmological simulation ever carried out, following the formation of about 20 million galaxies in a representative volume of the Universe more than 2 billion light-years across.

However, this dark matter universe cannot be compared to that seen through telescopes. This requires that we add a component of baryonic ("normal") matter, principally neutrons, protons and electrons, and introduce the equations that allow this fluid to cool, condense and form stars. Attempts to connect this theoretical model to current observations have thrown up several fundamental issues. Firstly, it is hard to understand why galaxies are so rare in our universe. Secondly, there is a hierarchy problem. Gravity would be expected to create small galaxies first and then assemble these into larger and larger systems. This does not match what is seen when the universe is studied with the largest telescopes. Such observations demonstrate that the largest galaxies are formed at surprisingly early times, perhaps only 1 billion years after the big bang. Only the formation of small galaxies continues over the remaining 15 billion years of cosmic time. Recent analytic work has suggested that these problems can be solved by including "feedback" from supernovae (the death of massive stars) and, crucially, AGN in the simulation (eg. Bower et al 2006). The challenge is now to demonstrate these simplified models in a fully numerical simulation.

A number of other, related problems exist when the properties of galaxies are examined in detail. For example, it is debated whether the observed shallow rotation curves of low-luminosity dwarf galaxies are consistent with the cuspy dark matter halos predicted by today's N-body simulations (Navarro, Frenk & White 1996, 1997; Navrarro et al. 2004). Whether or not the shape of the central dark matter profile can be significantly modified by dissipative baryonic processes, or by dynamical effects in stellar disks such as bar formation, is equally contested. Crucially, the way forward is to simulate the formation of galaxies in sufficient detail that their properties can be directly compared with images from the world's largest telescopes.

We therefore propose to carry out high-resolution simulations of the formation of galaxies (such as the Milky Way) that are fully embedded in their proper cosmological setting. We will include the most sophisticated hydrodynamical techniques that presently exist for the direct modelling of star formation and black hole accretion on galactic scales, including their back-reaction in the form of feedback processes. We are hence aiming to form Galaxies on a supercomputer, including the stellar distributions in the disk and bulge, and in the satellite systems that orbit the outer halo, and also accounting for the properties of condensed gas in the galactic disk as well as in the hot and diffuse component residing in the halo, and then to compare these simulated galaxies directly with their observed counter parts.

Besides being able to address the important small-scale problems of CDM, the proposed simulations will lead to new insights into two of the most fundamental puzzles in contemporary galaxy formation physics. What is the role of active galactic nuclei for galaxy formation? Why are galaxies so rare? And why have previous generations of low-resolution hydrodynamic simulations of galaxy formation in general failed to produce realistic galaxies?

Finally, our simulations will also address the long-standing problem of disk galaxy formation within the CDM cosmology. When gases cool and condense in a dark matter halo, they should retain their specific angular momentum and eventually settle into a cold, rotationally supported disk, which then transforms into stars once the gas becomes locally dense enough. However, hydrodynamical simulations that attempted to follow this process within the CDM model have consistently been rather unsuccessful, over producing stars and creating, at best, anaemic disks which are much too small compared with observations. This is often interpreted as being caused by inaccurate treatment of galaxy winds, driven by a combination of supernovae and AGN. Even though some gradual progress towards more realistic disk sizes has been achieved recently (Okamoto et al. 2005, Libeskind et al.

2006, Governato et al 2009), the inability to form disk galaxies with realistic morphologies from LCDM initial cosmology has been a major embarrassment for the theory of galaxy formation for more than a decade. Our proposed simulations have the potential to make significant progress on the galaxy formation question. Their drastically improved resolution compared to previous work allows a proper representation of star forming clouds, and combined with the sophisticated modelling of star formation and feedback processes in our code, we expect that our simulation models will be able to account for galaxy formation more successfully than ever before.

The code we will use is based on GADGET, a cosmological N-body/SPH simulation code designed for massively parallel computers with distributed memory (Springel, White & Yoshida 2001, Springel 2005). This code has been developed by Volker Springel, in close collaboration with the other team members of this proposal. A basic version of GADGET has been made publicly available and is presently the most widely used code in studies of cosmological structure formation. Roughly 400 publications using it are already in the refereed scientific literature.

GADGET computes gravitational forces with a hierarchical tree algorithm (optionally in combination with a particle-mesh scheme for long-range gravitational forces) and represents fluids by means of smoothed particle hydrodynamics (SPH). The code can be used for studies of isolated systems, or for simulations that include the cosmological expansion of space, both with or without periodic boundary conditions. In all these types of simulations, GADGET follows the evolution of a self-gravitating collisionless N-body system, and allows gas dynamics to be optionally included. Both the force computation and the time stepping of GADGET are fully adaptive, with a dynamic range which is, in principle, unlimited.

We are requesting 15 Million AU, which would already make it possible to carry out a small number of hydrodynamical simulations of Milky Way-sized halos with state-of-the-art resolution of a 1-2 million gas particles in a Milky-way halo with the existing code. However our aim is to improve the code performance to enable us to make a step change in particle number up to 5-10 million gas particles for a Milky-way halo which would be world leading. Because additional particles are needed to represent the cosmological environment and tidal field, the total particle number for the highest resolution simulations we want to do would be around 80 million particles. The resulting mass resolution is then 10^5Msun/h in the dark matter, and $2x10^4$Msun/h in the gas component, while the spatial resolution would reach down to 0.05 kpc/h, which sufficient to start resolving the individual star forming clouds. For runs of this size the wall-clock run times would be many months to years on small numbers of cores, which makes improving the scaling of the code to large core counts imperative.

Although it will probably not be feasible with the resources we request here to do more than one large Milky-way simulation, in the longer term we intend to compare simulations with and without black holes, and simulations with or without an explicit feedback model that accounts for galactic winds, all at extremely high resolution. The next upgrade of the COSMA supercomputer in Durham towards the end of this year would make this larger-scale project feasible by providing sufficient AUs, provided significant gains can be made in the code's scalability through this proposal.

Technical Challenges and Proposed Work

In the current version of GADGET3 there are some severe load balance issues with the current datasets which prevent scaling of the code to large core counts. The simulations proposed, generating sparse but strongly clustered systems with a relatively modest number of particles, are not what GADGET3 is optimised for. For example, initial Craypat analysis shows that runs on 64 cores of Phase 2a HECToR spend 50% of the time waiting to synchronize at MPI collectives. Our initial analysis also shows that the wallclock time per particle for individual timesteps decreases with particle count up to some threshold after which there is a large jump in this value. A decrease in time per particle makes sense as calculating more particles per timestep should be more efficient due to increase the amount of local compute work relative to communication required. However, the jump in time per particle at larger particle counts was not expected, and understanding this key to unlocking the scalability of GADGET3 for these particular simulations. It should be noted that the core communication in GADGET3 is point-to-point MPI, and is taking a small fraction of the total time in current simulations, so the potential for good scaling with these datasets does exist.

As such a program to better understand, and therefore remedy, the load balance issues in these particular simulations is proposed. Possible solutions to the load balance issue include tuning or modifying of the existing decomposition method, which seems to have difficulty dealing with the clustered systems arising from the current datasets. The current decomposition method does not adequately take into account the cost of particles in the individual partitions, so this is an obvious area for investigation. Another possible solution would also be including OpenMP in the code which would have the effect of making the local computational domains larger, and therefore more homogenous, thus reducing the potential for load imbalance. As such we propose the following plan of work.

Task	***Groups Involved***	***Time***	***Deliverable***
Continued profiling of the code at a range of core counts	***Cray CoE***	***2 weeks***	*A complete profile of GADGET performance and scaling*
Work to address load-balance issues in GADGET3	***Cray CoE jointly with Durham***	***6 weeks***	*Optimised version of GADGET3*
Perform GADGET3 run(s) with optimised code	***Durham jointly with Cray CoE***	***6 weeks***	*Cosmological datasets*
Post-processing of data and further analysis	***Durham jointly with Cray CoE***	***4 weeks***	*Final results*

Resources

The timetable for this project is outlined above. Initial investigation into optimization opportunities in GADGET3 have already begun and will be completed in the next few weeks. Work to address the load-balance issues will begin after this, and will take several weeks to implement. Based on our initial investigations, 15 million AU's will enable us to perform the above setup and a large production run. The output from this run will initially create 1-2 TB of cosmological data. Post-processing of this data we will allow us to reduce this substantially. In order to complete the optimization work and assist *when necessary with the running the of the production jobs, two person months of CoE time is requested.*

Deliverables

The specific outcomes and long-term benefits of this proposal will be:

- *A highly optimised version of GADGET3 running on Cray XT machines and configured for ongoing investigations of very high resolution cosmological simulations of individual galaxies.*
- *A new load-balancing scheme for GADGET3 that takes advantage of multi-core technology in its domain decomposition and work-load model.*
- *The highest-resolution simulation of the formation of Milky Way-like galaxy ever, enabling comparison to future datasets.*
- *Leading on to detailed set of models to investigate whether the standard cosmological model is able to produce Milky Way-like galaxies and an investigation into the effects of the energy emitted by super massive black holes on the formation of Milky Way-like galaxies.*