EXAMAG: 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.

V. Springel (✉) · R. Pakmor

Max-Planck-Institute for Astrophysics, München, Germany

e-mail: vspringel@mpa-garching.mpg.de; rpakmor@mpa-garching.mpg.de

C. Klingenberg

Universität Würzburg, Würzburg, Germany

e-mail: klingen@mathematik.uni-wuerzburg.de

T. Guillet

University of Exeter, Exeter, UK

e-mail: T.A.Guillet@exeter.ac.uk

P. Chandrashekar

TIFR Centre for Applicable Mathematics, Bengaluru, India

e-mail: praveen@tifrbng.res.in

© The Author(s) 2020

H.-J. Bungartz et al. (eds.), *Software for Exascale Computing - SPPEXA*

2016–2019, Lecture Notes in Computational Science and Engineering 136,

https://doi.org/10.1007/978-3-030-47956-5_12

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 Sect. 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 Sect. 3, we turn to our recent developments of discontinuous Galerkin hydro- and magnetohydrodynamics codes. 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 Sect. 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 Sect. 5.

2 The IllustrisTNG and Auriga Simulations

The AREPO code [31] 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 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 [35] demonstrated the utility of the approach for simulations 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 project¹ built on the technical and scientific achievements of its predecessor and pushed this line of research further. In particular, it has improved upon Illustris by including our newly developed accurate solver for ideal magnetohydrodynamics [22, 24], 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 24,000 cores, required 100 TB RAM, and produced 660 TB of scientific data. A visual impression of one of these simulation is given in Fig. 1. As an illustrative result, we show the clustering statistics of different matter components at the present epoch in TNG300 in Fig. 2 [33]. 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. 19, 20, 27, 33, 34], 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 [21], with TNG50 to follow in a year’s time, which will further amplify the scientific use of the simulations.

The Auriga simulations [11] 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

---

¹<http://www.tng-project.org>.

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

Fig. 2 The matter autocorrelation function for different mass components in our high-resolution TNG300 simulation at redshift $z = 0$, see [33]. 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

well as their build-up over cosmic time. In [25] 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% of the turbulent energy density with a typical strength of 10–50 μG . 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 [26], for different observer positions within and outside the simulated galaxies. We find that the strength of the Faraday

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 [26]. These predictions compare well to observations of real galaxies such as M51

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 Fig. 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 the approximate solution is not required to 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 [29].

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 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 Sect. 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 Sect. 3.2, the discontinuous Galerkin method on a moving mesh is discussed. Finally, in Sect. 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 [9].

In [10], 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 [28]. This version of the code supports adaptive mesh refinement (AMR) and is able to maintain high-order accuracy at AMR refinement

Fig. 4 A Kelvin-Helmholtz simulation with fourth order DG and adaptive mesh refinement (see [28]). 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

boundaries, unlike finite-volume approaches that typically fall-back to first order there. As an illustrative example, Fig. 4 depicts a simulation of a Kelvin-Helmholtz instability using TENET. In [28] we have also shown that DG has 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 [5] 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 [6] 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 [10], where these techniques were introduced for the hydrodynamical case. We assessed a different approach for the divergence constraint in [15], where post- and pre-processing methods are suggested in order to numerically maintain the divergence free constraint.

Finally, in [12] 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. [12] 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 Fig. 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 Fig. 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 2–6, both for a locally divergence-free basis and a Legendre basis. Note that for increasing order progressively more degrees of freedom per cell are needed, and thus more storage and computational work per element are required. This enlargement of the cost per element is a constant factor, however, leaving the slopes of the convergence order unaffected. The high-order methods hence always win in overall efficiency

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 [12]

and accuracy once the mesh resolution lies above a finite cross-over point (see also [12, 28] for an analysis of the accuracy of different DG-order as a function of invested CPU-time).

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 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 [16] and [17] for a semidiscrete arbitrary Lagrangian-Eulerian discontinuous Galerkin method. In [18] 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 [1], 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, Fig. 7 shows an isentropic vortex that is also advecting. Due to the

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 [12])

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

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 [1]

significant challenge for DG, although some guidance may also be obtained from high-order curvilinear finite element methods (e.g. [7]). Note that a higher order method is in particular high 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 [8], for the moment in two-dimensions only. But it appears conceptually straightforward to generalize this

solution to three 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 [28] to the simulation of three-dimensional hydrodynamic subsonic turbulence [4]. 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 [12] 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 Fig. 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 (Gpc^3 volumes), targeting cos-

Fig. 8 Growth rate of the magnetic field strength in units of the eddy-turnover time $t_{\text{ed}} = L/(2M c_s)$ (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. N is the number of mesh cells per spatial dimension. The turbulence is driven by purely solenoidal driving and saturates at a Mach number of $M \sim 0.3$. 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

mological applications. GADGET-4 (Springel et al., 2019, in prep) represents a complete rewrite of the successful and widely used GADGET code [30], using C++ and numerous refined algorithms. For example, it supports a variety of additional gravity solvers, among them a high-order fast multipole method (FMM), 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.

Figure 9 shows the effectiveness of this approach on the SuperMUC-NG machine at the Leibniz Supercomputing Centre (LRZ) in a weak scaling test, where the TreePM force computation algorithm [2, 30] is used. This approach assembles the total force as the sum of a short-range gravitational force computed with a hierarchical multipole expansion (here done with a one-sided tree algorithm, [3]), and a long-range gravity computed in Fourier space. This approach is particularly efficient for periodic cosmological simulations at high redshift where the density fluctuations are small and the large-scale residual forces nearly vanish. In this regime, pure Tree- or FMM-algorithms need to open many more nodes to accurately recover the near cancellation of most of the large-scale forces.

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 behavior of the communication heavy long-range gravity and domain decomposition

The most expensive part of the calculation, the Tree-based computation of the short-range gravity scales *perfectly* to 49,152 cores (1024 nodes on SuperMUC-NG), thanks to the MPI-3 parallelization scheme, which is able to eliminate a mid-step synchronization point that was still necessary in our older conventional communication scheme for this algorithm. Because the amount of work per MPI rank stays constant in the weak scaling regime for the short-range part of the TreePM algorithm, and the amount of data that needs to be imported from neighboring ranks stays approximately constant, too, a near perfect scaling should in principle be reachable. However, because the communication pattern is complex and highly irregular, this has normally become a bottleneck for the scalability to very large numbers of MPI ranks. This is here successfully eliminated thanks to our one-sided (and fully portable) communication approach.

In contrast, 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 and OpenMP yields very good thread scalability, see Fig. 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. 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 [32], where of order 100 time slices were produced and merger trees were made in

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

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. [36] 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 [23], thermal conduction [13] and radiative transport [14].

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.

Acknowledgments 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

1. Badwaik, J., Chandrashekar, P., Klingenberg, C.: Single-step arbitrary Lagrangian-Eulerian discontinuous Galerkin method for 1-D Euler equations. Commun. Appl. Math. Comput. (2019, submitted). Preprint arXiv:1602.09079
2. Bagla, J.S.: TreePM: a code for cosmological N-body simulations. J. Astrophys. Astron. **23**, 185–196 (2002)
3. Barnes, J., Hut, P.: A hierarchical $O(N \log N)$ force-calculation algorithm. Nature **324**(6096), 446–449 (1986)
4. Bauer, A., Schaal, K., Springel, V., Chandrashekar, P., Pakmor, R., Klingenberg, C.: Simulating turbulence using the astrophysical discontinuous Galerkin Code TENET. In: Proceedings of the SPPEXA Symposium. Lecture Notes in Computational Science and Engineering. Springer, Berlin (2016). arXiv:1602.09079
5. Chandrashekar, P., Klingenberg, C.: Entropy stable finite volume scheme for ideal compressible MHD on 2-D Cartesian meshes. SIAM J. Numer. Anal. **54**(2), 1313–1340 (2016)
6. Chandrashekar, P., Gallego, P., Klingenberg, C.: A Runge-Kutta discontinuous Galerkin scheme for the ideal Magnetohydrodynamical model. In: Klingenberg, C., Westdickenberg, M. (eds.) Theory, Numerics and Applications of Hyperbolic Problems, Springer Proceedings in Mathematics & Statistics (PROMS), vol. 236, pp. 1203–1232 (2018)
7. Dobrev, V.A., Kolev, T.V., Rieben, R.N.: High-order curvilinear finite element methods for Lagrangian hydrodynamics. SIAM J. Sci. Comput. **34**(5), B606–B641 (2012)

8. Gaburro, E., Boscheri, W., Chiochetti, S., Klingenberg, C., Springel, V., Dumbser, M.: High order direct Arbitrary-Lagrangian-Eulerian schemes on moving Voronoi meshes with topology changes. *J. Comput. Phys.* **407**, 109167 (2020). <https://doi.org/10.1016/j.jcp.2019.109167>
9. Gallego, P., Löbber, J., Bastian, P., Muthing, S., Klingenberg, C., Xia, Y.: Implementing a discontinuous Galerkin method for the compressible, inviscid Euler equations in the DUNE framework. In: *Proceedings in Applied Mathematics and Mechanics*, vol. 14, 1 pp. 25–62 (2014)
10. Gallego, P., Klingenberg, C., Chandrashekar, P.: On limiting for higher order discontinuous Galerkin methods for 2D Euler equations. *Bull. Braz. Math. Soc. (New Series)* **47**(1), 335–345 (2016)
11. Grand, R.J.J., Gómez, F.A., Marinacci, F., Pakmor, R., Springel, V., Campbell, D.J.R., Frenk, C.S., Jenkins, A., White, S.D.M.: The Auriga Project: the properties and formation mechanisms of disc galaxies across cosmic time. *Mon. Not. R. Astron. Soc.* **467**(1), 179–207 (2017)
12. Guillet, T., Pakmor, R., Springel, V., Chandrashekar, P., Klingenberg, C.: High-order magnetohydrodynamics for astrophysics with an adaptive mesh refinement discontinuous Galerkin scheme. *Mon. Not. R. Astron. Soc.* **485**(3), 4209–4246 (2019)
13. Kannan, R., Springel, V., Pakmor, R., Marinacci, F., Vogelsberger, M.: Accurately simulating anisotropic thermal conduction on a moving mesh. *Mon. Not. R. Astron. Soc.* **458**(1), 410–424 (2016)
14. Kannan, R., Vogelsberger, M., Marinacci, F., McKinnon, R., Pakmor, R., Springel, V.: AREPO-RT: radiation hydrodynamics on a moving mesh. *Mon. Not. R. Astron. Soc.* **485**(1), 117–149 (2019)
15. Klingenberg, C., Pörner, F., Xia, Y.: An efficient implementation of the divergence free constraint in a discontinuous Galerkin method for magnetohydrodynamics on unstructured meshes. *Commun. Comput. Phys.* **21**(2), 423–442 (2017)
16. Klingenberg, C., Schnücke, G., Xia, Y.: An arbitrary Lagrangian-Eulerian local discontinuous Galerkin method for Hamilton–Jacobi equations. *J. Sci. Comput.* **86**(305), 1203–1232 (2017)
17. Klingenberg, C., Schnücke, G., Xia, Y.: Arbitrary Lagrangian-Eulerian discontinuous Galerkin method for conservation laws: analysis and application in one dimension. *Math. Comput.* **86**, 335–345 (2017)
18. Klingenberg, C., Schnücke, G., Xia, Y.: An arbitrary Lagrangian-Eulerian discontinuous Galerkin method for conservation laws: entropy stability. In: Klingenberg, C., Westdickenberg, M. (eds.) *Theory, Numerics and Applications of Hyperbolic Problems*, Springer Proceedings in Mathematics & Statistics (PROMS), vol. 236, pp. 1203–1232 (2018)
19. Marinacci, F., Vogelsberger, M., Pakmor, R., Torrey, P., Springel, V., Hernquist, L., Nelson, D., Weinberger, R., Pillepich, A., Naiman, J., Genel, S.: First results from the IllustrisTNG simulations: radio haloes and magnetic fields. *Mon. Not. R. Astron. Soc.* **480**, 5113–5139 (2018)
20. Nelson, D., Pillepich, A., Springel, V., Weinberger, R., Hernquist, L., Pakmor, R., Genel, S., Torrey, P., Vogelsberger, M., Kauffmann, G., Marinacci, F., Naiman, J.: First results from the IllustrisTNG simulations: the galaxy colour bimodality. *Mon. Not. R. Astron. Soc.* **475**, 624–647 (2018)
21. 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.: The IllustrisTNG simulations: public data release. *Comput. Astrophys. Cosmol.* **6**(1), 2 (2019)
22. Pakmor, R., Marinacci, F., Springel, V.: Magnetic fields in cosmological simulations of disk galaxies. *Astrophys. J.* **783**(1), L20 (2014)
23. Pakmor, R., Pfrommer, C., Simpson, C.M., Kannan, R., Springel, V.: Semi-implicit anisotropic cosmic ray transport on an unstructured moving mesh. *Mon. Not. R. Astron. Soc.* **462**(3), 2603–2616 (2016)
24. Pakmor, R., Springel, V., Bauer, A., Mocz, P., Munoz, D.J., Ohlmann, S.T., Schaal, K., Zhu, C.: Improving the convergence properties of the moving-mesh code AREPO. *Mon. Not. R. Astron. Soc.* **455**(1), 1134–1143 (2016)

25. Pakmor, R., Gómez, F.A., Grand, R.J.J., Marinacci, F., Simpson, C.M., Springel, V., Campbell, D.J.R., Frenk, C.S., Guillet, T., Pfrommer, C., White, S.D.M.: Magnetic field formation in the Milky Way like disc galaxies of the Auriga project. *Mon. Not. R. Astron. Soc.* **469**(3), 3185–3199 (2017)
26. Pakmor, R., Guillet, T., Pfrommer, C., Gómez, F.A., Grand, R.J.J., Marinacci, F., Simpson, C.M., Springel, V.: Faraday rotation maps of disc galaxies. *Mon. Not. R. Astron. Soc.* **481**(4), 4410–4418 (2018)
27. Pillepich, A., Nelson, D., Hernquist, L., Springel, V., Pakmor, R., Torrey, P., Weinberger, R., Genel, S., Naiman, J.P., Marinacci, F., Vogelsberger, M.: First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. *Mon. Not. R. Astron. Soc.* **475**, 648–675 (2018)
28. Schaal, K., Springel, V.: Shock finding on a moving mesh - I. Shock statistics in non-radiative cosmological simulations. *Mon. Not. R. Astron. Soc.* **446**(4), 3992–4007 (2015)
29. Shu, C.: Discontinuous Galerkin method for time-dependent problems: survey and recent developments. In: Feng, X., Karakashian, O., Xing, Y. (eds.) *Recent Developments in Discontinuous Galerkin Finite Element Methods for Partial Differential Equations*, pp. 25–62. Springer, Berlin (2014)
30. Springel, V.: The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.* **364**, 1105–1134 (2005)
31. Springel, V.: E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Mon. Not. R. Astron. Soc.* **401**(2), 791–851 (2010)
32. Springel, V., White, S.D.M., Jenkins, A., Frenk, C.S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., Peacock, J.A., Cole, S., Thomas, P., Couchman, H., Evrard, A., Colberg, J., Pearce, F.: Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature* **435**, 629–636 (2005)
33. Springel, V., Pakmor, R., Pillepich, A., Weinberger, R., Nelson, D., Hernquist, L., Vogelsberger, M., Genel, S., Torrey, P., Marinacci, F., Naiman, J.: First results from the IllustrisTNG simulations: matter and galaxy clustering. *Mon. Not. R. Astron. Soc.* **475**, 676–698 (2018)
34. Torrey, P., Vogelsberger, M., Hernquist, L., McKinnon, R., Marinacci, F., Simcoe, R.A., Springel, V., Pillepich, A., Naiman, J., Pakmor, R., Weinberger, R., Nelson, D., Genel, S.: Similar star formation rate and metallicity variability time-scales drive the fundamental metallicity relation. *Mon. Not. R. Astron. Soc.* **477**, L16–L20 (2018)
35. 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. *Nature* **509**, 177–182 (2014)
36. Weinberger, R., Springel, V., Pakmor, R.: The Arepo public code release (2019). Preprint. arXiv:1909.04667

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.