

# Towards exascale simulations of the magnetic universe

Volker Springel, Christian Klingenberg, Rüdiger Pakmor, Thomas Guillet, and Praveen Chandrashekar

**Abstract** Simulations of cosmic structure formation address multi-scale, multi-physics problems of vast proportions. These calculations are presently at the forefront of today's use of supercomputers, and are important scientific drivers for the future use of exaflop computing platforms. However, continued success in this field requires the development of new numerical methods that excel in accuracy, robustness, parallel scalability, and physical fidelity to the processes relevant in galaxy and star formation. In the EXAMAG project, we have worked on improving and applying the astrophysical moving-mesh code AREPO with the goal to extend its range of applicability. We have also worked on developing new, powerful high-order discontinuous Galerkin schemes for astrophysics, on more efficient solvers for gravity, and on improvements of the accuracy of the treatment of ideal magnetohydrodynamics. In this context, we have also studied the applied mathematics required for higher-order discretization on dynamically moving meshes, thereby providing the foundations for much more efficient and accurate methods than are presently in use. Finally, we have worked towards publicly releasing two major community codes, AREPO and GADGET-4, which represent the state-of-the-art in the field.

---

Volker Springel • Rüdiger Pakmor

Max-Planck-Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany,  
e-mail: vspringel@mpa-garching.mpg.de; rpakmor@mpa-garching.mpg.de

Christian Klingenberg

Institut für Mathematik, Universität Würzburg, Emil-Fischer-Str. 40, 97074 Würzburg, Germany,  
e-mail: klingen@mathematik.uni-wuerzburg.de

Thomas Guillet

College of Engineering, Mathematics and Physical Sciences, Physics Building, Streatham Campus,  
University of Exeter, North Park Road, Exeter EX4 4QL, UK, e-mail: T.A.Guillet@exeter.ac.uk

Praveen Chandrashekar

TIFR Centre for Applicable Mathematics, Bangalore-560065, India, e-mail: praveen@tifrbng.res.in

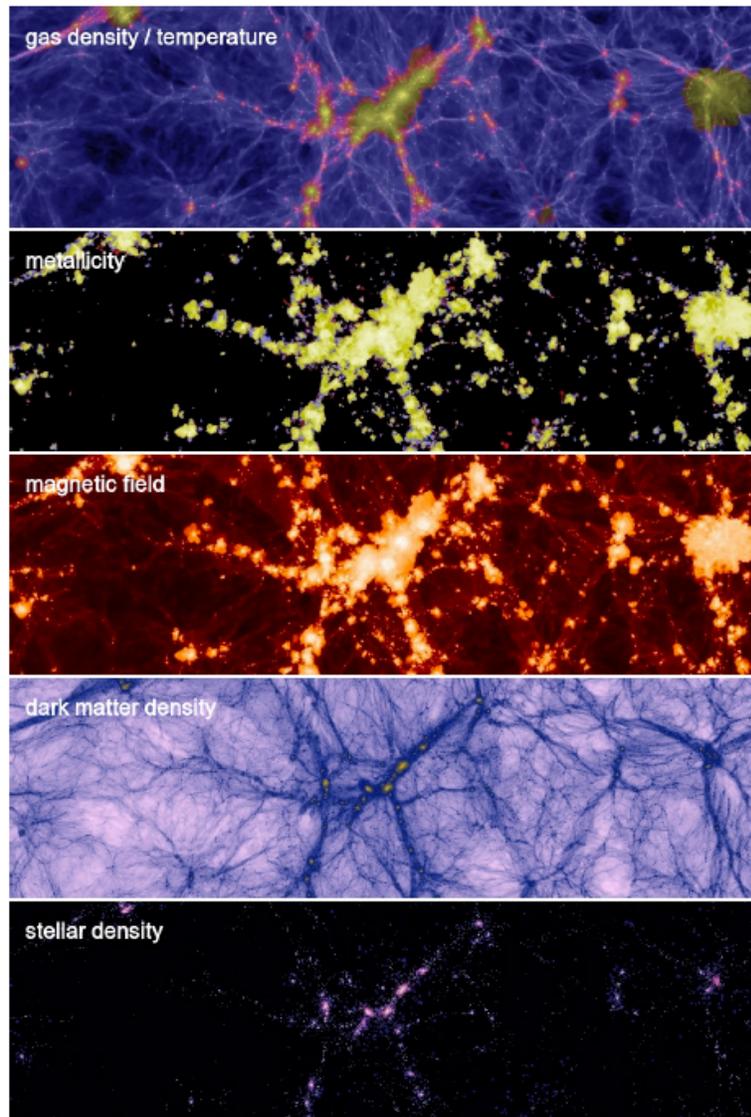
## 1 Introduction

Hydrodynamical simulations of galaxy formation have significantly matured over recent years and now enable successful predictions of the build-up of the galaxy population starting from cosmological initial conditions left behind by the Big Bang. These simulations can track the non-linearly coupled evolution of both baryons and dark matter, in principle fully accounting for their mutual influence on each other and yielding rich predictions for galaxy properties, the diffuse gas in the circumgalactic and intergalactic media, and cosmic dark matter clustering. Provided that such hydrodynamic simulations can be pushed to sufficiently large volumes, they provide the most powerful approach for forecasting non-linear cosmological observables related to clustering in different regimes and at different epochs. However, it is extremely challenging to include all the relevant physics, and to make the simulations accurate, fast, and scalable enough to be able to exploit the full capacity of today's supercomputers. The EXAMAG project has been aiming to advance novel numerical methodologies for astrophysical simulations, and to apply them right away to timely research questions in astrophysics. A particular focus has been on magnetic field predictions, and the development of higher-order methods.

In this project review, we report a subset of the results obtained within the EXAMAG project. In Section 2 we describe the IllustrisTNG simulations, the currently most advanced set of magnetohydrodynamical simulations of galaxy formation, as well as the Auriga simulations, which focus on predictions for our own Milky Way galaxy, in particular on the structure and origin of its magnetic field. In Section 3, we turn to our recent developments of discontinuous Galerkin hydro- and magnetohydrodynamics. We also give a short description of some of our methodological advances in combining the idea of fully dynamic, unstructured meshes with high-order discontinuous Galerkin approaches to hydrodynamics. We also report on the application of these higher order methodologies to the problem of driven isothermal turbulence, where these methods prove to be particularly powerful. We then recount in Section 4 some of our work on the performance and accuracy of the two cosmological hydrodynamical simulation codes GADGET-4 and AREPO, both of which we have prepared for public release to the community as part of this project. Finally, we summarize and give an outlook in Section 5.

## 2 The IllustrisTNG and Auriga simulations

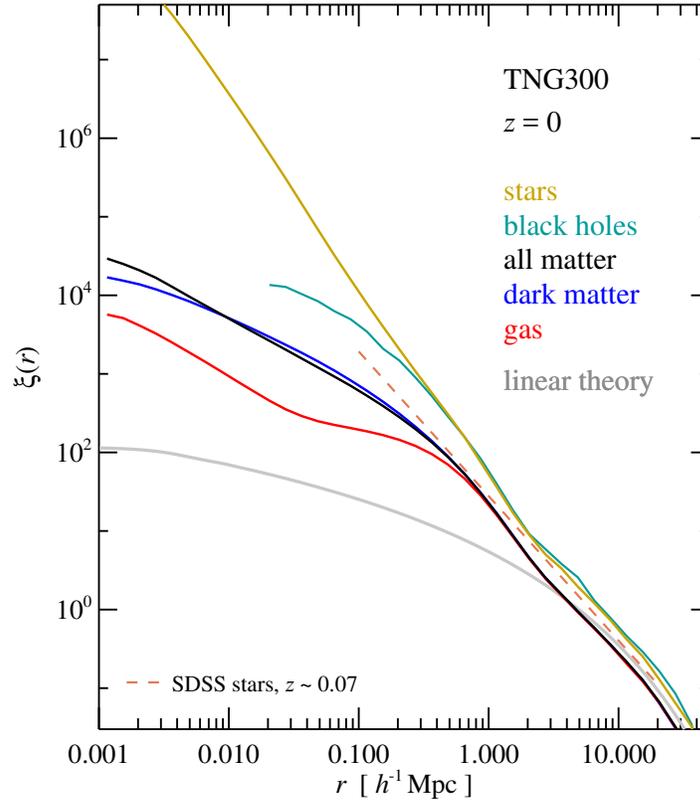
The AREPO code (Springel, 2010) introduced a different approach from the ones so far commonly adopted in astrophysics to evolve gas on a computer (smoothed particle hydrodynamics, SPH, and Eulerian mesh-based methods, typically utilizing adaptive mesh refinement, AMR). It employs a moving, unstructured mesh where, like AMR, the volume of space is discretized into many individual cells, but similar to SPH, these cells move with time, adapting to the flow of gas in their vicinity. As a result, the mesh itself, constructed through a Voronoi tessellation of space, has no



**Fig. 1** Thin projections through the TNG100 simulation, showing (from top to bottom) the gas density field, the metallicity, the magnetic field strength, the dark matter density, and the stellar density.

preferred directions or regular grid-like structure, and is highly spatially adaptive, making it ideal, in particular, for studying galaxy formation.

Our ground-breaking “Illustris” hydrodynamical calculation of galaxy formation (Vogelsberger et al., 2014) demonstrated the utility of the approach for simulations

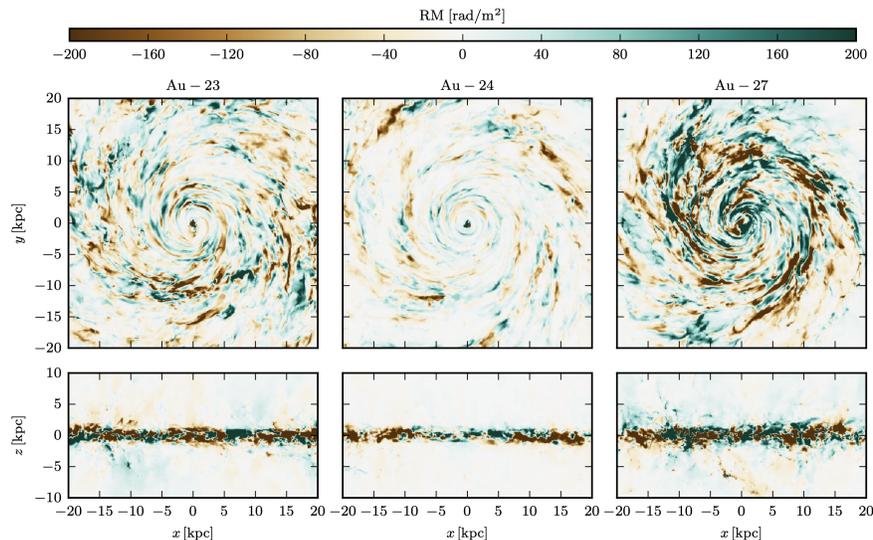


**Fig. 2** The matter autocorrelation function for different mass components in our high-resolution TNG300 simulation at redshift  $z = 0$ , see Springel et al. (2018). We show results for stellar matter, gas, dark matter, black holes, and all the matter, as labelled. The linear theory correlation function is shown in grey for comparison.

of structure formation. For the first time ever, it reproduced the observed morphological mix of galaxies and its dependence on stellar mass. Over the past years we have undertaken significant efforts to improve the underlying physics models (especially with respect to magnetic fields, and the processes regulating star formation through energetic feedback from supernovae and black holes), and the accuracy and scalability of the numerical algorithms (for example by developing a hierarchical local time-stepping algorithm for gravity). These efforts culminated in the simulation projects IllustrisTNG and Auriga.

The Next Generation (TNG) Illustris project<sup>1</sup> built on the technical and scientific achievements of its predecessor and pushed this line of research further. In particu-

<sup>1</sup> <http://www.tng-project.org>



**Fig. 3** Faraday rotation maps as seen by an external observer for three of the Auriga simulations. The top and bottom panels show face-on and edge-on Faraday rotation maps due to the magnetic fields in three different simulated disk galaxies Au-23, Au-24, and Au-27 (Pakmor et al., 2018). These predictions compare well to observations of real galaxies such as M51.

lar, it has improved upon Illustris by including our newly developed accurate solver for ideal magnetohydrodynamics (Pakmor et al., 2014, 2016b), and by extending the dynamic range and resolution of the simulated galaxies and haloes significantly through an ambitious suite of simulations carried out on the Hazel-Hen supercomputer at the High-Performance Computing Center Stuttgart (HLRS) with the help of two large compute-time grants by the Gauss Centre for Supercomputing (GCS). We have considered three different box sizes, roughly 300, 100, and 50 Mpc on a side, and computed extensive resolution studies for each of them, yielding three series or runs, entitled TNG300, TNG100, and TNG50. The calculations used up to 24000 cores, required 100 TB RAM, and produced 660 TB of science data. A visual impression of one of these simulation is given in Figure 1. As an illustrative result, we show the clustering statistics of different matter components at the present epoch in TNG300 in Figure 2 (Springel et al., 2018). The calculation is able to probe deeply into the non-linear regime, over a very large dynamic range, thereby allowing predictions both for the internal structure of galaxies as well as their large-scale clustering patterns.

The scientific analysis of the IllustrisTNG simulations has started in 2018 and has already led to many important results for a vast range of scientific questions (e.g. Springel et al., 2018; Nelson et al., 2018a; Pillepich et al., 2018; Torrey et al., 2018; Marinacci et al., 2018), demonstrating the great utility of these methods. At the time of this writing, the TNG simulations have already produced 67 journal

publications, and about 80 additional papers are currently in preparation by a network of international collaborators. In December 2018, we publicly released the data of TNG100 and TNG300 (Nelson et al., 2018b), with TNG50 to follow in a year’s time, which will further amplify the scientific use of the simulations.

The Auriga simulations (Grand et al., 2017) use a very similar physics and numerical model as IllustrisTNG, but “zoom-in” on individual Milky Way-sized galaxies that are studied with much higher resolution. A particular focus here has been on understanding the origin of the magnetic fields in galaxies, and on predicting their present structure as well as their built-up over cosmic time. In Pakmor et al. (2017) we have shown that the magnetic fields grow exponentially at early times owing to a small-scale dynamo with an e-folding time of roughly 100 Myr in the centre of haloes until saturation occurs around redshift  $z = 2 - 3$ , when the magnetic energy density reaches about 10 per cent of the turbulent energy density with a typical strength of  $10 - 50 \mu\text{G}$ . Outside the galactic centres, differential rotation in the discs leads to linear amplification of the magnetic fields that typically saturates around  $z = 0.5 - 0$ . The final radial and vertical variations of the magnetic field strength can be well described by two joint exponential profiles, and are in good agreement with observational constraints.

We have extended the observational comparisons by computing synthetic Faraday rotation maps due to the magnetic fields (Pakmor et al., 2018), for different observer positions within and outside the simulated galaxies. We find that the strength of the Faraday rotation of our simulated galaxies for a hypothetical observer at the solar circle is broadly consistent with the Faraday rotation seen for the Milky Way. The same holds for an observer outside the galaxy and the observed signal of the nearby spiral galaxy M51, see Figure 3. However, we also find that the structure of the synthetic all-sky Faraday rotation maps vary strongly with azimuthal position along the solar circle. This represents a severe obstacle for attempts to reconstruct the global magnetic field of the Milky Way from Faraday rotation maps alone without including additional observables.

### 3 Discontinuous Galerkin hydrodynamics for astrophysical applications

The AREPO code is based on a second order finite volume method on a moving mesh. An important part of the EXAMAG project was to develop a higher order method for AREPO, for which we have chosen to use a discontinuous Galerkin (DG) method due to the significant promise this class of methods holds for high performance computing. Here the hydrodynamic or magnetohydrodynamic partial differential equations are written in a weak form to construct a finite element method. On each cell, the solution and the test functions are approximated by appropriate polynomials. Note that in neighbouring cells this approximation need not be continuous across cell boundaries. The jumps between neighboring cells are taken into account by a numerical flux function based on an approximate Riemann solver which provides

stability to the method. The time discretisation is performed using a Runge-Kutta method. For more details, see Shu (2014).

Using the discontinuous Galerkin (DG) method brings with it the following advantages:

- DG works naturally on unstructured meshes which makes it suitable for adaptive meshes with hanging nodes and for a moving mesh methods.
- DG can be made to work for any order of accuracy with a compact stencil and provides spectral-type accuracy.
- DG is extremely local in data communication, thus ideally suited for efficient parallelization on current computer hardware.

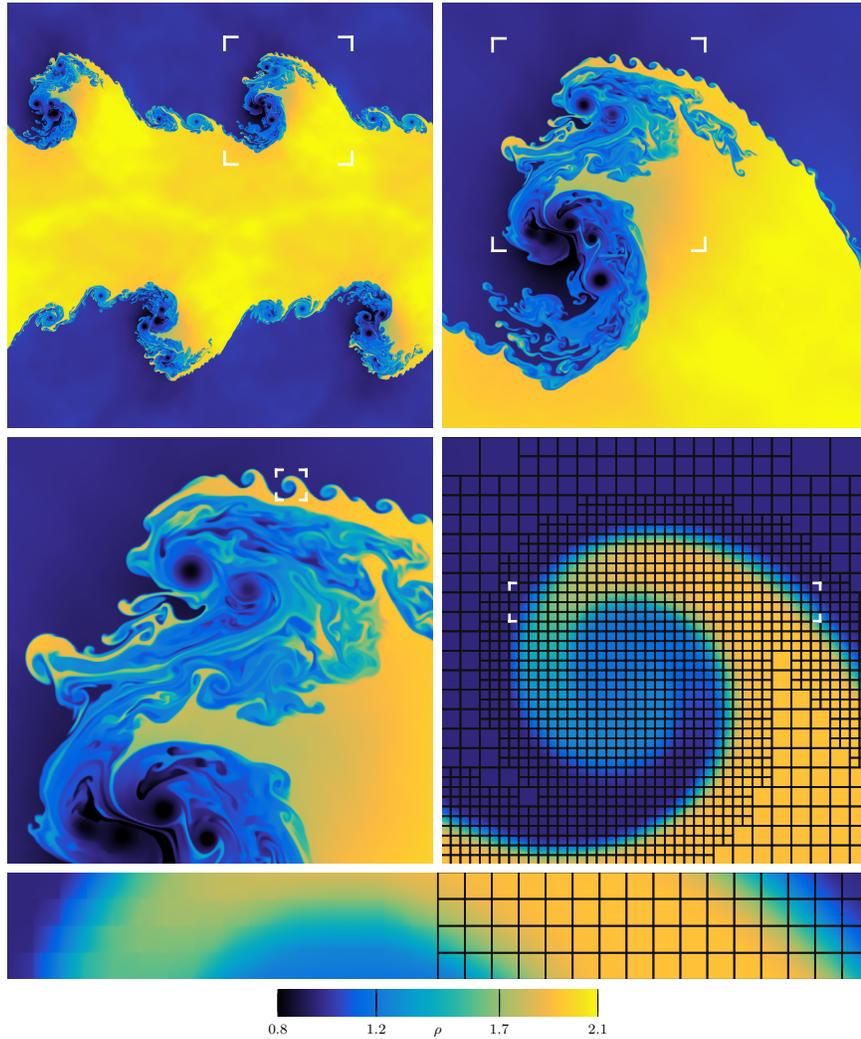
Due to these advantages and also the ability of DG to compute convection dominated problems in a stable and accurate manner, the EXAMAG project worked towards adopting DG methods for astrophysical applications. Our efforts in this direction followed two main paths that will be presented in the following subsections. First, in Section 3.1, we describe the discontinuous Galerkin method on a Cartesian mesh with automatic mesh refinement for both the Euler and magnetohydrodynamics equations. Then, in Section 3.2, the discontinuous Galerkin method on a moving mesh is discussed. Finally, in Section 3.3 we present some results on turbulence simulations with higher order numerics.

### **3.1 The discontinuous Galerkin method on a Cartesian mesh with automatic mesh refinement**

In a collaboration with the EXA-DUNE project (Peter Bastian, Heidelberg) we implemented a two-dimensional hydrodynamics code in the DUNE framework with total variation bounded and positivity preserving limiters. A mesh refinement triggered by the limiting criteria was also built in, see Gallego et al. (2014).

In Gallego et al. (2016), we developed a code based on the DG method for compressible flows to incorporate and test shock indicators that can determine which cells need limiting. We showed that the choice of variables that are limited can have a major influence on accuracy; limiting the characteristic variables was compared to limiting the conserved variables, with the former being better able to control oscillations. These limiters were then combined with a shock indicator, yielding the ability to solve complex flow problems in a more efficient and accurate manner since the costly limiters need not be applied everywhere.

These investigations provided the foundations for our implementation of the discontinuous Galerkin approach in a new branch of the AREPO code called TENET, as presented in Schaal et al. (2015). This version of the code supports adaptive mesh refinement (AMR) and is able to maintain high-order accuracy at AMR refinement boundaries, unlike finite-volume approaches that typically fall-back to first order there. As an illustrative example, Figure 4 depicts a simulation of a Kelvin-Helmholtz instability using TENET. In Schaal et al. (2015) we have also shown that DG has



**Fig. 4** A Kelvin-Helmholtz simulation with fourth order DG and adaptive mesh refinement (see Schaal et al., 2015). The simulation starts with  $64^2$  cells, and refines in selected regions to an effective resolution of  $4096^2$ . As can be seen in the bottom panel (white markers indicate regions that are enlarged in subsequent panels), the solution within every cell contains rich information, consisting of a third order polynomial.

significant advantages over a finite volume method. Given the same accuracy target, a higher order DG method requires fewer grid points than a finite volume method, allowing for much faster run times, especially for smooth solutions. The higher order DG method of TENET is also better in computing discontinuous solutions, although the numerical techniques for identifying regions where a limiter has to be

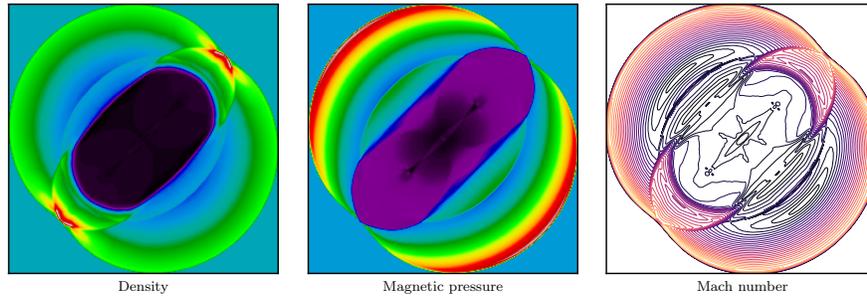
applied are intricate and still not fully mature. We also showed that the DG approach automatically conserves angular momentum in smooth regions which is beneficial for many astrophysical problems involving rotating objects.

We next worked towards adding magnetic fields, leading to the model of ideal magnetohydrodynamics in fully compressible flows. The system of ideal magnetohydrodynamical equations poses additional challenges for numerical simulations, mainly due to the need to preserve the divergence-free condition on the magnetic field, which requires specialized techniques. In Chandrashekar and Klingenberg (2016) we developed an entropy stable finite volume scheme based on a symmetrized version of the MHD equations. A numerical flux is given which allows for the construction of an entropy conservative and entropy stable scheme. It is demonstrated how this new scheme is robust for MHD simulations due to its entropy framework, in spite of the divergence condition not being explicitly satisfied. In Chandrashekar et al. (2018) this approach was extended to an explicit high order Runge-Kutta discontinuous Galerkin method. This methodology is then combined with techniques used to control oscillations near discontinuities, similar to Gallego et al. (2016), where these techniques were introduced for the hydrodynamical case. We assessed a different approach for the divergence constraint in Klingenberg et al. (2017a), where post- and pre-processing methods are suggested in order to numerically maintain the divergence free constraint.

Finally, in Guillet et al. (2019) we implemented high-order MHD in the AREPO code using two different approaches for maintaining the divergence constraint, a locally divergence-free basis combined with Powell terms for stability, and a hyperbolic divergence cleaning method. Two new numerical ingredients were introduced in the DG scheme: a non-linear limiting procedure for the magnetic field, and a different discretization of the Powell terms, which was found to be a key aspect for stability and accuracy of the method. The beneficial properties of the DG method found for hydrodynamical simulations were also confirmed by Guillet et al. (2019) for the MHD case. The resulting scheme shows lower advection errors and better Galilean invariance than a finite volume scheme, and hence constitutes a very promising approach for more realistic applications in an astrophysical context. In Figure 5, we show as an example the DG simulation of a two-dimensional MHD blast wave, which is a particularly challenging test case in terms of maintaining positive solutions. In Figure 6, we demonstrate the convergence of our DG algorithms when applied to a smooth isodensity MHD vortex problem that is advected at different angles through a periodic domain. Our numerical solutions reproduce to good accuracy the expected convergence orders for orders 2 to 6, both for a locally divergence-free basis and a Legendre basis.

### 3.2 The discontinuous Galerkin method on a moving mesh

In a parallel endeavour, we have investigated in EXAMAG the general problem of extending higher-order methods to unstructured meshes that move along with the



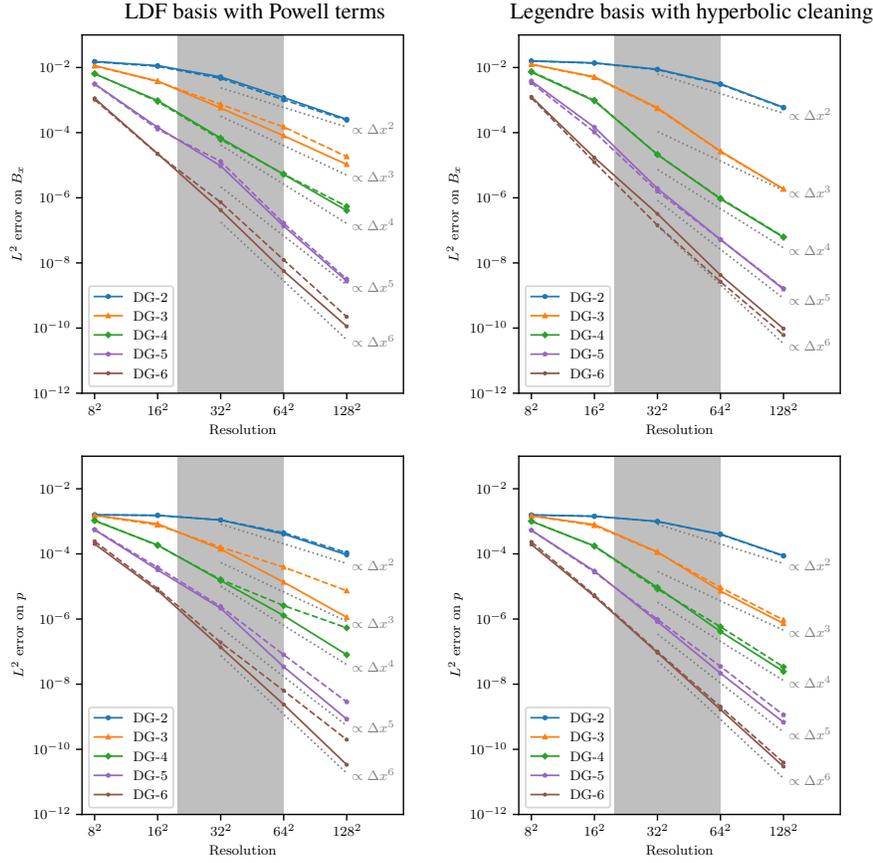
**Fig. 5** A two-dimensional MHD blast wave test problem. The density, magnetic pressure and Mach number contours are shown on a  $256^2$  grid using a third-order discontinuous Galerkin scheme where the MHD equations are written in symmetrized form using so-called Powell terms, see Guillet et al. (2019).

flow. The original AREPO approach employed a second order finite volume scheme for this purpose. As this works quite well and is able to improve the resolution of flow features by minimizing numerical dissipation from advection, we pursued extensions of this approach in three different directions.

The DG method can be proven to be entropy stable and convergent on a fixed grid for a scalar conservation law. In a first line of investigation we have shown that the entropy stability is maintained on a moving mesh. This was demonstrated in Klingenberg et al. (2017b) and Klingenberg et al. (2017c) for a semidiscrete arbitrary Lagrangian-Eulerian discontinuous Galerkin method. In Klingenberg et al. (2018) these ideas were extended to a fully discrete method. Numerical experiments have confirmed these properties also for a multi-dimensional implementation of the hydrodynamical equations.

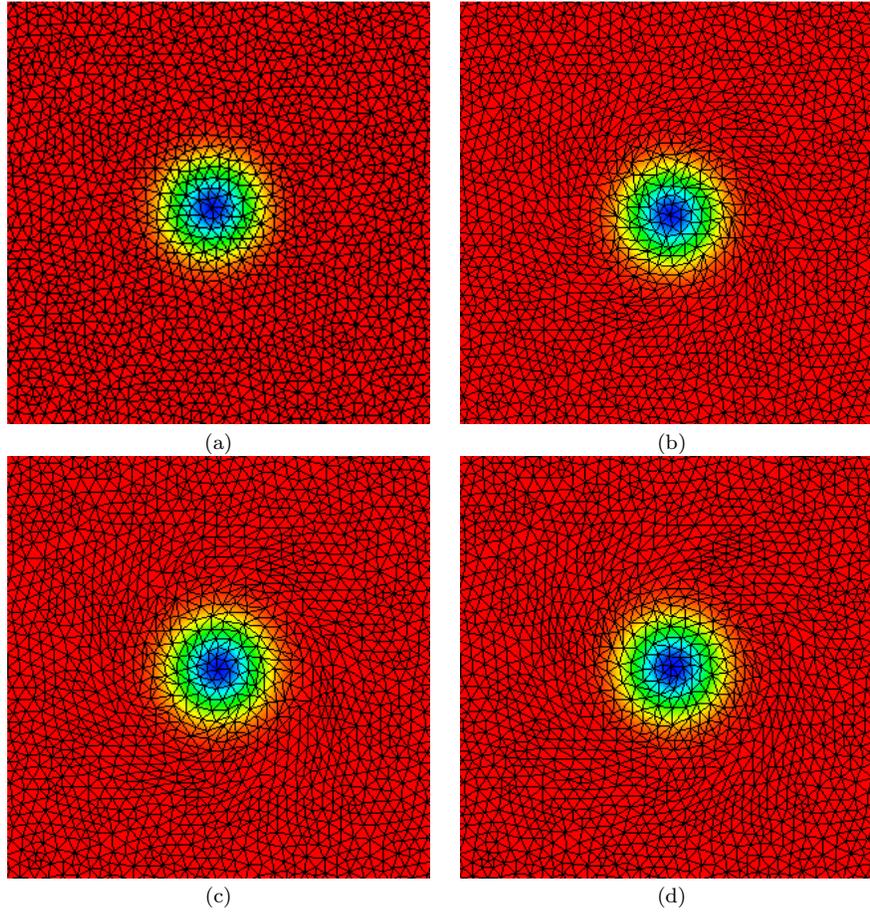
The above studies were restricted to arbitrary Lagrangian-Eulerian discontinuous Galerkin methods, where the grid needs no remeshing. In practice, one needs to remesh the grid once in a while since otherwise the mesh quality degrades to such an extent that the computations can break down. In Badwaik et al. (2019), a DG method for 2-D Euler equations was developed on triangular grids and the mesh vertices are moved in an almost Lagrangian manner. To maintain good mesh quality, only a local remeshing was performed in regions where the mesh quality has become poor. As an example, Figure 7 shows an isentropic vortex that is also advecting. Due to the high shearing inside the vortex, the mesh would become heavily skewed with time, but our remeshing scheme is able to maintain good mesh quality over long time intervals.

Finally, we pursued a third direction. Would it be possible to use a moving mesh composed of Voronoi cells, as employed by AREPO for a second order finite volume method, also for a high order DG method? The difficulty here is that this type of mesh is effectively remeshed everywhere at every time step, implying that the cell connectivity and their topology can change during a timestep. This represents a significant challenge for DG. Note that a higher order method is in particular high



**Fig. 6** Convergence of the MHD isodensity vortex problem computed for different DG orders. Solution  $L^2$  errors are plotted for  $B_x$  (top row) and pressure (bottom row), for both a locally divergence-free (LDF) basis with Powell terms (left column), and a Legendre basis with hyperbolic cleaning (right column). Errors are measured at time  $t = 20$  after the vortex has crossed the whole computational domain. Dotted lines show theoretical slopes for convergence orders 2 to 6. Solid and dashed lines correspond to errors for an advection angle  $\alpha = 45^\circ$  and  $\alpha = 30^\circ$ , respectively. The shaded area corresponds to a range of resolutions for which the vortex is resolved but not over-resolved (see Guillet et al., 2019).

order in time, meaning it needs predictor steps between two consecutive time steps. Thus the connectivity of a cell at one time step with the cell at the next time step is needed, but a moving Voronoi cell may change the number of its sides during the step. This leads to significant geometric complications in formulating a consistent time evolution of the basis function expansion. The associated challenges were successfully overcome in Gaburro et al. (2019), for the moment in two-dimensions only. But it appears conceptually straightforward to generalize this solution to three

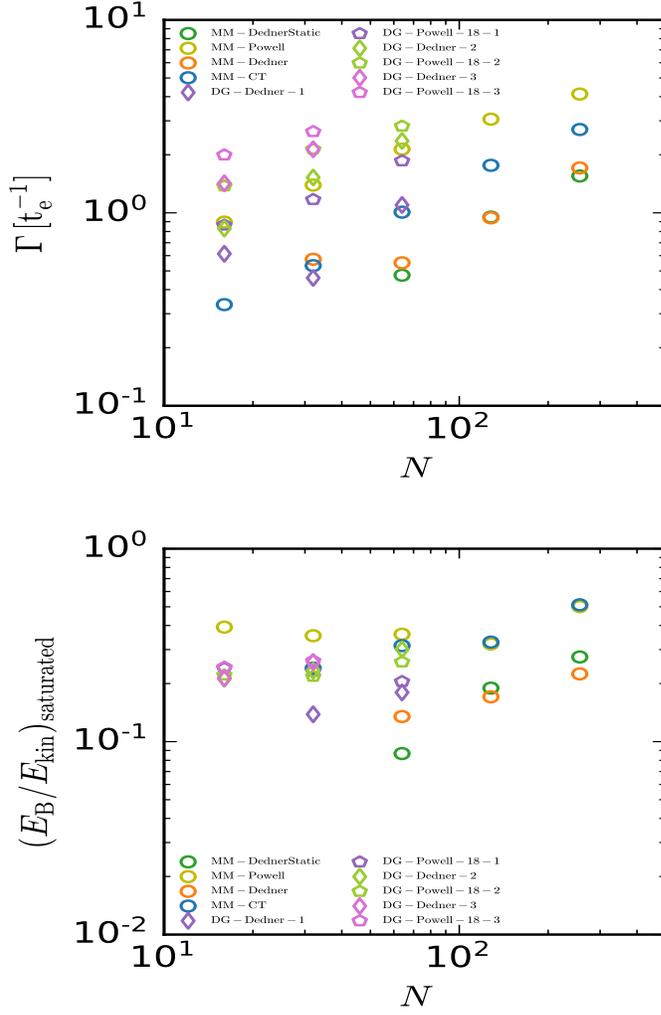


**Fig. 7** A rotating isentropic vortex in two space dimensions: we show the mesh and the pressure of the solution to the Euler equations at various times: (a)  $t = 0$ , (b)  $t = 6$ , (c)  $t = 12$ , and (d)  $t = 20$ . Notice how the mesh quality does not deteriorate thanks to the local remeshing technique described in Badwaik et al. (2019).

dimensions, therefore the path to a high order DG method in 3D using the moving Voronoi mesh of AREPO is now open.

### 3.3 Turbulence simulations with higher order numerics

We have applied our hydrodynamic DG code developed in Schaal et al. (2015) to the simulation of three-dimensional hydrodynamic subsonic turbulence (Bauer et al.,



**Fig. 8** Growth rate of the magnetic field strength in units of the eddy-turnover time  $t_{\text{ed}} = L/(2\mathcal{M}c_s)$  (top panel) and the ratio of the magnetic to the kinetic energy in the saturated state (bottom panel) for isothermal turbulence in a uniform box for different numerical schemes and spatial resolution. The turbulence is driven by purely solenoidal driving and saturates at a Mach number of  $\mathcal{M} \sim 0.3$ . We compare different numerical schemes to control the divergence of the magnetic field, including constrained transport (CT), Dedner cleaning, and Powell terms. Also, we compare finite volume on a moving mesh (MM) with high-order DG on a Cartesian mesh.

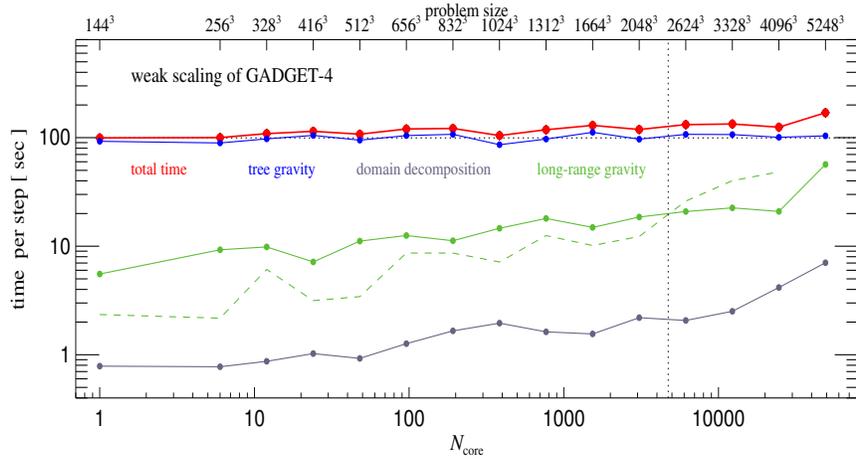
2016). This allowed us to demonstrate that this DG implementation gives accurate results at noticeably less computational cost than a finite volume method.

Driven magnetohydrodynamical turbulence is an even richer physics problem, which is of particular importance in a variety of astrophysical contexts, including star formation in the interstellar medium, stellar atmospheres, and the X-ray emitting gas in clusters of galaxies. In Pakmor et al. (2019, in prep), we have applied our DG-MHD code developed in Guillet et al. (2019) to the problem of isothermal turbulence in a uniform box, with the goal to test different numerical schemes for preserving the MHD constraint, and for validating the effectiveness of the high-order DG approach. Of particular interest is whether there are any systematic differences between a constrained transport (CT) finite-volume MHD approach, which is able to guarantee the divergence free constraint to machine precision at all times, and the Powell and Dedner approaches for divergence control, for which we also have high order DG formulations. It is sometimes suspected that CT may be required to obtain truly accurate solutions for this problem. Reassuringly, our results, summarized in Figure 8, do not support this view. The statistical properties of the quasi-stationary turbulent flow are very consistent between the different schemes. As expected, the growth rate of the turbulent dynamo increases with resolution and order of the scheme. While the slope of this relation is similar for all schemes, there are interesting differences in the absolute growth rate at a given resolution between the different schemes. If anything, here CT appears slightly more dissipative than the Powell approach. In contrast, the saturated magnetic energy does not seem to depend on the resolution or scheme, provided a minimum resolution is used that again depends on the scheme. Interestingly, here the Powell approach is able to numerically represent a working dynamo already at lower resolution than the CT approach.

#### **4 Performance and public release of the two cosmological hydrodynamical simulation codes GADGET-4 and AREPO**

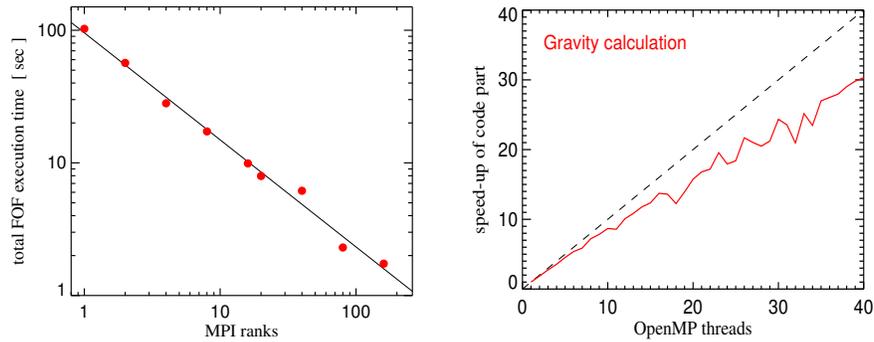
As part of EXAMAG, we have also developed a memory efficient and fast N-body/hydrodynamical code, GADGET-4, which is primarily intended for extremely large simulations of cosmic structure formation ( $\text{Gpc}^3$  volumes), targeting cosmological applications. GADGET-4 (Springel et al., 2019, in prep) represents a complete rewrite of the successful and widely used GADGET code (Springel, 2005), using C++ and numerous refined algorithms. For example, it supports a variety of additional gravity solvers, among them a high-order fast multipole method, as well as hierarchical local time-integration techniques.

The code is highly scalable, and can be run with two different approaches for hybrid parallelization, either a mix of MPI and OpenMP parallelization, or a novel shared memory parallelization model based on MPI-3 where one MPI rank is set aside on each shared memory node to respond to communication requests from MPI processes on remote nodes with minimum latency, thereby realizing truly one-sided communication independent of MPI progress engines. Within each node, the MPI layer can be bypassed entirely through shared-memory accesses in this method.



**Fig. 9** Scaling of the GADGET-4 code for cosmological simulations of structure formation on the SuperMUC-NG supercomputer. The scaling of the dominant tree calculation is excellent. The long-range gravity is done through very large FFT grids, which are communication bound. Our column-based parallel FFT (solid green) scales beyond the point where the number of MPI ranks exceeds the number of mesh planes (vertical dotted line), but requires more transpose operations than a slab-based algorithm (dashed green). For the largest number of nodes, two or more islands of the machine are needed, reducing the available cross-sectional bandwidth, which impacts the scaling behaviour of the communication heavy long-range gravity and domain decomposition.

Figure 9 shows the effectiveness of this approach on the SuperMUC-NG machine at the Leibniz Supercomputing Centre (LRZ) in a weak scaling test. The most expensive part of the calculation, the computation of the short-range gravity with a hierarchical multipole expansion (a tree algorithm) scales *perfectly* to 49152 cores (1024 nodes on SuperMUC-NG), thanks to the MPI-3 parallelization schemes, which is able to eliminate a mid-step synchronization point. The FFT-based calculation of the long-range gravity is communication-bandwidth bound and shows poorer scalability, as expected, but still stays subdominant overall. Here for the largest problem sizes an additional scaling bottleneck is resolved by GADGET-4. For a standard slab-based decomposition of the FFT (green dashed lines), there comes a point when there are more MPI ranks than mesh planes (marked by the vertical dotted line), at which point not only scalability ends, but also memory imbalance will quickly grow. This impasse is overcome in GADGET-4 with a column-based parallel FFT algorithm (green solid lines), which maintains scalability and memory balance up to the largest foreseeable problem sizes in cosmology. However, this algorithm requires twice as many transpose operations, making it more costly for small problem sizes where the slab-based approach is still viable. Like the parallel FFT, the domain decomposition algorithm is also communication bound and thus deviates from ideal scalability, but this part of GADGET-4 is fast enough to always stay subdominant. We note that the alternative hybrid parallelization through a combination of MPI



**Fig. 10** The left panel shows a strong scaling test of the built-in FOF group finder in GADGET-4, while the right panels shows the speed-up of the gravity calculation in GADGET-4 as a function of the number of OpenMP threads employed when run in MPI-OpenMP hybrid mode. In the latter example, a single MPI rank on a single node of Intel Xeon 6138 cluster was used. Even though the node has two processors with 20 physical cores each, the OpenMP scaling extends well into the second processor, despite the reduced memory bandwidth this entails. In practice, GADGET-4 is best run with at least 2 MPI ranks per processor (corresponding to up to 10 OpenMP threads on this cluster). In this regime, the OpenMP scaling is excellent.

and OpenMP yields very good thread scalability, see Figure 10 (right panel), but shows slightly poorer overall scalability when the number of shared-memory nodes becomes large due to losses in its MPI communication algorithm (which still contains a midstep synchronization point).

We have also developed new on-the-fly group finding and merger-tree building techniques for GADGET-4 (which also scale well, see Fig. 10, left panel), as well as sophisticated outputting strategies for light-cones and high angular resolution maps of line-of-sight projections of various quantities, such as the total mass (for weak lensing). These features are designed to support collisionless N-body simulations with extreme particle numbers in the regime of  $10^{12}$  particles and beyond, where the huge data volume necessitates that post-processing calculations are done during the simulation as much as possible to avoid the need for enormous disk storage capacity. In fact, GADGET-4 in principle removes the need to produce any time-slices of particle data, thereby eliminating a substantial obstacle to carry out semi-analytic galaxy formation on merger trees based on simulations in the trillion particle regime. Attempting this with the same approach as in the Millennium simulation (Springel et al., 2005), where of order 100 time slices were produced and merger trees were made in post-processing, would require 6 PB of particle storage, something that we can completely avoid with the new code.

As a legacy of EXAMAG, both the GADGET-4 and AREPO codes are scheduled for public release in 2019. Weinberger et al. (2019, in prep) introduces the community version of the AREPO code, and provides an overview of the available functionality as well an introduction of AREPO to new users. This is augmented with a suite of

examples of different complexity, a test suite, and a support forum that is hosted on the code's website as a platform to ask questions about the code. A similar release is planned for the GADGET-4 code (Springel et al., 2019, in prep). We hope that this will foster an active and supportive user community that will contribute to the further development of these highly parallel simulation codes, preparing for their eventual use on exascale class computers, as envisioned by SPPEXA and its EXAMAG subproject.

## 5 Summary and discussion

The primary research goals of the EXAMAG project have been to develop new mathematical methods and physics implementations in state-of-the-art hydrodynamical codes, allowing them to be used for groundbreaking astrophysical research that can make full use of the capabilities of current and emerging HPC platforms. Our approach consisted of leveraging a tight collaboration between applied mathematicians and numerical astrophysicists, thereby allowing a quick transfer of new mathematical ideas into applications at the forefront of today's supercomputer applications in astrophysics and cosmology.

In hindsight, we feel that our strategy to immediately apply new numerical procedures in large application projects has been successful. This provided immediate feedback on the most promising development directions, and thus allowed us to iteratively improve codes used for production science on powerful supercomputers. Our special focus on treatments of ideal magnetohydrodynamics has allowed us to make significant progress on the physical fidelity of simulations of galaxy formation, thereby making the IllustrisTNG and Auriga projects possible in the first place. The MHD capability also provided the foundations for new solvers we developed for anisotropic diffusive transport processes, relevant especially for cosmic rays (Pakmor et al., 2016a), thermal conduction (Kannan et al., 2016) and radiative transport (Kannan et al., 2019).

There is no shortage of ideas for developing the performance and capabilities of our AREPO, GADGET-4 and DG codes further in the future. Extending our DG-MHD techniques to high-order methods for self-gravity and source terms such as radiative cooling are an obvious direction. Other challenges lie in the technical aspects of parallelization, where especially the MPI-3 based shared-memory approach that we have introduced in GADGET-4 looks particularly promising for adoption in AREPO as well. It is clear that dedicated and sustained research efforts in numerical method development remain the basis for future scientific progress in computational astrophysics.

**Acknowledgements** The authors would like to thank the large user bases of the GADGET-2, GADGET-3, and AREPO codes for their continued use of these codes in a large number of scientific applications, thereby greatly motivating the development work carried out in the EXAMAG project. We are also grateful for computer time awarded through GCS at HLRS, LRZ and JSC.

## References

- Badwaik J, Chandrashekar P, Klingenberg C (2019) Single-step Arbitrary Lagrangian-Eulerian discontinuous Galerkin method for 1-D Euler equations. submitted to Communication on Applied Mathematics and Computation (CAMC)
- Bauer A, Schaal K, Springel V, Chandrashekar P, Pakmor R, Klingenberg C (2016) Simulating Turbulence Using the Astrophysical Discontinuous Galerkin Code TENET. Lecture Notes in Computational Science and Engineering (LNCSE), Springer Verlag
- Chandrashekar P, Klingenberg C (2016) Entropy stable finite volume scheme for ideal compressible MHD on 2-d cartesian meshes. *SIAM Journal on Numerical Analysis*, Vol 54, No 2
- Chandrashekar P, Gallego P, Klingenberg C (2018) A Runge-Kutta discontinuous Galerkin scheme for the ideal Magnetohydrodynamical model. C Klingenberg and M Westdickenberg (eds), *Theory, Numerics and Applications of Hyperbolic Problems*, Springer Proceedings in Mathematics Statistics (PROMS) 236
- Gaburro E, Boscheri W, Chiocchetti S, Klingenberg C, Springel V, Dumbser M (2019) High order direct Arbitrary-Lagrangian-Eulerian schemes on moving Voronoi meshes with topology changes. submitted to *Journal of Scientific Computing*
- Gallego P, Löbber J, Bastian P, Müthing S, Klingenberg C, Xia Y (2014) Implementing a discontinuous Galerkin method for the compressible, inviscid Euler equations in the DUNE framework. *Proceedings in Applied Mathematics and Mechanics*, Vol 14,1
- Gallego P, Klingenberg C, Chandrashekar P (2016) On limiting for higher order discontinuous Galerkin methods for 2D Euler equations. *Bulletin of the Brazilian Math Society, New Series*, March 2016, Volume 47, Issue 1
- Grand RJJ, Gómez FA, Marinacci F, Pakmor R, Springel V, Campbell DJR, Frenk CS, Jenkins A, White SDM (2017) The Auriga Project: the properties and formation mechanisms of disc galaxies across cosmic time. *Monthly Notices of the Royal Astronomical Society* 467(1):179–207, DOI 10.1093/mnras/stx071, 1610.01159
- Guillet T, Pakmor R, Springel V, Chandrashekar P, Klingenberg C (2019) High-order Magnetohydrodynamics for Astrophysics with an Adaptive Mesh Refinement Discontinuous Galerkin Scheme. *Monthly Notices of the Royal Astronomical Society*, Volume 485, Issue 3,
- Kannan R, Springel V, Pakmor R, Marinacci F, Vogelsberger M (2016) Accurately simulating anisotropic thermal conduction on a moving mesh. *Monthly Notices of the Royal Astronomical Society* 458(1):410–424, DOI 10.1093/mnras/stw294, 1512.03053
- Kannan R, Vogelsberger M, Marinacci F, McKinnon R, Pakmor R, Springel V (2019) AREPO-RT: radiation hydrodynamics on a moving mesh. *Monthly Notices of the Royal Astronomical Society* 485(1):117–149, DOI 10.1093/mnras/stz287, 1804.01987
- Klingenberg C, Pörner F, Xia Y (2017a) An efficient implementation of the divergence free constraint in a discontinuous Galerkin method for magnetohydrody-

- namics on unstructured meshes. *Communication in Computational Physics*, Vol 21, No 2
- Klingenberg C, Schnücke G, Xia Y (2017b) An arbitrary Lagrangian-Eulerian local discontinuous Galerkin method for Hamilton- Jacobi equations. *Journal of Scientific Computing*, vol 86, no 305
- Klingenberg C, Schnücke G, Xia Y (2017c) Arbitrary Lagrangian-Eulerian discontinuous Galerkin method for conservation laws: analysis and application in one dimension. *Mathematics of Computations*, Vol 86
- Klingenberg C, Schnücke G, Xia Y (2018) An arbitrary Lagrangian-Eulerian discontinuous Galerkin method for conservation laws: Entropy stability. C Klingenberg and M Westdickenberg (eds), *Theory, Numerics and Applications of Hyperbolic Problems*, Springer Proceedings in Mathematics Statistics (PROMS) 236
- Marinacci F, Vogelsberger M, Pakmor R, Torrey P, Springel V, Hernquist L, Nelson D, Weinberger R, Pillepich A, Naiman J, Genel S (2018) First results from the IllustrisTNG simulations: radio haloes and magnetic fields. 480:5113–5139, DOI 10.1093/mnras/sty2206, 1707.03396
- Nelson D, Pillepich A, Springel V, Weinberger R, Hernquist L, Pakmor R, Genel S, Torrey P, Vogelsberger M, Kauffmann G, Marinacci F, Naiman J (2018a) First results from the IllustrisTNG simulations: the galaxy colour bimodality. 475:624–647, DOI 10.1093/mnras/stx3040, 1707.03395
- Nelson D, Springel V, Pillepich A, Rodriguez-Gomez V, Torrey P, Genel S, Vogelsberger M, Pakmor R, Marinacci F, Weinberger R, Kelley L, Lovell M, Diemer B, Hernquist L (2018b) The IllustrisTNG Simulations: Public Data Release. arXiv e-prints, 181205609 1812.05609
- Pakmor R, Marinacci F, Springel V (2014) Magnetic Fields in Cosmological Simulations of Disk Galaxies. *The Astrophysical Journal* 783(1):L20, DOI 10.1088/2041-8205/783/1/L20, 1312.2620
- Pakmor R, Pfrommer C, Simpson CM, Kannan R, Springel V (2016a) Semi-implicit anisotropic cosmic ray transport on an unstructured moving mesh. *Monthly Notices of the Royal Astronomical Society* 462(3):2603–2616, DOI 10.1093/mnras/stw1761, 1604.08587
- Pakmor R, Springel V, Bauer A, Mocz P, Munoz DJ, Ohlmann ST, Schaal K, Zhu C (2016b) Improving the convergence properties of the moving-mesh code AREPO. *Monthly Notices of the Royal Astronomical Society* 455(1):1134–1143, DOI 10.1093/mnras/stv2380, 1503.00562
- Pakmor R, Gómez FA, Grand RJJ, Marinacci F, Simpson CM, Springel V, Campbell DJR, Frenk CS, Guillet T, Pfrommer C, White SDM (2017) Magnetic field formation in the Milky Way like disc galaxies of the Auriga project. *Monthly Notices of the Royal Astronomical Society* 469(3):3185–3199, DOI 10.1093/mnras/stx1074, 1701.07028
- Pakmor R, Guillet T, Pfrommer C, Gómez FA, Grand RJJ, Marinacci F, Simpson CM, Springel V (2018) Faraday rotation maps of disc galaxies. *Monthly Notices of the Royal Astronomical Society* 481(4):4410–4418, DOI 10.1093/mnras/sty2601, 1807.02113

- Pillepich A, Nelson D, Hernquist L, Springel V, Pakmor R, Torrey P, Weinberger R, Genel S, Naiman JP, Marinacci F, Vogelsberger M (2018) First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. 475:648–675, DOI 10.1093/mnras/stx3112, 1707.03406
- Schaal K, Bauer A, Chandrashekar P, Pakmor R, Klingenberg C, Springel V (2015) Astrophysical hydrodynamics with a high-order discontinuous Galerkin scheme and adaptive mesh refinement. *Monthly Notices Royal Astronomical Society*, Volume 453, Issue 4
- Shu C (2014) Discontinuous Galerkin method for time-dependent problems: survey and recent developments. Feng, X, Karakashian, O, Xing, Y (Eds), *Recent Developments in Discontinuous Galerkin Finite Element Methods for Partial Differential Equations*, Springer Verlag
- Springel V (2005) The cosmological simulation code GADGET-2. 364:1105–1134, DOI 10.1111/j.1365-2966.2005.09655.x, astro-ph/0505010
- Springel V (2010) E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society* 401(2):791–851, DOI 10.1111/j.1365-2966.2009.15715.x, 0901.4107
- Springel V, White SDM, Jenkins A, Frenk CS, Yoshida N, Gao L, Navarro J, Thacker R, Croton D, Helly J, Peacock JA, Cole S, Thomas P, Couchman H, Evrard A, Colberg J, Pearce F (2005) Simulations of the formation, evolution and clustering of galaxies and quasars. 435:629–636, DOI 10.1038/nature03597, astro-ph/0504097
- Springel V, Pakmor R, Pillepich A, Weinberger R, Nelson D, Hernquist L, Vogelsberger M, Genel S, Torrey P, Marinacci F, Naiman J (2018) First results from the IllustrisTNG simulations: matter and galaxy clustering. 475:676–698, DOI 10.1093/mnras/stx3304, 1707.03397
- Torrey P, Vogelsberger M, Hernquist L, McKinnon R, Marinacci F, Simcoe RA, Springel V, Pillepich A, Naiman J, Pakmor R, Weinberger R, Nelson D, Genel S (2018) Similar star formation rate and metallicity variability time-scales drive the fundamental metallicity relation. 477:L16–L20, DOI 10.1093/mnrasl/sly031, 1711.11039
- Vogelsberger M, Genel S, Springel V, Torrey P, Sijacki D, Xu D, Snyder G, Bird S, Nelson D, Hernquist L (2014) Properties of galaxies reproduced by a hydrodynamic simulation. 509:177–182, DOI 10.1038/nature13316, 1405.1418