



The Next Generation of Hydrodynamical Simulations of Galaxy Formation

Scientific background

Galaxies are comprised of up to several hundred billion stars and display a variety of shapes and sizes. Their formation involves a complicated blend of astrophysics, including gravitational, hydrodynamical and radiative processes, as well as dynamics in the enigmatic „dark sector“ of the Universe, which is composed of dark matter and dark energy. Dark matter is thought to consist of a yet unidentified elementary particle, making up about 85% of all matter, whereas dark energy opposes gravity and has induced an accelerated expansion of the Universe in the recent past. Because the governing equations are too complicated to be solved analytically, numerical simulations have become a primary tool in theoretical astrophysics to study cosmic structure formation. Such calculations connect the comparatively simple initial state left behind by the Big Bang some 13.6 billion years ago with the complex, evolved state of the Universe today. They provide detailed predictions for testing the cosmological paradigm and promise transformative advances in our understanding of galaxy formation.

One prominent example of such a simulation is our Illustris simulation from 2014 [1, 2]. It tracked the small-scale evolution of gas and stars within a representative portion of the Universe, using more than 6 billion hydrodynamical cells and an equally large number of dark matter particles. Illustris yielded for the first time a reasonable morphological mix of thousands of well-resolved elliptical and spiral galaxies. The simulation reproduced the observed distribution of

galaxies in clusters and the characteristics of hydrogen on large scales, and at the same time matched the metal and hydrogen content of galaxies on small scales. Indeed, the virtual universe created by Illustris resembles the real one so closely that it can be adopted as a powerful laboratory to further explore and characterize galaxy formation physics. This is underscored by the nearly 100 publications that have been written using the simulation thus far.

However, the Illustris simulation also showed some tensions between its predictions and observations of the real Universe, calling for both, improvements in the physical model as well as in the numerical accuracy and size of the simulations used to represent the cosmos. For example, one important physical ingredient that was missing are magnetic fields. In fact, our Universe is permeated with magnetic fields – they are found on Earth, in and around the Sun, as well as in our home galaxy, the Milky Way. Often the magnetic fields are quite weak, for example, the magnetic field on Earth is not strong enough to decisively influence the weather on our planet. In galaxies like the Milky Way, the field is however so strong that its pressure on the interstellar gas in the galactic disc is of the same size as the thermal pressure. This suggests that magnetic fields could play an important role in regulating star formation, but their origin still remains mysterious.

Another challenge that had become clear in our past work is that the regulation of star formation in massive galaxies through the energy output of growing supermassive black holes was not

adequately described by Illustris. Even though the adopted model was so strong and violent that it caused an excessive depletion of the baryon content of galaxy groups and low-mass galaxy clusters, it proved insufficient to reduce the star formation in the central galaxies in these systems to the required degree, causing these galaxies to become too massive. In essence, this showed a serious failure of the underlying theoretical model. Fixing it requires replacing it with something considerably different.




Fig. 1: Thin projections through one of our Illustris TNG simulations, showing different physical quantities as labelled.

We thus felt compelled to work on a new generation of simulations with the goal of advancing the state-of-art on all of these fronts by using a new comprehensive model for galaxy formation physics and updated numerical treatments. At the same time, we aimed for higher numerical resolution, larger volume covered and hence better statistics, as well as an improved accuracy in our hydrodynamical solvers. Using Hazel-Hen and a GCS grant for computer time, we were able to succeed on many of these aims, and to produce a novel, scientifically very interesting set of simulation models, which we now call “The Next Generation Illustris Simulations” (IllustrisTNG).

New physics modelling

For IllustrisTNG, we developed a new kinetic feedback model for AGN driven winds [3], motivated by recent theoretical suggestions that conjecture advection dominated inflow-outflow solutions for the accretion flows onto black holes in the low accretion rate regime. In terms of energetic feedback, we distinguish between a quasar mode for high accretion rates where the feedback is purely thermal, and a kinetic mode for low accretion rate states where the feedback is purely kinetic. The distinction between the two feedback modes is based on the Eddington ratio of the black hole accretion. In the kinetic feedback state, strong quenching of cooling flows and star formation in the host halo is possible, such that the corresponding galaxy can quickly redden.

Another important change we made relates to the modelling of galactic winds and outflows



[4], which now scale differently with the Hubble rate, and also take metallicity-dependent cooling losses better into account. The net effect of this is a stronger suppression of star formation in small galaxies, yielding an improved faint-end of the galaxy luminosity function.

Importantly, we have also added magnetic fields to our simulations, using a new implementation of ideal magnetohydrodynamics in our AREPO code [5, 6]. This opens up a rich new area of predictions that are still poorly explored, given that the body of cosmological magneto-hydrodynamic simulations is still very small. In particular, it allows us to study the strength of magnetic field amplification through structure formation as a function of halo mass and galaxy type.

For IllustrisTNG, we have also improved our modelling of chemical enrichment, both by using updated yield tables that account for the most recent results of stellar evolution calculations, and by making the tracking of different chemical elements (H, He, C, N, O, Ne, Mg, Si, Fe) more accurate and informative. For example, we developed a special chemical tagging method that separately accounts for metals produced by asymptotic giant branch stars, type-II supernovae, and type-Ia supernovae. This has not been done before in such hydrodynamical simulations.

Finally, we also developed a novel hierarchical timestepping scheme in our AREPO code that solves this in a mathematically clean fashion. This is done by recursively splitting the

Hamiltonian describing the dynamics into a slow and a fast system, with the fast system being treated through sub-cycling. An important feature of this time integration scheme is that the split-off fast system is self-contained, i.e. its evolution does not rely on any residual coupling with the slow part. This means that poorly populated short timesteps can be computed without touching any parts of the system living on longer timesteps, making these steps very fast so that they not slow down the main calculation significantly.




Fig. 2: Two disk galaxies in halos of mass $8.3 \times 10^{11} M_{\odot}$ (left) and $2.0 \times 10^{12} M_{\odot}$ (right).

Computational challenge

The AREPO code [5] we developed for cosmological hydrodynamics uses a finite-volume approach on a three-dimensional, fully dynamic Voronoi tessellation. The moving mesh is particularly well suited to the high dynamic range in space and time posed by the galaxy formation problem. The very low advection errors of AREPO are very helpful for the highly

supersonic flows occurring in cosmology and for treating subsonic turbulence within the gas of virialized halos. These properties make it superior to smoothed particle hydrodynamics and adaptive mesh refinement codes that use a stationary Cartesian mesh. AREPO also follows the dynamics of dark matter with high accuracy, as required to compute cosmic structure growth far into the non-linear regime.

The simulations carried out in the IllustrisTNG project represent a significant challenge not only in terms of size and spatial dynamic range, but also in terms of the dynamic range in timescales. In particular, the strong kinetic feedback by black holes, which couples to the densest gas in galaxies, induces very small timesteps for a small fraction of the mass. Over the course of 13 billion years of cosmic evolution that we cover, we needed to do up to 10^7 timesteps in total. This would be completely infeasible with time integration schemes that employ global timesteps, but even for the new individual timestepping we have used in AREPO, this represents a formidable problem. It could only be tackled by making the computation of sparsely populated timesteps extremely fast so that they do not dominate the total CPU time budget.

In addition to the challenging dynamic range in timescales, we also aim for a larger number of resolution elements, and a larger simulation volume than realized previously. This is necessary to study the regime of galaxy clusters better (which are rare and can only be found in a sufficiently large volume), and to allow a sampling

of the massive end of the galaxy and black hole mass functions. The primary science runs of IllustrisTNG consist of two large full-physics calculations (and a third one targeting dwarf galaxies is underway in follow up work), each significantly more advanced and also larger than the older Illustris simulation. This is complemented with matching dark matter only simulations, as well as a series of lower resolution calculations to assess numerical convergence. The calculations include magneto-hydrodynamics and adopted the newest cosmological models as determined by the Planck Satellite.

We have used between 10752 and 24000 cores on Hazel-Hen, benefitting in full from the large memory, high communication bandwidth, high floating point performance and high I/O bandwidth of this machine. The peak memory consumption of our largest run is about 95 TB RAM, and each of our simulation time slices weighs in with several TB. In fact, we have already transferred more than 300 TB of final production data to the Heidelberg Institute of Theoretical Studies, in part by using fast gridftp services offered by HLRS.

First results and outlook

In Figure 1, we illustrate the large-scale distribution of different physical quantities in one of our IllustrisTNG simulations. From top to bottom, we show projections of the gas density field, the mean mass-weighted metallicity, the mean magnetic field strength, the dark matter density, and the stellar density. The displayed regions are about 350 million lightyears across from left

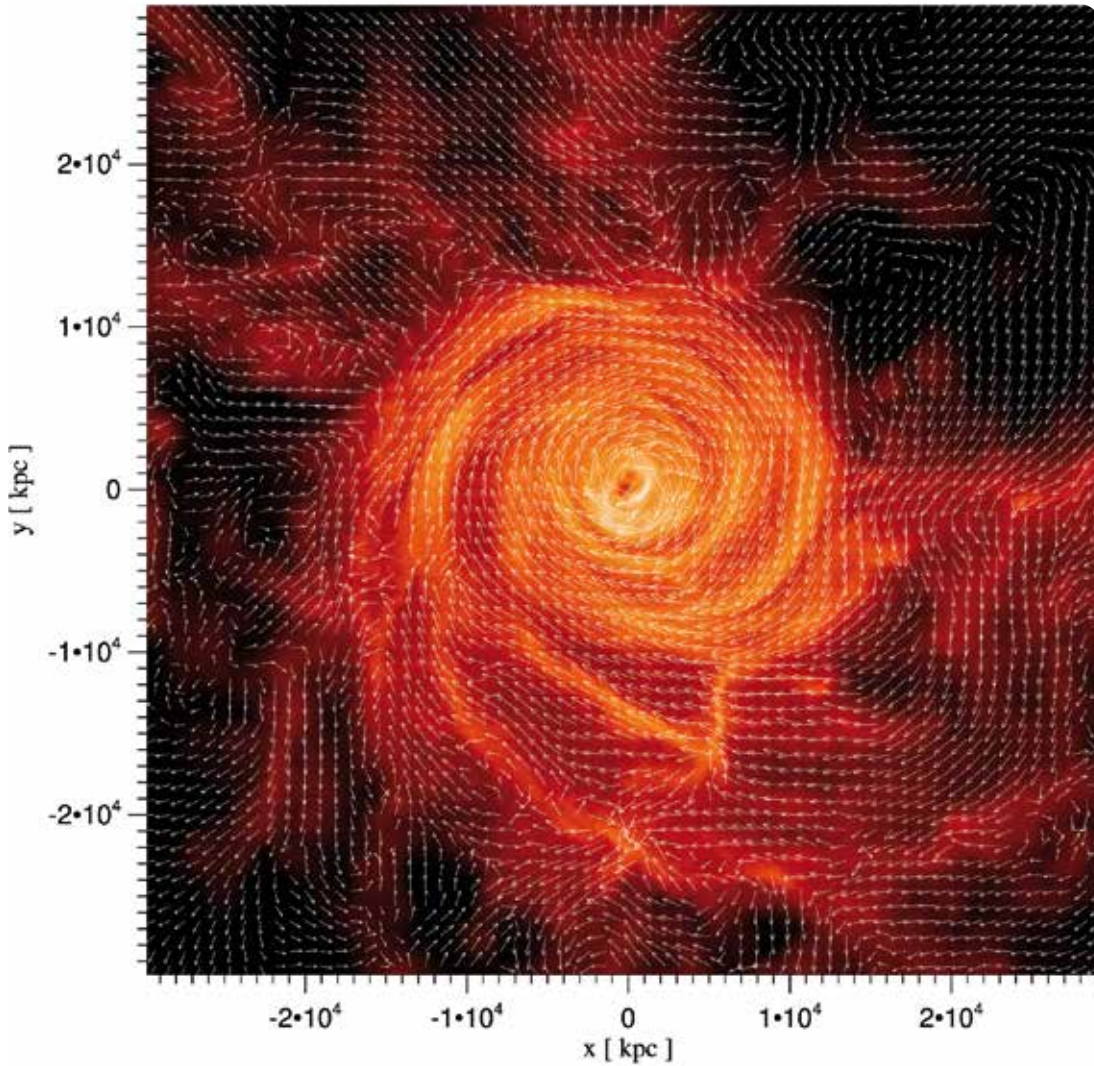


Fig. 3: Magnetic field direction (vectors) in the disk galaxy displayed on the left in Fig. 2, with the gas density structure shown in the background (color scale).

there is clear evidence for very strong outflows in them, causing widespread heating as they impinge on the gas in the intergalactic medium.

The rightmost panel in Figure 1 displays the stellar mass density. Clearly, on the scales shown in this image, the individual galaxies appear as very small dots, illustrating that the stellar component fills only a tiny fraction of the volume. However, our simulations have enough resolution and dynamic range to actually

to right. On large scales, the dark matter and the diffuse gas trace out the so-called cosmic web that emerges through gravitational instability. The color in the gas distribution encodes the mass-weighted temperature across the slice. The largest halos are filled with hot plasma, and

resolve the internal structure of these galaxies in remarkable detail. This is shown in Figure 2, which zooms in on two disk galaxies formed in our simulations. The one on the right hand panel is in a more massive halo and has a more massive black hole. This in fact has made it start to

transition into the quenched regime, which here begins by a reduced star formation in the center as a result of kinetic black hole feedback. The outskirts of the galaxy still support some level of star formation, causing blue spiral arms.

In Figure 3, we plot the magnetic vector field of this galaxy, overlaid on a rendering of the gas density in the background. We see that the field is ordered in the plane of the disk, where it has been amplified by shearing motions to sizable strength. Interestingly, there are multiple field reversals and a complicated topology of the field surrounding the disk. The realistic field topologies predicted here should be very useful for studying the propagation of cosmic rays in the Milky Way. Already now we can say that our calculations demonstrate that an extremely tiny magnetic field left behind by the Big Bang is sufficient to explain the orders of magnitude larger field strengths observed today. In fact, the field strengths we measure in our galaxies agree quite well with observational constraints.

In Figure 4, we show a break down of the total metal content in the gas phase of IllustrisTNG at different times as a function of gas density. The histograms are normalized to the total metal content in the gas at the corresponding epoch, so that the distributions inform about the question at which gas densities the majority of the metals can be found. Most of the metals are actually stored at gas densities that correspond to the circumgalactic medium, whereas only a smaller fraction is contained in the star-forming interstellar medium, and relatively little in the

low-density intergalactic medium. These distributions are shaped by the galactic winds in the simulation, and determining them observationally will provide powerful constraints on our theoretical models.

The scientific exploitation of the IllustrisTNG simulations has just begun. We expect that they will significantly expand the scientific possibilities and predictive power of hydrodynamical simulations of galaxy formation, thereby forming the ideal basis for the comparison with real data. Obtaining this valuable data has only been possible thanks to the power of the Hazel-Hen supercomputer, and mining the rich scientific results this data has in store will keep us and

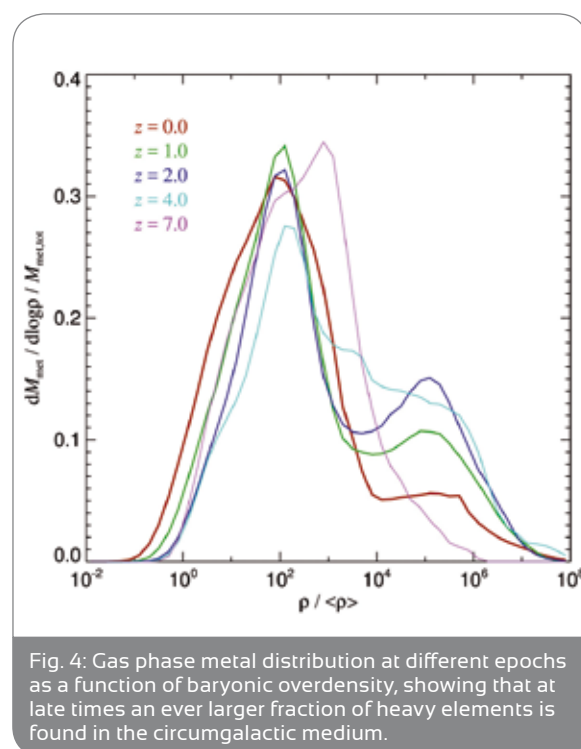


Fig. 4: Gas phase metal distribution at different epochs as a function of baryonic overdensity, showing that at late times an ever larger fraction of heavy elements is found in the circumgalactic medium.



many of our colleagues in the field busy for years to come.

Acknowledgements

We gratefully acknowledge support by the High Performance Computing Center in Stuttgart, and computer time through project GCS-ILLU on Hazel-Hen. We also acknowledge financial support through subproject EXAMAG of the Priority Programme 1648 SPPEXA of the German Science Foundation, and through the European Research Council through ERC-StG grant EXA-GAL-308037, and we would like to thank the Klaus Tschira Foundation.

References

- [1] **Vogelsberger M., Genel S., Springel V., Torrey P., Sijacki D., Xu D., Snyder G., Bird S., Nelson D., Hernquist L.:** Properties of galaxies reproduced by a hydrodynamic simulation, 2014, *Nature*, 509, 177
- [2] **Genel S., Vogelsberger M., Springel V., Sijacki D., Nelson D., Snyder G., Rodriguez-Gomez V., Torrey P., Hernquist L.:** Introducing the Illustris project: the evolution of galaxy populations across cosmic time, 2014, *Monthly Notices of the Royal Astronomical Society*, 445, 175
- [3] **Weinberger R., Springel V., Hernquist L., Pillepich A., Marinacci F., Pakmor R., Nelson D., Genel S., Vogelsberger M., Naiman J., Torrey P.:** Simulating galaxy formation with black hole driven thermal and kinetic feedback, 2017, *Monthly Notices of the Royal Astronomical Society*, 465, 3291
- [4] **Pillepich A., Springel V., Nelson D., Genel S., Naiman J., Pakmor R., Hernquist L., Torrey P., Vogelsberger M., Weinberger R., Marinacci F.:** Simulating Galaxy Formation with the IllustrisTNG Model, 2017, *Monthly Notices of the Royal Astronomical Society*, submitted, arXiv:1703.02970
- [5] **Springel V.:** E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh, 2014, *Astrophysical Journal Letters*, 401, 791
- [6] **Pakmor R., Marinacci F., Springel V.:** Magnetic Fields in Cosmological Simulations of Disk Galaxies, 2017, *Astrophysical Journal Letters*, 783, L20

Written by Volker Springel, Rainer Weinberger, Rüdiger Pakmor

Heidelberg Institute for Theoretical Studies and Heidelberg University, Germany

Annalisa Pillepich

Max-Planck Institute for Astronomy, Heidelberg, Germany

Dylan Nelson

Max-Planck-Institute for Astrophysics, Garching, Germany

Mark Vogelsberger, Federico Marinacci, Paul Torrey

Massachusetts Institute of Technology, Cambridge, USA

Lars Hernquist, Jill Naiman

Harvard University, Cambridge, USA

Shy Genel

Center for Computational Astrophysics, Flatiron Institute, New York, USA

Contact: Volker Springel, volker.springel@h-its.org