The Aquarius Project: Cold Dark Matter under a Numerical Microscope

A major puzzle in Cosmology is that 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 Cold Dark Matter (CDM) interacts extremely weakly with regular atoms and photons, so that gravity alone has affected its distribution since very early times, when the Universe was in a nearly uniform state.

When the effects of the baryons can be neglected, the nonlinear growth of structure is a well-posed problem where both the initial conditions and the evolution equations are known. In fact, this is an N-body problem par excellence. The faithfulness of late-time predictions is limited purely by numerical technique and by the available computing resources. Over the past two decades, simulations have already been of tremendous importance for establishing the viability of the CDM paradigm. In particular, simulation predictions for the distribution of matter on large scales have been compared directly with a wide array of observations; so far the paradigm has passed with flying colors.

Given CDM's success in reproducing the main aspects of the large-scale structure of the Universe, it is important to test its predictions on smaller scales, both to test it further and to seek clues to the nature of dark matter. In the Aquarius Project carried out by the international Virgo Consortium on the HLRB II supercomputer at LRZ, we aim to do this by studying the highly nonlinear structure of CDM halos in unprecedented detail. We are especially interested in the innermost regions of these halos and in their substructures, where the density contrast exceeds 10^6 and the astrophysical consequences of the nature of dark matter may be most clearly apparent. Quantifying these consequences reliably through simulations is, however, an acute challenge to numerical technique.

The numerical Challenge

In the Aquarius Project, we have performed the first ever one-billion particle simulation of a Milky Way-sized dark matter halo, improving resolution by a factor of more than 15 relative to previously published simulations of this type. The achieved mass resolution of ~1,700 solar masses is nearly a million times better than that of the largest cosmological simulation of large-scale structure formation carried out to date (the "Millennium Simulation"). Our spatial resolution reaches 20 parsec, which implies a dynamic range of close to 10^7 per dimension within the simulated box-size of more than 400 million lightyears across. This huge dynamic range makes our simulation a microscope for the phase-space structure of dark matter and enables dramatic advances in our understanding of the structure and substructure of dark matter in our galaxy. However, formidable challenges had to be overcome to make this calculation possible. Gravitational timescales are inversely proportional to the square root of the density, so simulating a CDM halo means dealing with a system where different regions evolve on timescales which may differ by factors of thousands. Codes with spatially-dependent, adaptive timestepping are mandatory otherwise the most rapidly evolving regions, which usually include only a tiny fraction of the mass, force timesteps so short that the calculation grinds to a halt.

A second challenge stems from the highly clustered spatial distribution of matter which affects in particular the scalability of parallel algorithms. A CDM halo is a near-monolithic, highly concentrated structure with a well-defined centre and no obvious geometrical decomposition which can separate it into the large number of computationally equivalent domains required for optimal exploitation of the many processors available in high-performance parallel architectures. In addition, gravity couples the dynamics of matter throughout the halo and beyond, requiring efficient communication between all parts of the simulated region. In our calculation, the clustering is so extreme that about one third of all simulation particles collect in a region that encompasses less than a fraction of 10^{-8} of the simulated volume!

Calculation Method and Parallelization Techniques

To make the Aquarius Project possible on the HLRB II, we have developed a major new version of our simulation code, GADGET-3, in order to improve scalability and performance for this extremely tightly coupled problem. GADGET uses a hierarchical multipole expansion (organized in a "tree") to calculate gravitational forces. In this method, particles are hierarchically grouped, multipole moments are calculated for each node, and then the force on each particle is obtained by approximating the exact force with a sum over multipoles. A great strength of the tree algorithm is the near insensitivity of its performance to clustering of matter, its ability to adapt to arbitrary geometries of the particle distribution, and the absence of any intrinsic resolution limitation.

However, there are actually faster methods to obtain the gravitational fields on large scales. In particular, the particle-mesh (PM) approach based on Fourier techniques is probably the fastest approach to calculate the gravitational field on a homogeneous mesh. The obvious limitation of this method is that the force resolution cannot be better than the size of one mesh cell, and the latter cannot be made small enough to resolve all the scales of interest in cosmological simulations. In fact, in our application we would need a mesh with $(10{,}000{,}000)^3$ cells to deliver the desired resolution with a single PM mesh - storing such a mesh would require several million petabytes!

GADGET therefore uses a compromise between the two methods. The gravitational field on large scales is calculated with a Particle-Mesh (PM) algorithm, while the short-range forces are delivered by the tree, such that a very accurate and fast gravitational solver results. A central role in our parallel code is played by the domain decomposition. It has to split the problem into smaller parts without data duplication, in a way that ensures a good balance of the computational work induced for each processor. GADGET uses a space-filling self-similar fractal, a Peano-Hilbert curve, for this purpose, which is made to become finer in high-density regions.

The domains themselves are then generated by cutting the one-dimensional space-filling curve into N_{cpu} pieces that approximately induce the same computational work-load. The domains are of nearly arbitrary shape but always have a relatively small surface-to-volume ratio.

First Results

As part of our Aquarius Project on the HLRB II, we have carried out extensive numerical resolution tests where we systematically increased the particle number used to simulate the same galaxy. Our primary simulation series culminated in our largest production calculation, which we refer to as C02-2400. This simulation followed about 4,5 billion particles, from a time briefly after the Big Bang to the present epoch, over more than 13 billion years of cosmic evolution. About 1,3 billion particles end up in the virialized region of a single, Milky-Way-sized galaxy, opening up a qualitatively new regime for studying the non-linear phase-space structure of dark matter halos.

An impression of the dynamic range of the simulation is given in Figure 1. What is readily apparent is the fascinating richness of dark matter substructure revealed by the simulation, which resolves several hundred thousand gravitational bound clumps of matter that orbit within the galaxy's potential. However, this extraordinary dynamic range comes at a price. More than 100,000 timesteps on 1,024 cores of the HLRB II and about 3 TB of RAM were required to carry out the simulation. In sum, the total CPU-time needed for completion of the C02-2400 calculation was nearly 4 million hours. The output produced forms a rich dataset of 45 terabytes in size, and will provide an extremely valuable scientific resource for many years to come.

Figure 2 shows spherically averaged density profiles obtained for the different resolutions that we calculated for our "C02" Milky Way halo. The convergence is excellent over the entire range where convergence can be expected based on the numerical parameters of the simulations. For the first time, our simulation series probes directly into a regime where the local logarithmic slope of the density profile of the dark matter cusp becomes shallower than -1. The structure of the cusp is of fundamental importance for our understanding of the CDM model, but has remained a highly contentious issue up to now. Our results demonstrate convincingly that an asymptotic power law of fixed slope apparently does not exist. Instead, the profile continues to become gradually shallower at ever smaller radii.

Our simulations also provide the first accurate and numerically converged measurement of the density profile of dark matter substructures, and they deliver precise predictions for the abundance of dark matter substructures. Using the simulation, we can obtain accurate determinations of the dark matter annihilation signal expected from the Milky Way's halo, which becomes potentially measurable with the launch of the GLAST gamma-ray satellite later this year. The simulation will also help to improve our understanding of galaxy formation and, in particular, the role and physics of satellite galaxies in the Milky Way's halo.

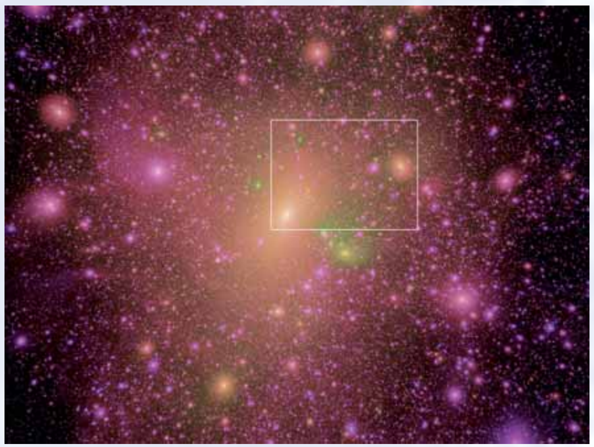


Figure 1: Dark matter distribution in the C02-2400 halo at the present time, showing the virialized region of the main halo and its immediate surroundings. The smaller picture enlarges the marked region by a factor of 4.

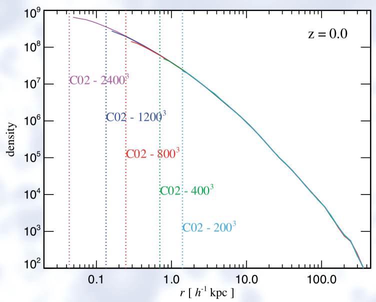

Figure 2: Spherically averaged density profiles of simulations of the same object carried out at different numerical resolution within the Aquarius Project. The dashed vertical lines mark the gravitational resolution limit of the individual calculations.

Thanks to the powerful HLRB II computer, the Aquarius Project produced the best resolved calculation of the Milky Way's halo carried out worldwide, and the ongoing analysis delivers many new insights for theoretical astrophysics. However, computational astrophysics in this area has still many exciting challenges in store for the future. Ultimately we would like to repeat our calculation by not only including the dark matter, but also the ordinary baryons that make up the stars, at similar or even better resolution. Carrying out such ultra-highly resolved hydrodynamical cosmological simulations will require further progress in the scalability of our codes, and the use of yet larger numbers of compute cores.

References

[1] **Barnes, J., Hut, P.** A Hierarchical O(N logN) Force-Calculation Algorithm, 1986, Nature, 324, 446

[2] **Davis, M., Efstathiou, G., Frenk, C. S., White, S. D. M.** The evolution of large-scale structure in a universe dominated by cold dark matter, 1985, Astrophysical Journal, 292, 371

[3] **Diemand J., Kuhlen M., Madau P.** Dark Matter Substructure and Gamma-Ray Annihilation in the Milky Way Halo, Astrophysical Journal, 657, 262, 2007

[4] **Navarro J. F., Frenk C. S., White S. D. M.** The Structure of Cold Dark Matter Halos, Astrophysical Journal, 462, 563, 1996

[5] **Springel, V.** The cosmological simulation code GADGET-2, Monthly Notices of the Royal Astronomical Society, 364, 1105, 2005

[6] **Springel, V., White S. D. M., Jenkins A., et al.** Simulations of the formation, evolution and clustering of galaxies and quasars, Nature, 435, 629, 2005

- Volker Springel[1]
- Simon D.M. White[1]
- Julio Navarro[2]
- Adrian Jenkins[3]
- Carlos S. Frenk[3]
- Amina Helmi[4]

[1] Max-Planck-Institute for Astrophysics, Garching

[2] University of Victoria

[3] University of Durham, Institute for Computational Cosmology

[4] University of Groningen, Kapteyn Astronomical Institute

Applications