The Millennium-XXL Project: Simulating the Galaxy Population of dark Energy Universes

Applications

Modern cosmology as encoded in the leading ΛCDM model confronts astronomers with two major puzzles. One 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. The other is a mysterious dark energy force field, which dominates the energy content of today's Universe and has led to its accelerated expansion in recent times. In the standard model, this component is described in terms of Einstein's cosmological constant ('Λ'). Uncovering the nature of dark energy has become one of the most actively pursued goals of observational cosmology.

In particular, the arrival of the largest galaxy surveys ever made is imminent, offering enormous scientific potential for new discoveries. Experiments like SDSSIII/BOSS or PanSTARRS have started to scan the sky with unprecedented detail, considerably improving the accuracy of existing cosmological probes. This will likely lead to challenges of the standard ΛCDM paradigm for cosmic structure formation, and perhaps even discover new physics.

One of the most important aims of these galaxy surveys is to shed light on the nature of the dark energy via measurements of the redshift evolution of its equation of state. However, the ability of these surveys to achieve this major scientific goal crucially depends on an accurate understanding of systematic effects and on a precise way

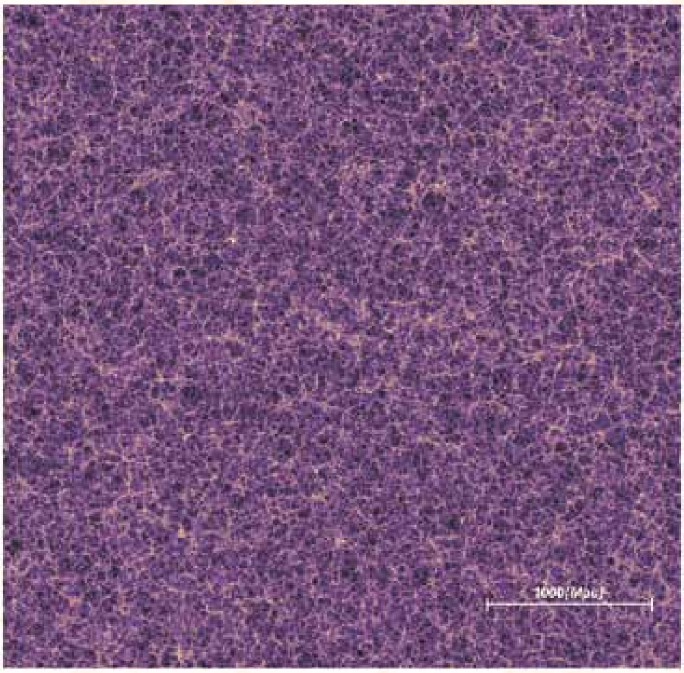

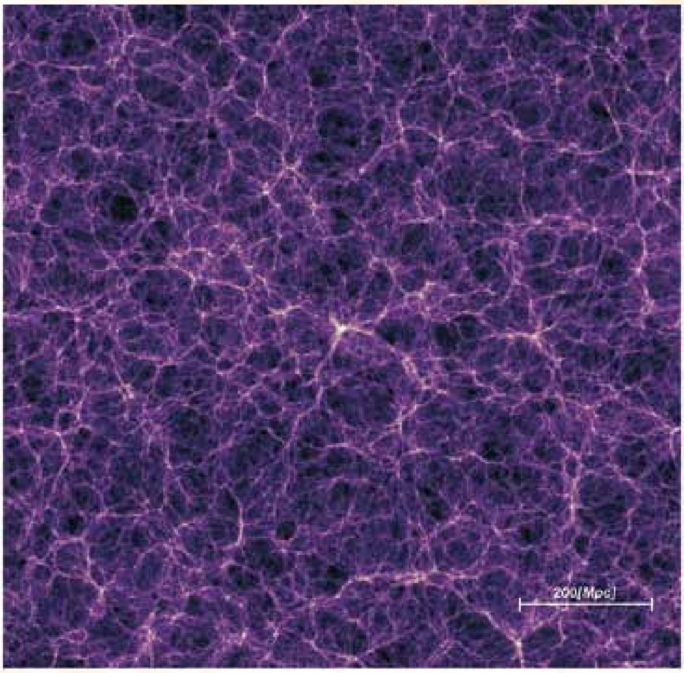

to physically model the observations, in particular the scale-dependent bias between luminous red galaxies and the underlying dark matter distribution, or the impact of mildly non-linear evolution on the so-called baryonic acoustic oscillations (BAOs) measured in the power spectrum of galaxy clustering.

Simulations of the galaxy formation process are arguably the most powerful technique to accurately quantify and understand these effects. However, this is an extremely tough computational problem, because it requires ultra-large volume N-body simulations with sufficient mass resolution to identify the halos likely to host the galaxies seen in the surveys, and a realistic model to populate these halos with galaxies. Given the significant investments involved in the ongoing galaxy surveys, it is imperative to tackle these numerical challenges to ensure that accurate theoretical predictions become available both to help to quantify and understand the systematic effects, and to extract the maximum amount of information from the observational data.

The State of the Art

The N-body method for the collisionless dynamics of dark matter is a long-established computational technique used to follow the growth of cosmic structure through gravitational instability. The Boltzmann-Vlasov system of equations is here discretized in terms of N fiducial simulation particles, whose motion is followed under their mutual gravitational forces in an expanding background space-time. While conceptually simple, calculating the long-range gravitational forces exactly represents an N^2-problem, which quickly becomes prohibitively expensive for interesting problem sizes. However, it is fortunately sufficient to calculate the forces approximately, for which a variety of algorithms have been developed over the years.

This allowed the sizes of cosmological simulations to steadily increase since the early 1980s, roughly doubling the particle number every 17 months and hence providing progressively more faithful models for the real Universe [1, 2, 3, 4]. Such simulations have proven

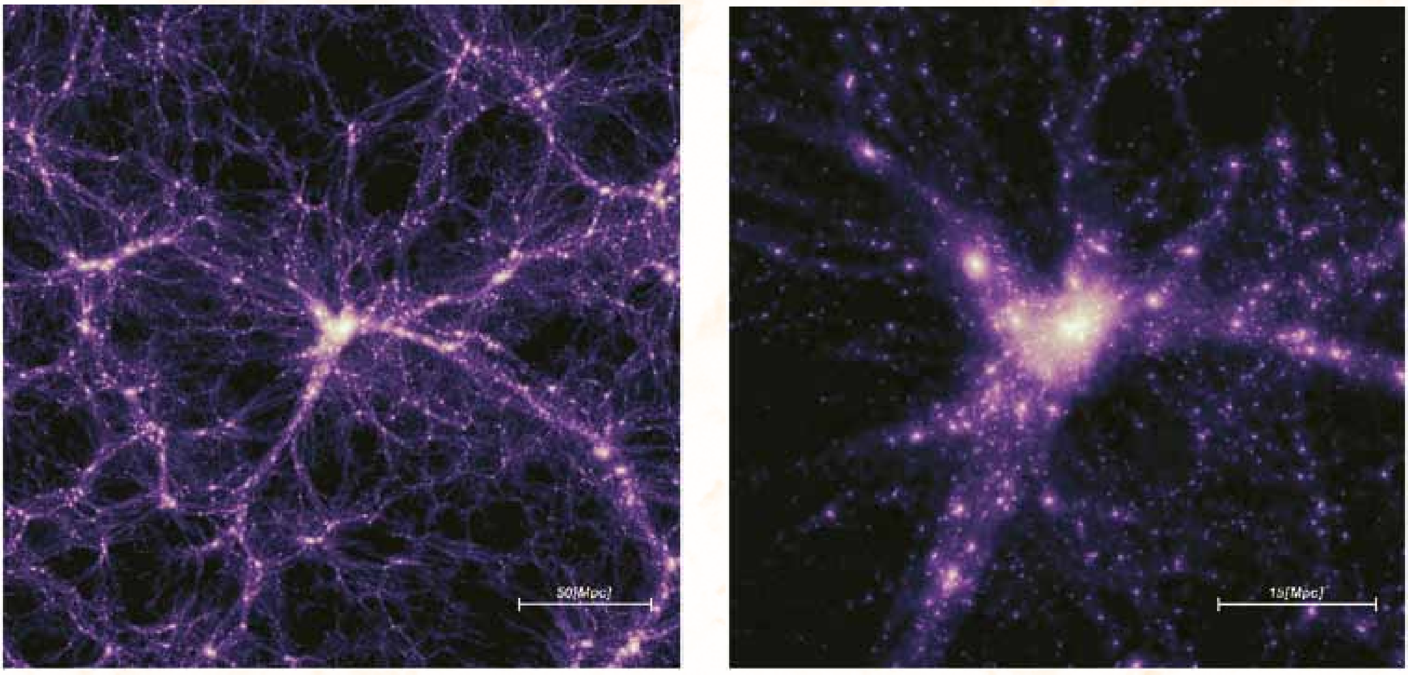

Figure 1: Dark matter distribution in the MXXL simulation on different scales. Each panel shows the projected density of dark matter in a slice of thickness 20 Mpc.

to be an indispensable tool to understand the low- and high-redshift Universe by comparing the predictions of CDM to observations, since these calculations are the only way to accurately calculate the outcome of non-linear cosmic structure formation.

A milestone in the development of cosmological simulations was set by the Millennium Run (MR), performed by our Virgo Consortium group in 2004 [3]. This simulation was the first, and for many years the only run with more than 10^{10} particles, exceeding the size of previous simulations by almost an order of magnitude. Its success was not only computational but most importantly scientific – more than 300 research articles in the fields of theoretical and observational cosmology have used the MR data-products since. The MR has an exquisite mass resolution and accuracy but, unfortunately, its volume is insufficient to get reliable statistics on large scales at the level needed for future surveys.

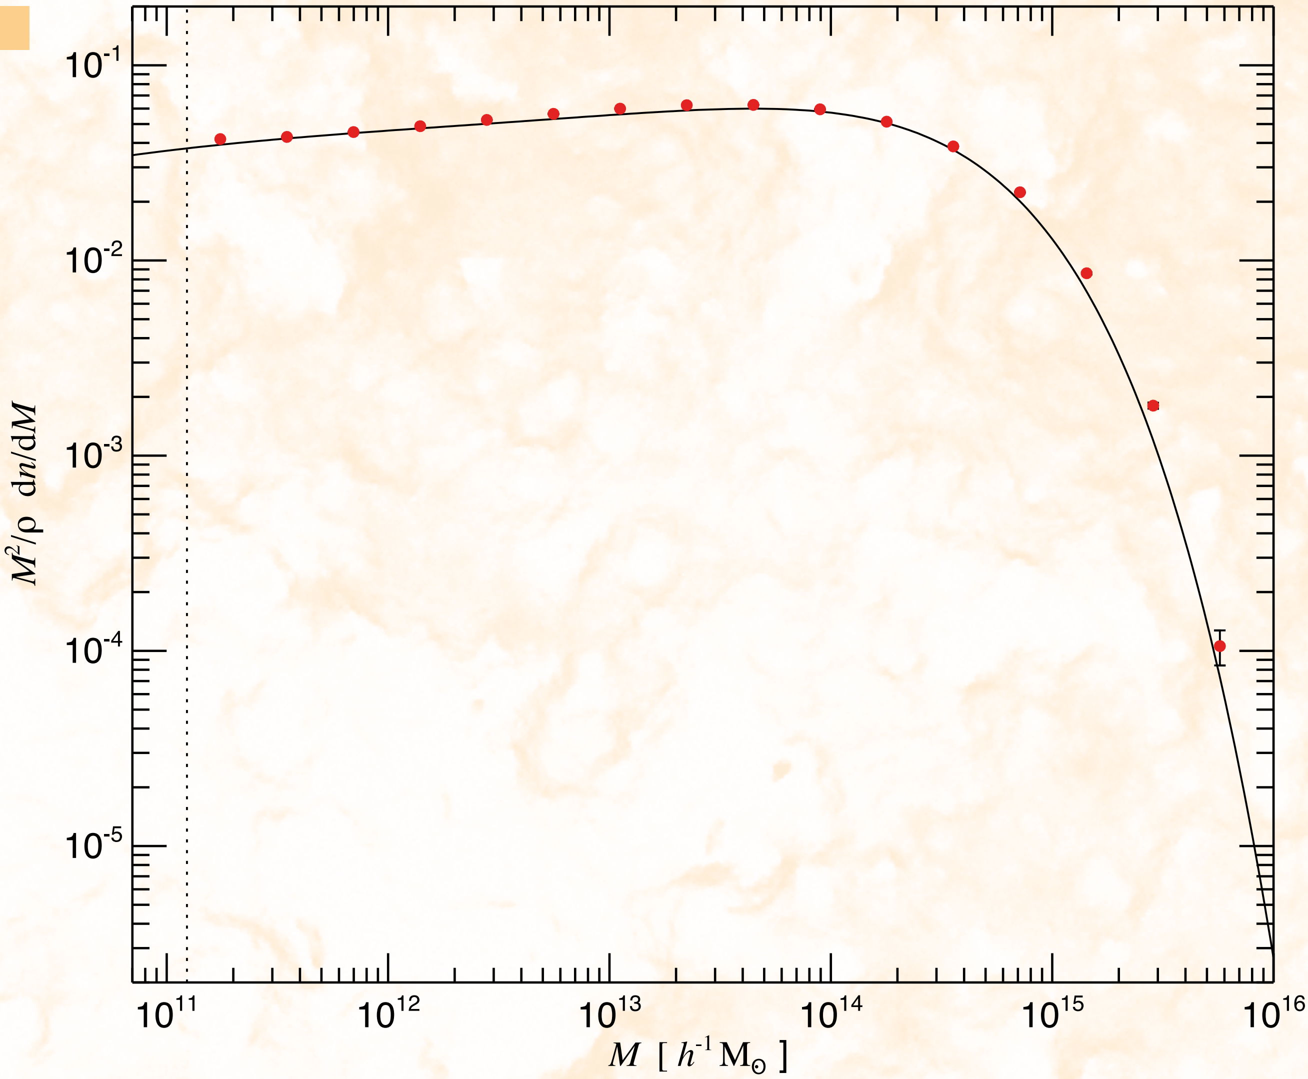

Figure 2: Differential halo abundance as a function of mass at the present epoch in the MXXL simulation. Note the good sampling far into the exponential tail. Apart from the last point, the error bars from counting statistics are smaller than the plotted symbols.

Recently, French and Korean collaborations [5, 6] have successfully carried out simulations containing almost 70 billion particles, but at considerably worse mass and spatial resolution than the MS, which did not allow them to robustly identify Milky-Way sized halos. Also, the need to manipulate and analyze a huge volume of data has proven to be a non-trivial challenge in working with simulations of this size, a fact that has made scientific exploitation of these projects difficult.

We have therefore set out to perform a new ultra-large N-body simulation of the hierarchical clustering of dark matter, featuring a new strategy for dealing with the data volume, and combining it with semi-analytical modelling of galaxy formation, which allows a prediction of all the luminous properties of the galaxies that form in the simulation. We designed the simulation project, dubbed Millennium-XXL (MXXL), to follow more than 303 billion particles (6720^3) in a cosmological box of size 4.2 Gpc across, resolving the cosmic web with an unprecedented combination of volume and resolution. While the particle mass of the MXXL is slightly worse than that of the MR, its resolution is sufficient to accurately measure dark matter merger histories for halos hosting luminous galaxies, within a volume more than 200 times larger than that of the MS. In this way the simulation can provide extremely accurate statistics of the large-scale structure of the Universe by resolving around 500 million galaxies at the present epoch, allowing for example highly detailed clustering studies based on galaxy or quasar samples selected in a variety of different ways. This comprehensive information is indispensable for the correct analysis of future observational datasets.

The computational Challenge

However, it was clear from the start that performing a simulation with these characteristics poses severe challenges, involving raw execution time, scalability of the algorithms employed, as well as their memory consumption and the disk space required for the output data. For example, simply storing the positions and velocities of the simulation particles in single precision consumes of order 10 TB of RAM memory. This figure, of course, is greatly enlarged by the extra memory required by the complex data structures and algorithms employed in the simulation code for the force calculation, domain decomposition, and halo and subhalo finding.

The code we used is a novel version of GADGET-3 we wrote specifically for the MXXL project. GADGET computes short-range gravitational forces with a hierarchical tree algorithm, and long-range forces with a FFT-based particle-mesh scheme [7]. Both the force computation and the time stepping of GADGET are fully adaptive. The code is written in highly portable C and uses a spatial domain decomposition to map different parts of the computational domain to individual processors. In 2005, we publicly released GADGET-2, which presently is the most widely employed code for cosmic structure formation.

Ultimately, memory requirements are the most serious limiting factor for cosmological simulations of the kind studied here. We have therefore put significant effort in developing algorithms and strategies that minimize memory consumption in our code, while at the same time retaining high integration accuracy and calculational speed.
The special "lean" version of our code we developed for MXXL requires only 72 bytes per particle in the peak for the

ordinary dynamical evolution of MXXL. The sophisticated in-lined group and substructure finders add a further 26 bytes to the peak memory consumption. This means that the memory requirement for the target size of 303 billion particles of our simulation amounts to slightly more than 29 TB of RAM.

The JuRoPa machine at the Jülich Supercomputing Centre (JSC) appeared as one of the best suited supercomputers within Germany to fulfill our computing requirements, thanks to its available storage of 3 GB per compute core. Still, on JuRoPa the simulation demanded a partition of 1,536 nodes, each equipped with two quad-core X5570 processors and 24 GB of RAM, translating to 12,288 cores in total. This represents a substantial fraction (70%) of the whole supercomputer installation. In fact, normal user operation on such a large partition is still unusual and not done on a regular basis yet. It hence required substantial support on the side of the system administrators at JSC to nevertheless allow us carry out the MXXL production calculation on JuRoPa. In particular, severe problems with the memory usage of our simulation were encountered initially, as the code required essentially all the available physical memory of the compute nodes, leaving very little room for memory buffers allocated by the MPI library or the parallel Lustre filesystem attached to JuRoPa. Fortunately, with help from JSC and ParTec, the software maker behind the MPI software stack on JuRoPa, these problems could be overcome.

For the MXXL simulation, we decided to try a new hybrid parallelization scheme that we recently developed for our GADGET code. Instead of using 12,288 distributed-memory MPI tasks, we only employed one MPI task per processor (3,072 in total), exploiting the 4 cores of each socket via threads. This mixture of distributed and shared memory parallelism proved ideal to limit memory and work-load imbalance losses, because a smaller number of MPI tasks reduces the number of independent spatial domains in which the volume needs to be decomposed, as well as the data volume that needs to be shuffled around with MPI calls. The thread-based parallelization itself was partially done via explicit Posix thread programming, and partially via OpenMP directives for the simpler parallel constructs. The Posix-threads approach has allowed us to achieve effectively perfect parallelizability of the gravitational tree-walk in our code, which dominates the computational cost.

Still, the total CPU time required for the simulation was substantial. This originates from both, the huge number of particles and the very dissimilar time scales and matter densities involved in the problem. Very different density configurations are encountered in the clustered state produced by the simulation, and consequently there are structures featuring a large variety of timescales. Our code tackles this problem by using spatially and temporally adaptive time-stepping, so that short time-steps are used only when particles enter the localized dense regions where dynamical times are short. In addition, we use a new domain decomposition scheme that produces almost ideal scaling on massively parallel computers. In spite of these optimizations, the scope of the computational problem remained huge. The final production run of MXXL required more than 86 trillion force calculations and took slightly more than 2.7 million CPU hours (~300 years), corresponding to about 9.3 days of runtime on 12,288 cores.

A significant fraction of 15% of the time was however spent for running our sophisticated on-the-fly postprocessing software, notably the group finding, the substructure finding, and the power spectrum calculation, and another 14% were needed for the I/O operations. The long-range force calculations based on 9216^3 sized FFTs consumed only about 3% of the time. We note that the parallel "friends-of-friends" group finding for the 303 billion particles at the final output time took just 470 seconds, which we think is a remarkable achievement. Determining all gravitational bound substructures within these halos took a further 2,600 seconds. It is also during this substructure finding that the peak memory consumption of the code is reached.

Another demanding aspect of the MXXL project lies in the data handling. The production code is able to write data parallel to disk. Despite the huge size of our data sets, the restart files, amounting to 10 TB in total, could be written/read to disk in about 40 minutes on the JuRoPa system. However, of larger practical importance for us is the aggregated size of the simulation outputs produced by the simulation. If an approach similar to the MR was taken for the MXXL project, then the disk space requirement would approach 700 TB for the storage of the dark matter particle positions, velocities and groups catalogues, a prohibitively large amount of data. To cope with this problem, we have therefore carried out a significant part of the post-processing on-the-fly during the simulation, avoiding the need to store the particle phase-space variables for most of the output times on disk. This postprocessing included the identification and characterization of halos and subhalos, and the calculation of the dark matter density field and its time derivative. As a result, the total disk space required for the data products of the simulation could be shrunk to about 100 TB.

First Results and Outlook

In Figure 1, we show the projected density field of the MXXL simulation in a series of thin slices through the simulation volume. This gives a hint about the enormous statistical resolving power of the simulation, and its large dynamic range everywhere in the simulation cube. The gravitational spatial resolution of the simulation translates to a dynamic range of 300,000 per dimension, or formally more than 10^{16} resolution elements in 3D.

We have found 650 million halos at redshift zero in the simulation, binding 44% of all particles in non-linear structures. In total there are more than 25 billion halos in the group catalogues that we produced and that are linked together in our merger trees. In Figure 2, we show the halo mass function at the present epoch, which is one of the most basic and important simulation predictions, describing the abundance of objects as a function of mass and epoch.

The largest cluster of galaxies at z=0 has a mass of 9 x 10^{15} solar masses. Such extreme objects are so rare that they can only be found in volumes as large as that of the MXXL. In Figure 3, we show a map of the time derivative of the gravitational potential of the simulation at redshift z=12. We produced such maps on-the-fly during the simulation in order to study the integrated Sachs-Wolfe effect, which describes distortions in the cosmic microwave background by foreground structures.

The full analysis of the MXXL in terms of its predicted galaxy distributions has just begun, and will still take some time to complete. We can, for example, use the MXXL to predict the spatial distribution of galaxies in the past light-cone of a fiducial observer. Without any replication of the box, we will be able to construct all-sky synthetic galaxy catalogues up to redshift 0.57. These mock observations are crucial for future surveys to study cluster finding algorithms, the influence of projection effects, the evolution of the clustering along the line-of-sight, and most importantly,

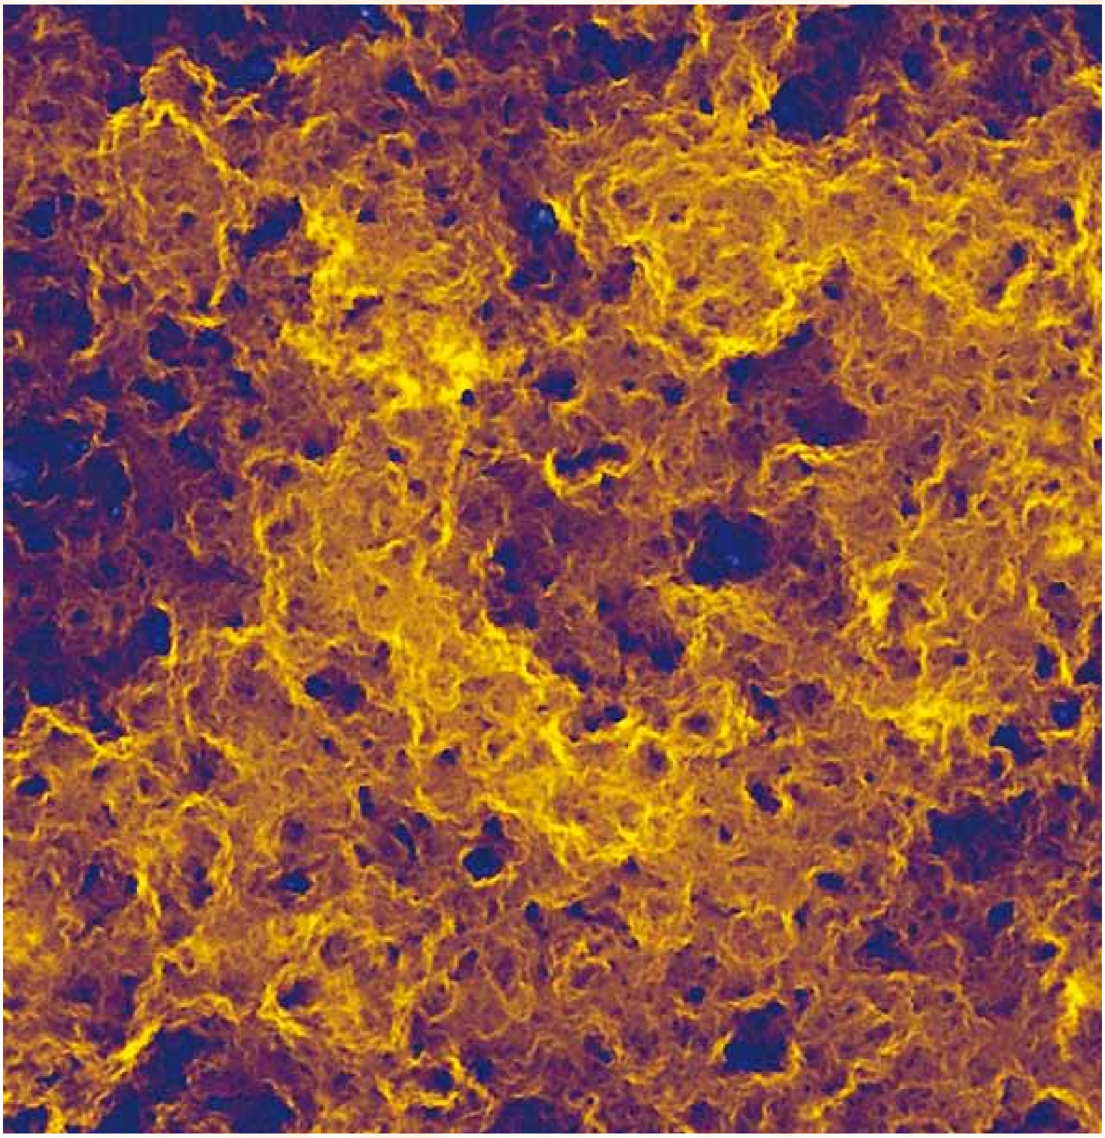

Figure 3: Map of the time-derivative of the gravitational potential at z=12 in the MXXL simulation. Such maps can be used to calculate the integrated Sachs-Wolfe effect along the backwards light-cone to the last scattering surface of the cosmic microwave background.

the systematic effects in the cosmological constraints derived from BAO, weak lensing, or cluster counts. Using a novel scaling technique we just developed [8], it is furthermore possible to cast the simulation results into theoretical models for the galaxy distribution that cover a significant range of cosmologies.

The Millennium XXL is by far the largest cosmological simulation ever performed and the first multi-hundred billion particle run. The scope of the simulation project has pushed the envelope of what is feasible on current world-class HPC resources, but the expected rich scientific return from the project makes it well worth the effort. At the same time, the successful scaling of the cosmological code to well beyond ten-thousand cores is an encouraging sign for computational research in cosmology, which will address yet more demanding problems on future HPC machines.

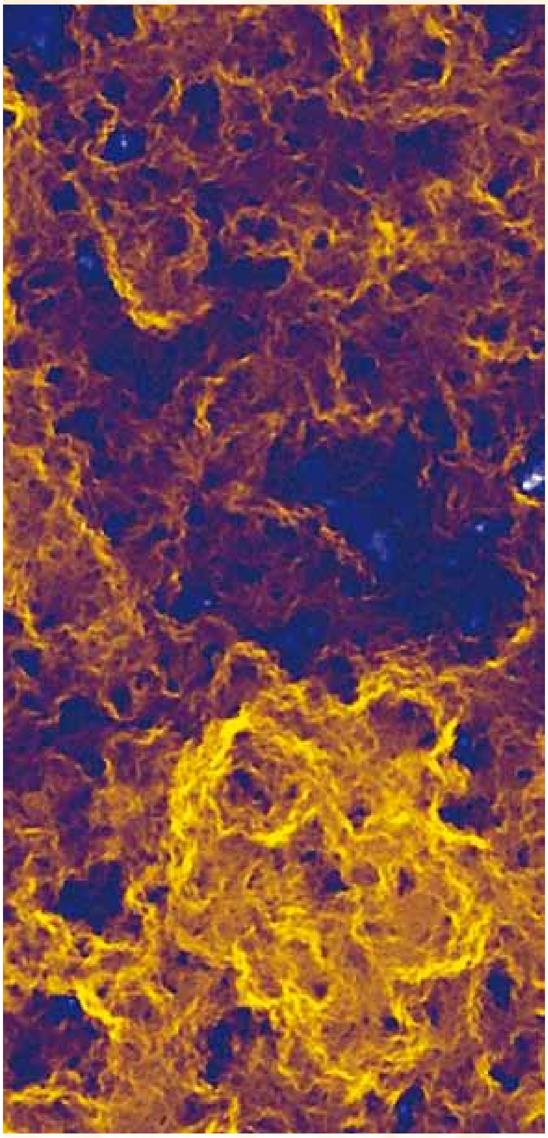

References

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

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

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

[4] Springel, V., Wang, J., Vogelsberger, M., et al.
2008, The Aquarius Project: the subhaloes of galactic haloes, Monthly Notices of the Royal Astronomical Society, 391, 1685

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

[6] Teyssier, R., Pires, S., Prunet, S., et al.
2009, Full-sky weak-lensing simulation with 70 billion particles, Astroparticle Physics, 497, 335

[7] Kim, J., Park, C., Gott, J. R., Dubinski, J.
2009, The Horizon Run N-Body Simulation: Baryon Acoustic Oscillations and Topology of Large-scale Structure of the Universe, Astrophysical Journal, 701, 1547

[8] Angulo, R. E., White, S. D. M.
2010, One simulation to fit them all - changing the background parameters of a cosmological N-body simulation, Monthly Notices of the Royal Astronomical Society, 405, 143

- Volker Springel[1]
- Raul Angulo, Simon D.M. White[2]
- Adrian Jenkins, Carlos S. Frenk, Carlton Baugh, Shaun Cole[3]

[1] Heidelberg Institute for Theoretical Studies, Heidelberg

[2] Max-Planck-Institute for Astrophysics, Garching, Germany

[3] Institute for Computational Cosmology, University of Durham, United Kingdom