Astron. Nachr. / AN **333**, No. 5/6, 515 – 522 (2012) / **DOI** 10.1002/asna.201211701

Larger, faster, better: Current trends in cosmological simulations

V. Springel[1,2]⋆

[1] Heidelberger Institut für Theoretische Studien, Schloss-Wolfsbrunnenweg 35, 69118 Heidelberg, Germany
[2] Zentrum für Astronomie der Universität Heidelberg, Astronomisches Recheninstitut, Mönchhofstr. 12–14, 69120 Heidelberg, Germany

Received 2012 Apr 30, accepted 2012 May 2

Key words cosmology: theory – galaxies: formation – methods: numerical

Numerical simulations of galaxy formation provide a powerful technique for calculating the non-linear evolution of cosmic structure formation. In fact, they have played an instrumental role in establishing the current standard cosmological model known as ΛCDM. In this brief review, I highlight some of largest calculations carried out recently in the field, and I discuss the challenges involved in making further progress with simulations in the future. These entail the need to scale up the simulations to ever larger sizes and much better resolution, the need to develop more comprehensive numerical models of the physics of galaxy formation, and the need to improve the accuracy of the numerical techniques currently employed. The present epoch of petaflop supercomputers, and the exaflop machines that are foreseen towards the end of the decade, demand novel types of simulation codes that yet have to be developed.

1 Introduction

Galaxies exist in a bewildering variety of shapes and sizes, ranging from dwarf satellites in our own Milky Way, over grand design spirals to huge elliptical galaxies. At present, our understanding of galaxy formation is still sketchy, even though a basic paradigm for it exists – the theory of hierarchical galaxy formation within the ΛCDM cosmology. The fundamental challenge is that galaxy formation involves a complicated blend of different physics that is non-linearly coupled on a wide range of scales, leading to extremely complex dynamics. Of primary importance are self-gravity of dark matter and gas, collisionless dynamics of dark matter, hydrodynamical shocks in the gas, high Reynolds number turbulence, radiative cooling and heating processes, radiative transfer, star formation and evolution, non gravitational energy input into diffuse gas by supernovae or black hole accretion, magnetohydrodynamics, and possibly cosmic rays.

The set of partial differential equations describing this mix of physics is for the most part well known, but largely inaccessible by analytic techniques. For this reason, high-performance computing (HPC) techniques have become a primary avenue for theoretical research in galaxy formation. This is also helped by the fact that the current standard model of cosmology precisely specifies the initial conditions of cosmic structure formation at a time briefly after the Big Bang. It becomes then a computational problem par excellence to try to evolve this initial state forward in time, staying as faithful to the physics as possible.

As long as only the dark matter is followed in such cosmological simulations, the relevant computational methods have reached a high degree of maturity and led to substantial progress in understanding the non-linear regime of structure formation in dark matter. Projects such as the Millennium (Boylan-Kolchin et al. 2009; Springel et al. 2005) and Bolshoi simulations (Klypin et al. 2011), or the Via Lactea (Diemand et al. 2007), Aquarius (Springel et al. 2008) and GHALO (Stadel et al. 2009) zoom simulations have made dark matter simulations a precision science. They progressed to a state where they reliably predict scales much smaller than the ones affected by baryonic physics.

This predictive power of high-resolution dark matter simulations has however revealed a number of tensions in the ΛCDM framework, such as the satellite problem or the presence of central dark matter cusps. At the same time, these simulations have put our still inadequate understanding of galaxy formation into the limelight. This is for example revealed by comparing the predicted dark matter halo mass function with the accurately determined stellar mass function of the SDSS survey. Based on only the assumption that stellar mass monotonically increases with halo mass, an abundance matching analysis (Guo et al. 2010) shows that the stellar mass to halo mass ratio has a well defined peak at around $\sim 10^{12}\,\mathrm{M}_{\odot}$, and sharply declines towards lower and higher masses. But even at the peak, the efficiency of star formation is low; only a small fraction of the expected universal amount of baryons is turned into stars. In halos of lower mass, atomic cooling is even more efficient, yet a much smaller fraction of stars is turned into stars. Likewise, for more massive halos, at the group and cluster scales, we would in principle also expect much larger central galaxies.

⋆ Corresponding author: volker.springel@h-its.org

We have still a rather poor understanding of why the galaxy formation efficiency is so extremely low, and where all the missing baryons are at low redshift. It is commonly believed that a solution of this riddle lies in a more complete and accurate accounting of baryonic feedback processes, such as energy input by evolving stars or supernovae, radiative heating effects, quasar activity, or non-thermal processes associated with magnetic fields and cosmic rays, to name just the most common suspects.

However, current state-of-the-art hydrodynamical simulations still miss to reproduce this low star formation efficiency (Guo et al. 2010; Sawala et al. 2011), even though strong feedback processes are included. This is for example seen in the results for so-called zoom simulations on individual objects (e.g. Dolag et al. 2009; Governato et al. 2007; Scannapieco et al. 2009; Stinson et al. 2009). It is fair to say that none of these simulation models has yet been able to present a convincing model of feedback that could explain the low star-formation efficiency required to reconcile observations with ΛCDM, although incremental progress has clearly been made. The best results in terms of producing late-type galaxies akin to a Milky Way system (Guedes et al. 2011) have been obtained with a delayed-cooling feedback model. However, the recent "Aquila" comparison project (Scannapieco et al. 2011) has also shown limitations of this feedback model as it is not universally able to produce reasonable galaxies in all dark matter halos.

In comparison, semi-analytic models of galaxy formation have been able to produce a globally much more successful description of galaxy formation (e.g. Croton et al. 2006; Guo et al. 2011), based on rather drastic simplifications of the baryonic physics and ad-hoc parameterizations of feedback efficiencies. While highly useful for testing the overall validity of the cosmological paradigm, these models give only very limited insight into the detailed hydrodynamical processes that regulate star formation in galaxies and are unable to describe the diffuse gas in and between galaxies. In order to transcend the predictive power of these theoretical models, it is necessary to directly simulate the baryonic physics as well.

Broadly speaking, future progress in simulations of galaxy formation requires advances in three important areas:

1. Resolution and size of the simulations.
2. Accuracy and efficiency of the numerical algorithms.
3. Realistic accounting of all the relevant physics.

In this brief review, I discuss a few of the challenges involved and highlight a number of recent works along this line, largely in an exemplary fashion. In Sect. 2, I first introduce two recent large cosmological simulations. I then discuss issues of hydrodynamical code accuracy in Sect. 3, and mention some aspects of physics extensions in Sect. 4. I conclude with a brief summary in Sect. 5.

2 Larger and faster simulations

2.1 Extremely large N-body simulations

Large galaxy surveys that are currently underway or are to commence in this decade (such as Pan-STARRS, BigBOSS, EUCLID, etc.) will drastically improve the statistical constraints on galaxy formation and evolution. One of their main goals is to relate the observed galaxy distribution to the underlying matter concentration in order to reconstruct the expansion history of the Universe, and thus to constrain dark energy. The required percent level accuracy in the interpretation of these missions can only be reached through a much better understanding of galaxy formation in direct cosmological simulations. However, in order to accurately represent the features of the baryonic acoustic simulations, extremely large simulation volumes of order $3\,{\rm Gpc}$ or more are required, pushing the resolution requirements to several hundred billion particles.

A first large N-body simulation able to meet these requirements has recently been presented in terms of the "Millennium XXL" (MXXL) calculation (Angulo et al. 2012). The simulation adopts a ΛCDM cosmology with the same cosmological parameters and output times as the previous two Millennium simulations (Boylan-Kolchin et al. 2009; Springel et al. 2005) in order to facilitate their joint use in building models for the galaxy population. The total particle number was 6720^3, slightly more than 303 billion, in a periodic simulation cube $3\,h^{-1}\,{\rm Gpc}$ on a side. This simulation has been carried out in the late summer 2010 on the JuRoPa machine at the Jülich Supercomputing Centre (JSC). A partition of 1536 computer nodes was used, each equipped with two quad-core Intel X5570 processors and 24 GB of RAM. The code was run in a hybrid MPI/shared memory setup on 12888 cores, placing one MPI task per processor (3072 in total), and employing all four cores of each socket via threads. A highly memory optimized version of the GADGET-3 code (last described in Springel 2005) was employed for the calculation, which needed a total of 2.86 million CPU hours, including the on-the-fly postprocessing done by the code, which primarily consisted of group and subhalo identification.

At the present day the MXXL contains more than 700 million nonlinear structures with at least 20 particles, binding $44\,\%$ of all the simulation particles. Among these objects, 23 million have a value of M_{200} larger than that of the Milky-Way halo ($M_{200} = 2{\times}10^{12}\,{\rm M}_{\odot}$) and 464 a value in excess of that of the Coma cluster ($M_{200} = 2{\times}10^{15}\,{\rm M}_{\odot}$). These statistics make the simulation ideal for precision studies of large-scale clustering and the identification of rare objects such as extremely massive clusters of galaxies.

In Fig. 1, power spectra for the mass density field at the present epoch are shown and compared to results for the MS and MS-II simulations. Clearly, only the MXXL simulation probes scales significantly beyond the maximum of the power spectrum. The MXXL is the only one of the three runs that provides good sampling of the baryonic acoustic

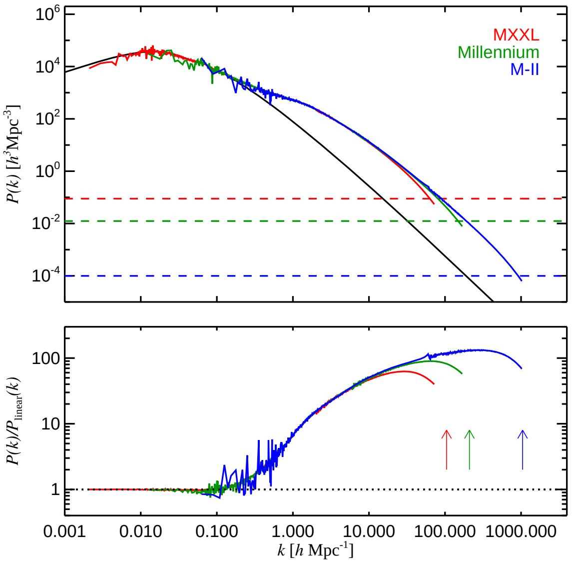

Fig. 1 Power spectrum at $z = 0$ in the Millennium, Millennium-II, and MXXL simulations (adopted from Angulo et al. 2012). The runs agree remarkable well over jointly resolved scales.

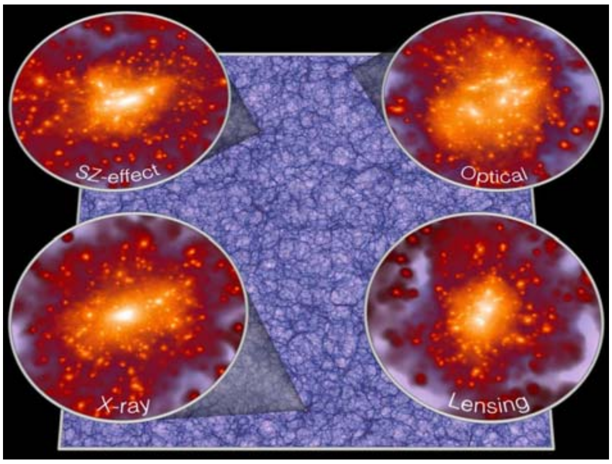

Fig. 2 Clusters of galaxies detected in the Millennium-XXL simulation through surrogate variables that model different physical probes (adopted from Angulo et al. 2012). The density slice shown in the background has an extension of $1500\,h^{-1}\,\mathrm{Mpc}$ and a thickness of $15\,h^{-1}\,\mathrm{Mpc}$. The four clusters shown are the ones with the largest signals in each of the observables. The slice has been cut such that it contains three of these objects.

oscillations (BAO) feature, which lies on scales that already show clear signs of nonlinear evolution. Throughout the nonlinear regime, the power spectra of the three simulations show excellent agreement with each other, until eventually the individual spatial resolution limits kick in, resulting in power spectra which fall below those of higher resolution counterparts on scales slightly larger than the nominal resolution limit, which is here identified with the gravitational softening length.

A visual impression of the large-scale structure in the simulation and some of its largest clusters is given in Fig. 2. The clusters highlighted here are the four largest clusters according to surrogate observables created to mimic cluster detections through observations of X-ray emission, gravitational lensing, Sunyaev-Zeldovich effect, or optical cluster counts. Interestingly, the strongest of these four signals identify different clusters, i.e. one would infer different "most massive" clusters in the simulation volume depending on what observational probe is used. This highlights the difficulties associated with using these observables to measure an absolute mass of galaxy clusters. It also provides a clear hint for the importance of selection effects when studying cluster scaling relations, an aspect that can be studied particularly well with this extremely large simulation. Indeed, Angulo et al. (2012) have recently been able to show that a combination of selection effects and systematic errors can explain the peculiar offset in the cluster scaling relations recently reported for an analysis of Planck early release data (Planck Collaboration 2011) between the mean thermal SZ signal measured for X-ray or optically selected clusters.

2.2 Extremely large hydrodynamic simulations

It has recently become clear that baryonic effects may modify the matter power spectrum on fairly large scales (Jing et al. 2006; Semboloni et al. 2011), sometimes reaching of order 30 % at wave numbers as small as $k \sim 10\,\mathrm{Mpc}^{-1}$. The magnitude of this effect has been found to depend on details of the baryonic feedback model, in particular on the inclusion of quasar-driven outflows (Semboloni et al. 2011). Effects on the power spectrum of this size must be accurately taken into account in order to achieve precise constraints on dark energy through missions such as EUCLID. So far, the impact of baryonic physics on the structure of dark matter halos and on large-scale clustering has often been neglected, or has only been coarsely estimated. Future hydrodynamic simulations face the task to calculate quantitatively more reliable predictions for these effects.

Another motivation for extremely large cosmological hydrodynamic simulations lies in the study of the emergence of the first galaxies and quasars, which are characterized by a high space density due to their rareness. An example for a recent work in this direction is the "MassiveBlack" simulation of Di Matteo et al. (2012). In this calculation, more than 65.5 billion resolution elements were used in a box of size $533\,h^{-1}\,\mathrm{Mpc}$, making it by far the largest cosmological SPH simulation with "full physics" of galaxy formation (meaning here an inclusion of radiative cooling, star formation and feedback physics) ever carried out. It also used an impressively large number of 10^5 cores, employing the entire Cray-XT5 Kraken supercomputer at the National Institute for Computational Sciences in the US.

Figure 3 shows a rendering of a sequence of zooms created with an interactive visualization technique developed for the simulation (Feng et al. 2011), based on terapixel im-

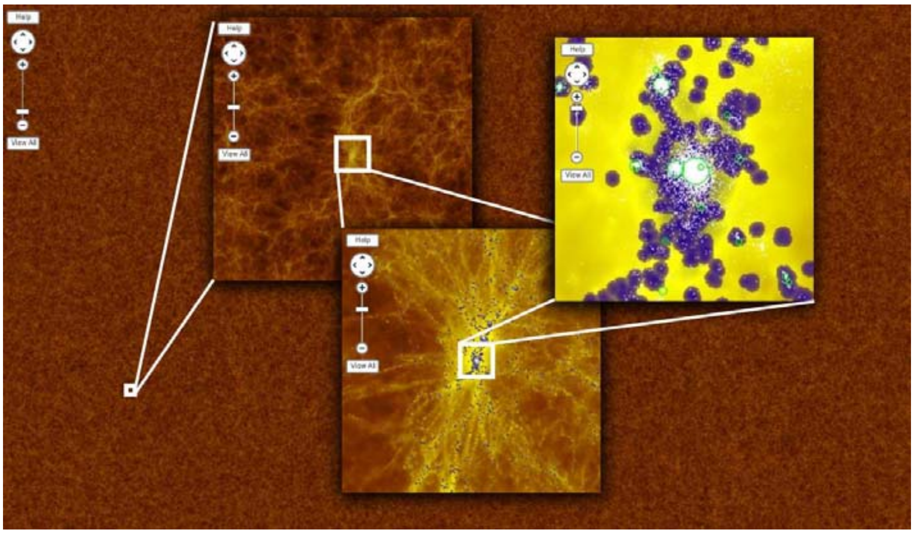

Fig. 3 (online colour at: www.an-journal.org) Visualization of the gas distribution in the MassiveBlack simulation (Di Matteo et al. 2012) at redshift $z = 4.75$ (adopted from Feng et al. 2011). The background image gives an overall view of the entire snapshot, while the insets show a zoom sequence onto a region around one of the most massive black holes (green circles) in the simulation.

ages that are streamed to a web-browser. As a first important science result for this simulation, Di Matteo et al. (2012) found that the early Sloan quasars appear to grow through cold gas flows that bring baryons quickly to the very central regions of large proto-galaxies, without first being held up in a galactic-scale disk that then becomes unstable and feeds the gas to the center through secular processes (Bournaud et al. 2011).

2.3 The quest for speed

In the past decades, computational cosmology has enjoyed a rapid expansion of what is possible thanks due to an exponential growth of the power of fastest supercomputers. While at present this trend is still in place, something quite fundamental has changed. The individual computing cores have stopped to become much faster with every generation. Instead, more and more of these cores are used in single CPUs, and more of them are combined in single compute nodes that are networked together via a fast communication fabric. In effect, the exponential growth of computing power is now coming from an exponential growth in the number of cores, creating an intense parallelization challenge for non-trivially coupled problems. The long-range nature of gravity makes cosmic structure growth a particularly difficult problem in this context. It should hence not be taken for granted that current petaflop supercomputers, yet alone the exaflop machines anticipated in 2018, can be used at scale for galaxy formation. Today's codes will certainly not be up for the task.

One strategy to extend the range of scalability of today's MPI parallelized codes is to use a mixed MPI, OpenMP, and Pthreads approach for parallelization. The resulting hybrid codes can take better advantage of shared memory nodes when available, helping to reduce communication times and scalability losses from work-load imbalances. In fact, hybrid approaches have already been successfully used in the currently largest cosmological simulations, in particular in the MXXL and MassiveBlack simulations.

Another possibility for improved performance lies in efforts to aggressively exploit vector instructions that have recently been introduced into the mainstream CPUs from Intel and AMD (and in similar fashion by IBM). In the form of SSE and AVX instructions (or AltiVec/QPX for IBM), these CPU features carry out parallel operations on up to 4 double- or 8 single-precision floating point numbers with just a single machine instruction.

A yet more radical approach for parallelization is to employ the streaming processors developed for Graphics Processing Units (GPUs). Here often more than 10^5 lightweight threads can be executed in parallel on these special devices, in principal offering a more cost efficient way to reach a certain performance goal, provided the problem at hand can be adjusted to make efficient use of such devices. Remarkably, even the relatively complicated Barnes & Hut tree algorithm has been successfully adapted to GPUs (e.g. Bédorf et al. 2012; Gaburov et al. 2010; Hamada et al. 2009), leading to substantial speed-ups compared to a solution running only on the host CPU. While the long-term technological development is uncertain, it appears clear that programming the supercomputers of the future will become considerably harder and likely involve heterogeneous architectures similar to the present CPU/GPU combinations.

3 Simulation accuracy

When only dark matter is considered, the newest generation of cosmological simulations yields a consensus view of the matter distribution in the Universe, even though the computational methods employed in these simulations vary

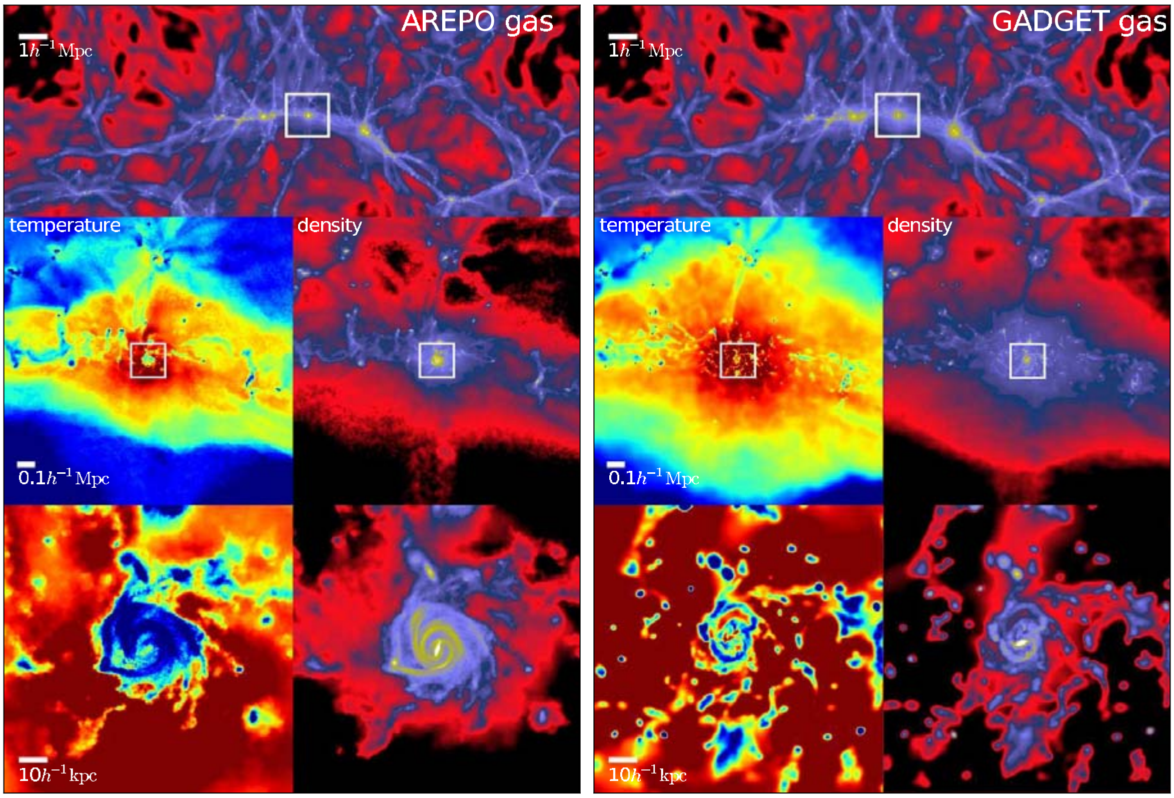

Fig. 4 Large-scale gas distribution in simulations of galaxy formation that compare the new AREPO code with the traditional SPH technique implemented in the GADGET code (adopted from Vogelsberger et al. 2011). In the large-scale distribution shown in the *top panels*, broad agreement is found. However, zooms into individual galaxies (*bottom panels*) reveal that the gaseous disks in the AREPO calculations are larger and more massive.

between the codes employed by different groups. The situation is rather different for cosmological hydrodynamic simulations, which also account for the baryons in the Universe and attempt to follow their evolution in concert with the dark matter. Because such more ambitious simulations track the luminous material directly, their results can be related to observations with a minimum of additional assumptions.

Over the last two decades, a number of hydrodynamic codes have been developed that are widely used in computational cosmology. Fundamentally, they all solve the equations of hydrodynamics coupled to gravity. However, the discretization schemes and numerical algorithms used to solve these equations differ strongly from code to code, and large differences are found in the simulation outcomes as well. Qualitative differences are present already at the simplest level of modeling where radiative gas cooling and energetic feedback processes owing to star formation are ignored (Frenk et al. 1999), emphasizing the need to more clearly understand which numerical methods give reliable results in the context of cosmic structure formation. In particular, the widely employed smoothed particle hydrodynamics (SPH) technique has come under severe scrutiny recently, as it has been shown to be inaccurate for certain fluid phenomena relevant for galaxy formation, such as Kelvin-Helmholtz instabilities (Agertz et al. 2007) or subsonic turbulence (Bauer & Springel 2011). On the other hand, adaptive mesh refinement (AMR) codes have also weaknesses that may impact their accuracy in cosmological applications, such as a dependence of advection errors on the bulk velocity, and problems to follow the growth of very small halos (Heitmann et al. 2008).

In this situation, Springel (2010) proposed a novel moving mesh scheme for hydrodynamics that is particularly well matched to the requirements of galaxy formation. The method uses a set of mesh-generating points to define a Voronoi tessellation on which a second-order accurate, finite-volume Godunov scheme is formulated. The mesh-generating points can in principle be moved arbitrarily. If they are moved with the flow, a pseudo-Lagrangian method results that retains the accuracy and high convergence rate of Eulerian hydrodynamics, as well as the adaptivity and geometric flexibility of SPH. In particular, the truncation error does not depend on the bulk motion of the system (making the results manifestly Galilean invariant), unlike in Eulerian mesh codes. Also, the reduced numerical diffusivity allows in principle a larger resolving power for a given number of resolution elements.

This new method, which has been implemented in the massively parallel AREPO code, appears ideally matched to the problem of galaxy formation because it features an automatic and continuous adaptivity to collapsing structures and it is free of preferred coordinate directions, just like SPH. In

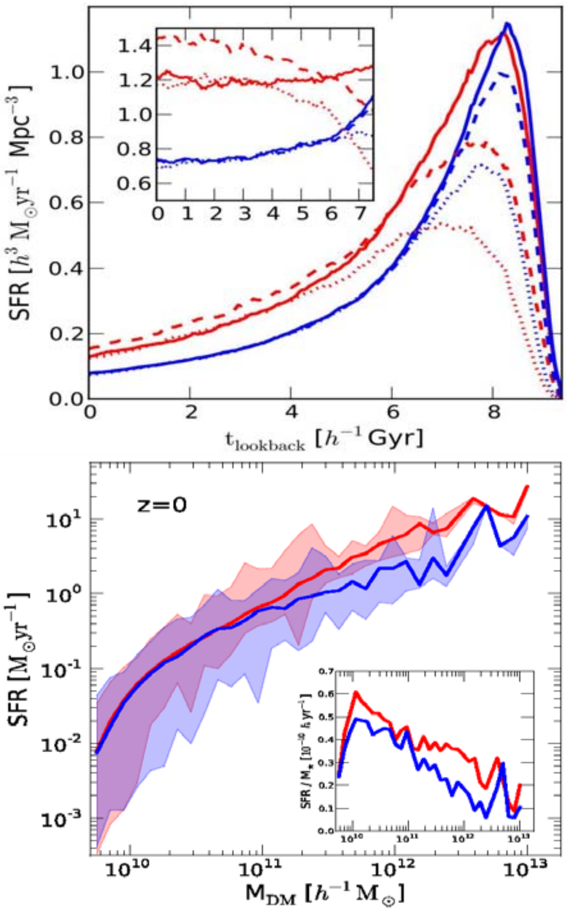

Fig. 5 Differences in the total star formation rate density (*top panel*) and in the mean specific star formation rate as a function of halo mass (*bottom panel*) in cosmological simulations carried out with the AREPO (shown in red) and GADGET (shown in blue) codes (adopted from Vogelsberger et al. 2011).

addition, the mesh can be adaptively refined and de-refined if desired, just like in Eulerian AMR codes, although this will only be rarely necessary if a roughly constant mass resolution is desired. At the same time, the AREPO method is free of the characteristic noise and gradient errors that plague SPH. Instead, it offers the high accuracy for fluid phenomena that can be achieved with Eulerian techniques. Last but not least, large supersonic translations of a system, which are common in cosmology, can be followed without loss of accuracy in AREPO.

In recent work (Sijacki et al. 2011), it has been shown that AREPO is considerably more accurate than SPH in all the examined problems, especially for multi-dimensional flows. Importantly, in a number of recent first applications of the code to the problem of galaxy formation (Keres et al. 2011; Vogelsberger et al. 2011), it was found that the new method leads to significant differences in cosmological simulations when compared with SPH. For an identical physics model and an equal gravity solver, larger and better defined disk galaxies are formed in AREPO compared with SPH. This is illustrated in Fig. 4, which compares the gas distribution in the two codes on large scales and in a zoom onto one of the galaxies formed.

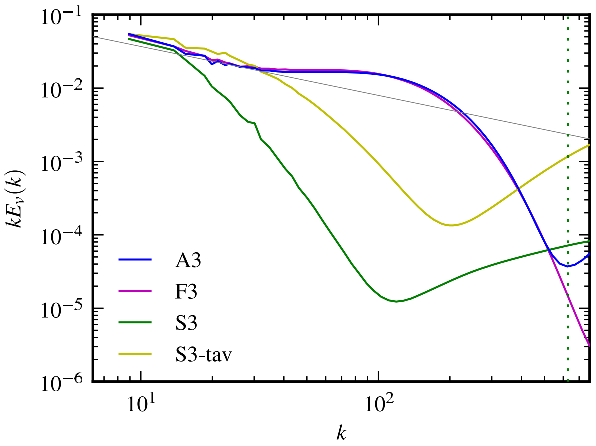

Fig. 6 Turbulent velocity power spectra in simulations of driven subsonic turbulence. Runs simulated at 256^3 resolution with AREPO (labeled A3), a fixed mesh (F3), and two variants of SPH, a "standard" one (S3) and one with a time-variable viscosity (S3-tav) are compared. The grey line indicates the power-law region of Kolmogorov turbulence (adopted from Bauer & Springel 2011).

We emphasize that this difference arises purely from the use of a more accurate treatment of hydrodynamics, stressing the need to constantly improve the numerical techniques used for galaxy formation. In particular, Vogelsberger et al. (2011) have shown that there is a different efficiency of cooling in large halos, manifesting itself in a change of the cosmic star formation history, as seen in Fig. 5. The cause of this numerical effect lies in the dissipation of subsonic turbulence and in viscous heating effects in the outer parts of halos. This is also supported by the work of Bauer & Springel (2011) who pointed out that AREPO produces the expected Kolmogorov turbulent cascade for subsonic turbulence whereas SPH gives problematic results in this regime (see Fig. 6), producing excessive dissipation on large scales due to numerical viscosity, and spurious dissipation on small scales due to the damping of small-scale velocity noise. The origin of the latter can be traced back to a lack of first-order consistency in SPH, causing large gradient errors (e.g. Abel 2011; Maron et al. 2011). Gradient errors are also one of the primary causes for the general problems of SPH in accurately following fluid instabilities, which were reported in recent years.

It remains to be seen whether improved versions of SPH can rectify these accuracy problems. In any case, the new moving-mesh approach is an interesting alternative which may well be the method of choice for the highly dynamic and adaptive regime of galaxy formation. The potential of AREPO has also become clear in first applications of the code to primordial star formation (Greif et al. 2011a,b), which traditionally has been a stronghold of AMR codes, albeit with notable SPH contributions as well.

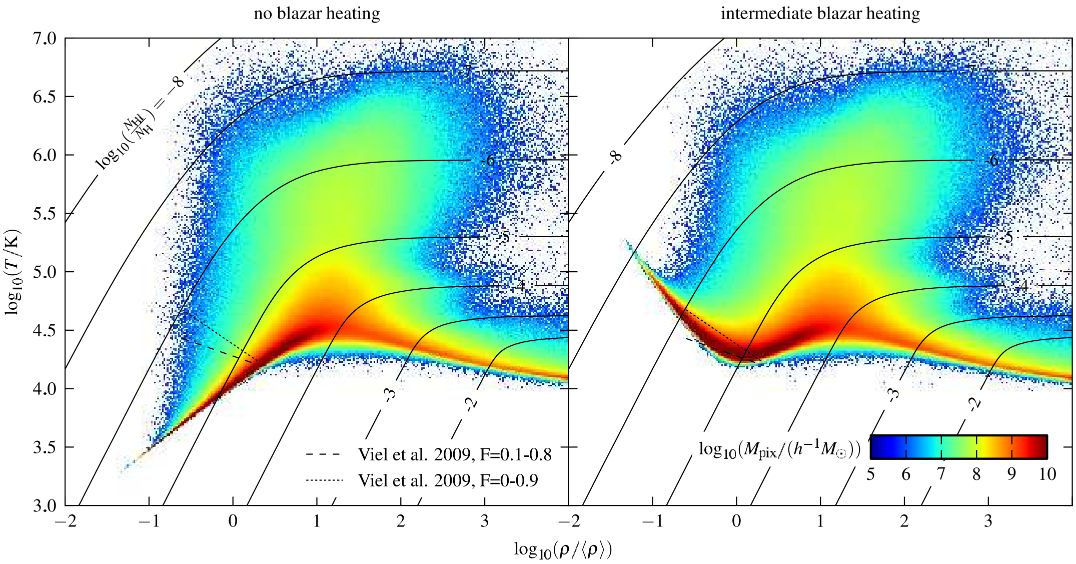

Fig. 7 Phase-space distribution of gas in the density-temperature plane in simulations without blazar heating (*left panel*) and with blazar heating (*right panel*) at redshift $z = 3$ (adopted from Puchwein et al. 2011). The contours indicate the neutral hydrogen fractions $\log(N_{\rm H\,I}/N_{\rm H})$ in the simulations.

4 Galaxy formation physics in simulations

4.1 Modeling star formation and feedback

Lack of sufficient numerical resolution has for a long time forced simulators to develop simple sub-resolution parameterizations for star formation and its regulation by energy injection (e.g. Springel & Hernquist 2003). The star formation rate itself has often been parameterized such that the observational Kennicutt law is reproduced, which empirically relates the total surface density of gas in disk galaxies to their star formation rate surface density. The feedback processes included in the models were typically restricted to supernovae feedback, and were often treated through “delayed cooling” or other implicit or explicit sub-resolution models. However, advances in numerical techniques and, more importantly, our understanding of the physics of star formation make this approach outdated by now.

Observationally, we know that the star formation law is really somewhat different, and appears to closely follow the molecular hydrogen mass and not the total gas density (e.g. Bigiel et al. 2008). Also, the efficiency of star formation in molecular gas seems to be constant and independent of density, at around $\sim$1 % of the local free fall time. This progress needs to be reflected in galaxy formation models through a more direct accounting for the multiphase structure of the ISM. Ideally, this should be done in global galaxy-wide simulations, as only such models can properly address the balance between all phases in the ISM, the relevance of gravitationally induced turbulence driving, the continuous formation and destruction of GMCs over a range of sizes, and the origin of Kennicutt like scaling relations. Recently, some promising first steps in this direction have been made (Gnedin et al. 2009; Hopkins et al. 2011; Robertson & Kravtsov 2008), but it remains a significant challenge to make such models applicable in cosmological simulations and successfully predict the properties of whole galaxy populations.

4.2 Including more “exotic” physics

Despite the recent progress in improving the modelling of star formation and feedback, it appears likely that none of the existing hydrodynamical simulation models accounts yet for all of the most relevant physics for galaxy formation. For example, cosmic ray particles are known to contribute substantially to the pressure of the ISM on our own Galaxy, yet they are often neglected in simulation works. Interestingly, in the small number of studies where they have been included, they were found to have quite strong effects, particularly in small dwarf galaxies and satellite systems (Jubelgas et al. 2008; Uhlig et al. 2012; Wadepuhl & Springel 2011). A similar statement could be made about magnetic fields. They are known to be an important factor in the ISM, but it is presently quite unclear whether they also regulate galaxy formation on larger scales.

Occasionally, also entirely new physical processes affecting galaxy formation are identified. For example, Chang et al. (2011) recently proposed that TeV blazar heating, as a result of plasma instabilities in the relativistic positron-electron beams created by interactions of TeV blazar photons with the background photon field, creates a volumetric heating of the intergalactic gas. Interestingly, this physical heating process produces an inverted temperature-density relation for the low density gas, something that had previously been proposed (Bolton et al. 2008) as a phenomenological model to resolve some nagging discrepancies in observations of the Lyα forest, which is otherwise in very good agreement with predictions by the CDM theory.

Indeed, using hydrodynamic simulations that accounted for the first time for this new physics of TeV blazar heating, Puchwein et al. (2011) have shown that the Lyα predictions resulting from this heating process are for the first time in essentially perfect agreement with data on the flux power spectrum, the redshift evolution of the mean transmission, and the flux one-point probability distribution function. As seen in Fig. 7, the heating produces an inverted equation of state, and it probably has also a profound influence on galaxies at the dwarf scale, something that needs to be further investigated in future work.

5 Conclusions

Cosmological simulations have become a primary driver for further advancing galaxy formation theory. They help us to validate the underlying cosmological paradigm and allow us to understand the connection between the physics of galaxy formation and the large-scale distribution of gas in the Universe. Furthermore, they provide crucial guidance for the interpretation of forthcoming observational data.

In the next years we can look forward to see ground-breaking hydrodynamical simulations of galaxy formation with an unprecedented combination of resolution, physical complexity and volume. This optimistic outlook is based in part also on the continued rapid increase in the capabilities of state-of-the-art supercomputers. However, one should be keenly aware of the fact that it is becoming rapidly more difficult to fully exploit high-performance computers at the leading edge. Only when the necessary efforts for developing ever more scalable and accurate simulation codes are made, this vision of a golden age in computational cosmology will come true.

Acknowledgements. VS acknowledges financial support from Research Center SFB 881, “The Milky Way System”, of the German Science Foundation.

References

Abel, T.: 2011, MNRAS 413, 271
Agertz, O., Moore, B., Stadel, J., et al.: 2007, MNRAS 380, 963
Angulo, R.E., Springel, V., White, S.D.M., Jenkins, A., Baugh, C.M., Frenk, C.S.: 2012, astro-ph/1203.3216
Bauer, A., Springel, V.: 2011, astro-ph/1109.4413
Bédorf, J., Gaburov, E., Portegies Zwart, S.: 2012, astro-ph/1204.2280
Bigiel, F., Leroy, A., Walter, F., et al.: 2008, AJ 136, 2846
Bolton, J.S., Viel, M., Kim, T.-S., Haehnelt, M.G., Carswell, R.F.: 2008, MNRAS 386, 1131
Bournaud, F., Dekel, A., Teyssier, R., et al.: 2011, ApJ 741, L33
Boylan-Kolchin, M., Springel, V., White, S.D.M., Jenkins, A., Lemson, G.: 2009, MNRAS 398, 1150
Chang, P., Broderick, A.E., Pfrommer, C.: 2011, astro-ph/1106.5504
Croton, D.J., Springel, V., White, S.D.M., et al.: 2006, MNRAS 365, 11
Diemand, J., Kuhlen, M., Madau, P.: 2007, ApJ 657, 262
Di Matteo, T., Khandai, N., DeGraf, C., et al.: 2012, ApJ 745, L29
Dolag, K., Borgani, S., Murante, G., Springel, V.: 2009, MNRAS 399, 497
Feng, Y., Croft, R.A.C., Di Matteo, T., et al.: 2011, ApJS 197, 18
Frenk, C.S., White, S.D.M., Bode, P., et al.: 1999, ApJ 525, 554
Gaburov, E., Bédorf, J., Portegies Zwart, S.: 2010, Procedia Computer Science 1, 1119
Gnedin, N.Y., Tassis, K., Kravtsov, A.V.: 2009, ApJ 697, 55
Governato, F., Willman, B., Mayer, L., et al.: 2007, MNRAS 374, 1479
Greif, T.H., Springel, V., White, S.D.M., et al.: 2011a, ApJ 737, 75
Greif, T.H., White, S.D.M., Klessen, R.S., Springel V.: 2011b, ApJ 736, 147
Guedes, J., Callegari, S., Madau, P., Mayer, L.: 2011, ApJ 742, 76
Guo, Q., White, S., Li, C., Boylan-Kolchin, M.: 2010, MNRAS 404, 1111
Guo, Q., White, S., Boylan-Kolchin, M., et al.: 2011, MNRAS 413, 101
Hamada, T., Nitadori, K., Benkrid, K., et al.: 2009, Computer Science – Research and Development 24, 21
Heitmann, K., Lukić, Z., Fasel, P., et al.: 2008, Computational Science and Discovery 1, 015003
Hopkins, P.F., Quataert, E., Murray, N.: 2011, MNRAS 417, 950
Jing, Y.P., Zhang, P., Lin, W.P., Gao, L., Springel, V.: 2006, ApJ 640, L119
Jubelgas, M., Springel, V., Enßlin, T.A., Pfrommer, C.: 2008, A&A 481, 33
Keres, D., Vogelsberger, M., Sijacki, D., Springel, V., Hernquist, L.: 2011, astro-ph/1109.4638
Klypin, A.A., Trujillo-Gomez, S., Primack, J.: 2011, ApJ 740, 102
Maron, J.L., McNally, C.P., Mac Low, M.-M.: 2011, astro-ph/1110.0835
Planck Collaboration: 2011, A&A 536, A12
Puchwein, E., Pfrommer, C., Springel, V., Broderick, A.E., Chang P.: 2011, astro-ph/1107.3837
Robertson, B.E., Kravtsov, A.V.: 2008, ApJ 680, 1083
Sawala, T., Guo, Q., Scannapieco, C., Jenkins, A., White, S.: 2011, MNRAS 413, 659
Scannapieco, C., White, S.D.M., Springel, V., Tissera, P.B.: 2009, MNRAS 396, 696
Scannapieco, C., Wadepuhl, M., Parry, O.H., et al.: 2011, astro-ph/1112.0315
Semboloni, E., Hoekstra, H., Schaye, J., van Daalen, M.P., McCarthy, I.G.: 2011, MNRAS 417, 2020
Sijacki, D., Vogelsberger, M., Keres, D., Springel, V., Hernquist, L.: 2011, astro-ph/1109.3468
Springel, V.: 2005, MNRAS 364, 1105
Springel, V.: 2010, MNRAS 401, 791
Springel, V., Hernquist, L.: 2003, MNRAS 339, 289
Springel, V., White, S.D.M., Jenkins, A., et al.: 2005, Nature 435, 629
Springel, V., Wang, J., Vogelsberger, M., et al.: 2008, MNRAS 391, 1685
Stadel, J., Potter, D., Moore, B., et al.: 2009, MNRAS 398, L21
Stinson, G.S., Dalcanton, J.J., Quinn, T., Gogarten, S.M., Kaufmann, T., Wadsley, J.: 2009, MNRAS 395, 1455
Uhlig, M., Pfrommer, C., Sharma, M., Nath, B.B., Enßlin, T.A., Springel, V.: 2012, astro-ph/1203.1038
Vogelsberger, M., Sijacki, D., Keres, D., Springel, V., Hernquist, L.: 2011, astro-ph/1109.1281
Wadepuhl, M., Springel, V.: 2011, MNRAS 410, 1975