



A High-Order Discontinuous Galerkin Solver with Dynamic Adaptive Mesh Refinement to Simulate Cloud Formation Processes

Lukas Krenz^(✉), Leonhard Rannabauer, and Michael Bader

Department of Informatics, Technical University of Munich, Munich, Germany
{lukas.krenz,rannabau,bader}@in.tum.de

Abstract. We present a high-order discontinuous Galerkin (DG) solver of the compressible Navier-Stokes equations for cloud formation processes. The scheme exploits an underlying parallelized implementation of the ADER-DG method with dynamic adaptive mesh refinement. We improve our method by a PDE-independent general refinement criterion, based on the local total variation of the numerical solution. While established methods use numerics tailored towards the specific simulation, our scheme works scenario independent. Our generic scheme shows competitive results for both classical CFD and stratified scenarios. We focus on two dimensional simulations of two bubble convection scenarios over a background atmosphere. The largest simulation here uses order 6 and 6561 cells which were reduced to 1953 cells by our refinement criterion.

Keywords: ADER-DG · Navier-Stokes · Adaptive mesh refinement

1 Introduction

In this paper we address the resolution of basic cloud formation processes on modern super computer systems. The simulation of cloud formations, as part of convective processes, is expected to play an important role in future numerical weather prediction [1]. This requires both suitable physical models and effective computational realizations. Here we focus on the simulation of simple benchmark scenarios [10]. They contain relatively small scale effects which are well approximated with the compressible Navier-Stokes equations. We use the ADER-DG method of [5], which allows us to simulate the Navier-Stokes equations with a space-time-discretization of arbitrary high order. In contrast to Runge-Kutta time integrators or semi-implicit methods, an increase of the order of ADER-DG only results in larger computational kernels and does not affect the complexity of the scheme. Additionally, ADER-DG is a communication avoiding scheme and reduces the overhead on larger scale. We see our scheme in the regime of already established methods for cloud simulations, as seen for example in [10, 12, 13].

Due to the viscous components of the Navier-Stokes equations, it is not straightforward to apply the ADER-DG formalism of [5], which addresses hyperbolic systems of partial differential equations (PDEs) in first-order formulation. To include viscosity, we use the numerical flux for the compressible Navier-Stokes equations of Gassner et al. [8]. This flux has already been applied to the ADER-DG method in [4]. In contrast to this paper, we focus on the simulation of complex flows with a gravitational source term and a realistic background atmosphere. Additionally, we use adaptive mesh refinement (AMR) to increase the spatial resolution in areas of interest. This has been shown to work well for the simulation of cloud dynamics [12]. Regarding the issue of limiting in high-order DG methods, we note that viscosity not only models the correct physics of the problem but also smooths oscillations and discontinuities, thus stabilizing the simulation.

We base our work on the ExaHyPE Engine (www.exahype.eu), which is a framework that can solve arbitrary hyperbolic PDE systems. A user of the engine is provided with a simple code interface which mirrors the parts required to formulate a well-posed Cauchy problem for a system of hyperbolic PDEs of first order. The underlying ADER-DG method, parallelization techniques and dynamic adaptive mesh refinement are available for simulations while the implementations are left as a black box to the user. An introduction to the communication-avoiding implementation of the whole numerical scheme can be found in [3].

To summarize, we make the following contributions in this paper:

- We extend the ExaHyPE Engine to allow viscous terms.
- We thus provide an implementation of the compressible Navier-Stokes equations. In addition, we tailor the equation set to stratified flows with gravitational source term. We emphasize that we use a standard formulation of the Navier-Stokes equations as seen in the field of computational fluid mechanics and only use small modifications of the governing equations, in contrast to a equation set that is tailored exactly to the application area.
- We present a general AMR-criterion that is based on the detection of outlier cells w.r.t. their total variation. Furthermore, we show how to utilize this criterion for stratified flows.
- We evaluate our implementation with standard CFD scenarios and atmospheric flows and inspect the effectiveness of our proposed AMR-criterion. We thus inspect, whether our proposed general implementation can achieve results that are competitive with the state-of-the-art models that rely on heavily specified equations and numerics.

2 Equation Set

The compressible Navier-Stokes equations in the conservative form are given as

$$\underbrace{\frac{\partial}{\partial t} \begin{pmatrix} \rho \\ \rho \mathbf{v} \\ \rho E \end{pmatrix}}_{\mathbf{Q}} + \underbrace{\nabla \cdot \left(\underbrace{\begin{pmatrix} \rho \mathbf{v} \\ \mathbf{v} \otimes \rho \mathbf{v} + I p \\ \mathbf{v} \cdot (I \rho E + I p) \end{pmatrix}}_{\mathbf{F}^h(\mathbf{Q})} + \underbrace{\begin{pmatrix} 0 \\ \boldsymbol{\sigma}(\mathbf{Q}, \nabla \mathbf{Q}) \\ \mathbf{v} \cdot \boldsymbol{\sigma}(\mathbf{Q}, \nabla \mathbf{Q}) - \kappa \nabla T \end{pmatrix}}_{\mathbf{F}^v(\mathbf{Q}, \nabla \mathbf{Q})} \right)}_{\mathbf{F}(\mathbf{Q}, \nabla \mathbf{Q})} = \underbrace{\begin{pmatrix} S_\rho \\ S_{\rho \mathbf{v}} \\ S_{\rho E} \end{pmatrix}}_{\mathbf{S}(\mathbf{Q}, \mathbf{x}, t)} \quad (1)$$

with the vector of conserved quantities \mathbf{Q} , flux $\mathbf{F}(\mathbf{Q}, \nabla \mathbf{Q})$ and source $\mathbf{S}(\mathbf{Q})$. Note that the flux can be split into a hyperbolic part $\mathbf{F}^h(\mathbf{Q})$, which is identical to the flux of the Euler equations, and a viscous part $\mathbf{F}^v(\mathbf{Q}, \nabla \mathbf{Q})$. The conserved quantities \mathbf{Q} are the density ρ , the two or three-dimensional momentum $\rho \mathbf{v}$ and the energy density ρE . The rows of Eq. (1) are the conservation of mass, the conservation of momentum and the conservation of energy.

The pressure p is given by the equation of state of an ideal gas

$$p = (\gamma - 1) \left(\rho E - \frac{1}{2} (\mathbf{v} \cdot \rho \mathbf{v}) - g z \right). \quad (2)$$

The term $g z$ is the geopotential height with the gravity of Earth g [10]. The temperature T relates to the pressure by the thermal equation of state

$$p = \rho R T, \quad (3)$$

where R is the specific gas constant of a fluid.

We model the diffusivity by the stress tensor

$$\boldsymbol{\sigma}(\mathbf{Q}, \nabla \mathbf{Q}) = \mu \left((2/3 \nabla \cdot \mathbf{v}) - (\nabla \mathbf{v} + \nabla \mathbf{v}^\top) \right), \quad (4)$$

with constant viscosity μ . The heat diffusion is governed by the coefficient

$$\kappa = \frac{\mu \gamma}{\text{Pr}} \frac{1}{\gamma - 1} R = \frac{\mu c_p}{\text{Pr}}, \quad (5)$$

where the ratio of specific heats γ , the heat capacity at constant pressure c_p and the Prandtl number Pr depend on the fluid.

Many realistic atmospheric flows can be described by a perturbation over a background state that is in hydrostatic equilibrium

$$\frac{\partial}{\partial z} \bar{p}(z) = -g \bar{\rho}(z), \quad (6)$$

i.e. a state, where the pressure gradient is exactly in balance with the gravitational source term $S_{\rho v} = -\mathbf{k}\rho g$. The vector \mathbf{k} is the unit vector pointing in z -direction. The momentum equation is dominated by the background flow in this case. Because this can lead to numerical instabilities, problems of this kind are challenging and require some care. To lessen the impact of this, we split the pressure $p = \bar{p} + p'$ into a sum of the background pressure $\bar{p}(z)$ and perturbation $p'(\mathbf{x}, t)$. We split the density $\rho = \bar{\rho} + \rho'$ in the same manner and arrive at

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\mathbf{v} \otimes \rho \mathbf{v} + \mathbf{I} p') + \mathbf{F}_\rho^v \mathbf{v} = -g \mathbf{k} \rho'. \quad (7)$$

Note that a similar and more complex splitting is performed in [10, 12]. In contrast to this, we use the true compressible Navier-Stokes equations with minimal modifications.

3 Numerics

The ExaHyPE Engine implements an ADER-DG-scheme and a MUSCL-Hancock finite volume method. Both can be considered as instances of the more general PNPm schemes of [5]. We use a Rusanov-style flux that is adapted to PDEs with viscous terms [7, 8]. The finite volume scheme is stabilized with the van Albada limiter [15]. The user can state dynamic AMR rules by supplying custom criteria that are evaluated point-wise. Our criterion uses an element-local error estimate based on the total variation of the numerical solution. We exploit the fact that the total variation of a numerical solution is a perfect indicator for edges of a wavefront. Let $\mathbf{f}(\mathbf{x}) : \mathbb{R}^{N_{\text{vars}}} \rightarrow \mathbb{R}$ be a sufficiently smooth function that maps the discrete solution at a point \mathbf{x} to an arbitrary indicator variable. The total variation (TV) of this function is defined by

$$\text{TV}[f(\mathbf{x})] = \left\| \int_C |\nabla f(\mathbf{x})| d\mathbf{x} \right\|_1 \quad (8)$$

for each cell. The operator $\|\cdot\|_1$ denotes the discrete L_1 norm in this equation. We compute the integral efficiently with Gaussian quadrature over the collocated quadrature points. How can we decide whether a cell is important or not? To resolve this conundrum, we compute the mean and the population standard deviation of the total variation of all cells. It is important that we use the method of [2] to compute the modes in a parallel and numerical stable manner. A cell is then considered to contain significant information if it deviates from the mean more than a given threshold. This criterion can be described formally by

$$\text{evaluate-refinement}(\mathbf{Q}, \mu, \sigma) = \begin{cases} \text{refine} & \text{if } \text{TV}(\mathbf{Q}) \geq \mu + T_{\text{refine}}\sigma, \\ \text{coarsen} & \text{if } \text{TV}(\mathbf{Q}) < \mu + T_{\text{coarsen}}\sigma, \\ \text{keep} & \text{otherwise.} \end{cases} \quad (9)$$

The parameters $T_{\text{refine}} > T_{\text{coarsen}}$ can be chosen freely. Chebyshev's inequality

$$\mathbb{P}(|X - \mu| \geq c\sigma) \leq \frac{1}{c^2}, \quad (10)$$

with probability \mathbb{P} guarantees that we neither mark all cells for refinement nor for coarsening. This inequality holds for arbitrary distributions under the weak assumption that they have a finite mean μ and a finite standard deviation σ [16]. Note that subcells are coarsened only if all subcells belonging to the coarse cell are marked for coarsening. In contrast to already published criteria which are either designed solely for the simulation of clouds [12] or computationally expensive [7], our criterion works for arbitrary PDEs and yet, is easy to compute and intuitive.

4 Results

In this section, we evaluate the quality of the results of our numerical methods and the scalability of our implementation. We use a mix of various benchmarking scenarios. After investigating the numerical convergence rate, we look at three standard CFD scenarios: the Taylor-Green vortex, the three-dimensional Arnold-Beltrami-Childress flow and a lid-driven cavity flow. Finally, we evaluate the performance for stratified flow scenarios in both two and three dimensions.

4.1 CFD Testing Scenarios

We begin with a manufactured solution scenario which we can use for a convergence test. We use the following constants of fluids for all scenarios in this section:

$$\gamma = 1.4, \quad \text{Pr} = 0.7, \quad c_v = 1.0. \quad (11)$$

Our description of the manufactured solution follows [4]. To construct this solution, we assume that

$$\begin{aligned} p(\mathbf{x}, t) &= 1/10 \cos(\mathbf{k}\mathbf{x} - 2\pi t) + 1/\gamma, \\ \rho(\mathbf{x}, t) &= 1/2 \sin(\mathbf{k}\mathbf{x} - 2\pi t) + 1, \\ \mathbf{v}(\mathbf{x}, t) &= \mathbf{v}_0 \sin(\mathbf{k}\mathbf{x} - 2\pi t), \end{aligned} \quad (12)$$

solves our PDE. We use the constants $\mathbf{v}_0 = 1/4(1, 1)^\top$, $\mathbf{k} = \pi/5(1, 1)^\top$ and simulate a domain of size $[10 \times 10]$ for 0.5s. The viscosity is set to $\mu = 0.1$. Note that Eq. (12) does not solve the compressible Navier-Stokes equations Eq. (1) directly. It rather solves our equation set with an added source term which can be derived with a computer algebra system. We ran this for a combination of

orders $1, \dots, 6$ and multiple grid sizes. Note that by order we mean the polynomial order throughout the entire paper and not the theoretical convergence order. For this scenario, we achieve high-order convergence (Fig. 1) but notice some diminishing returns for large orders.

After we have established that the implementation of our numerical method converges, we are going to investigate three established testing scenarios from the field of computational fluid mechanics. A simple scenario is the Taylor-Green vortex. Assuming an *incompressible* fluid, it can be written as

$$\begin{aligned}\rho(\mathbf{x}, t) &= 1, \\ \mathbf{v}(\mathbf{x}, t) &= \exp(-2\mu t) \begin{pmatrix} \sin(x) \cos(y) \\ -\cos(x) \sin(y) \end{pmatrix}, \\ p(\mathbf{x}, t) &= \exp(-4\mu t)^{1/4} (\cos(2x) + \cos(2y)) + C.\end{aligned}\tag{13}$$

The constant $C = 100/\gamma$ governs the speed of sound and thus the Mach number $\text{Ma} = 0.1$ [6]. The viscosity is set to $\mu = 0.1$.

We simulate on the domain $[0, 2\pi]^2$ and impose the analytical solution at the boundary. A comparison at time $t = 10.0$ of the analytical solution for the pressure with our approximation (Fig. 2a) shows excellent agreement. Note that we only show a qualitative analysis because this is not an exact solution for our equation set as we assume compressibility of the fluid. This is nevertheless a valid comparison because for very low Mach numbers, both incompressible and compressible equations behave in a very similarly. We used an ADER-DG-scheme of order 5 with a grid of 25^2 cells.

The Arnold-Beltrami-Childress (ABC) flow is similar to the Taylor-Green vortex but is an analytical solution for the three-dimensional *incompressible* Navier-Stokes equations [14]. It is defined in the domain $[-\pi, \pi]^3$ as

$$\begin{aligned}\rho(\mathbf{x}, t) &= 1, \\ \mathbf{v}(\mathbf{x}, t) &= \exp(-1\mu t) \begin{pmatrix} \sin(z) + \cos(y) \\ \sin(x) + \cos(z) \\ \sin(y) + \cos(x) \end{pmatrix}, \\ p(\mathbf{x}, t) &= -\exp(-2\mu t) (\cos(x) \sin(y) + \sin(x) \cos(z) + \sin(z) \cos(y)) + C.\end{aligned}\tag{14}$$

The constant $C = 100/\gamma$ is chosen as before. We use a viscosity of $\mu = 0.01$ and analytical boundary conditions. Our results (Fig. 2b) show a good agreement between the analytical solution and our approximation with an ADER-DG-scheme of order 3 with a mesh consisting of 27^3 cells at time $t = 0.1$ s. Again, we do not perform a quantitative analysis as the ABC-flow only solves our equation set approximately.


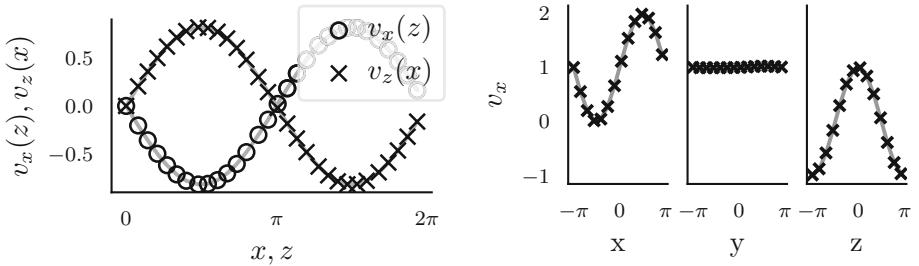


Fig. 1. Mesh size vs. error for various polynomial orders P . Dashed lines show the theoretical convergence order of $P + 1$.

As a final example of standard flow scenarios, we consider the lid-driven cavity flow where the fluid is initially at rest, with $\rho = 1$ and $p(\mathbf{x}) = 100/\gamma$. We consider a domain of size $1\text{ m} \times 1\text{ m}$ which is surrounded by no-slip walls. The flow is driven entirely by the upper wall which has a velocity of $v_x = 1\text{ m/s}$. The simulation runs for 10 s . Again, our results (Fig. 3) have an excellent agreement with the reference solution of [9]. We used an ADER-DG-method of order 3 with a mesh of size 27^2 .



(a) Our result (markers) of the Taylor-Green vortex vs. the analytical solution (lines) Eq. (13). The plot shows two velocity slices, the respective other coordinate is held constant at a value of π .

(b) Our approximation (markers) of the ABC-flow vs. analytical solution (lines, Eq. (14)). All other axes are held constant at a value of 0. We show every 6th value.

Fig. 2. Two-dimensional CFD scenarios

4.2 Stratified Flow Scenarios

Our main focus is the simulation of stratified flow scenarios. In the following, we present bubble convection scenarios in both two and three dimensions. With the constants

$$\gamma = 1.4, \quad \text{Pr} = 0.71, \quad R = 287.058, \quad p_0 = 10^5 \text{Pa}, \quad g = 9.8 \text{m/s}^2, \quad (15)$$

all following scenarios are described in terms of the potential temperature

$$\theta = T \left(\frac{p_0}{p} \right)^{R/c_p}, \quad (16)$$

with reference pressure p_0 [10,12]. We compute the initial background density and pressure by inserting the assumption of a constant background energy in Eq. (6). The background atmosphere is then perturbed. We set the density and energy at the boundary such that it corresponds to the background atmosphere. Furthermore, to ensure that the atmosphere stays in hydrostatic balance, we need to impose the viscous heat flux

$$\mathbf{F}_{\rho E}^v = \kappa \frac{\partial \bar{T}}{\partial z}. \quad (17)$$

at the boundary [10]. In this equation, $\bar{T}(z)$ is the background temperature at position z , which can be computed from Eqs. (2) and (6).

Our first scenario is the colliding bubbles scenario [12]. We use perturbations of the form

$$\theta' = \begin{cases} A & r \leq a, \\ A \exp\left(-\frac{(r-a)^2}{s^2}\right) & r > a, \end{cases} \quad (18)$$

where s is the decay rate and r is the radius to the center

$$r^2 = \|\mathbf{x} - \mathbf{x}_c\|_2^2, \quad (19)$$

i.e., r denotes the Euclidean distance between the spatial positions $\mathbf{x} = (x, z)$ and the center of a bubble $\mathbf{x}_c = (x_c, z_c)$ – for three-dimensional scenarios \mathbf{x} and \mathbf{x}_c also contain a y coordinate.

We have two bubbles, with constants

$$\begin{array}{llllll} \text{warm:} & A = 0.5 \text{ K}, & a = 150 \text{ m}, & s = 50 \text{ m}, & x_c = 500 \text{ m}, & z_c = 300 \text{ m}, \\ \text{cold:} & A = -0.15 \text{ K}, & a = 0 \text{ m}, & s = 50 \text{ m}, & x_c = 560 \text{ m}, & z_c = 640 \text{ m}. \end{array} \quad (20)$$

Similar to [12], we use a constant viscosity of $\mu = 0.001$ to regularize the solution. Note that we use a different implementation of viscosity than [12]. Hence, it is difficult to compare the parametrization directly. We ran this scenario twice: once without AMR and a mesh of size $1000/81 \text{ m} = 12.35 \text{ m}$ and once with AMR with two adaptive refinement levels and parameters $T_{\text{refine}} = 2.5$ and $T_{\text{coarsen}} = -0.5$.

For both settings we used polynomials of order 6. We specialize the AMR-criterion (Eq. 9) to our stratified flows by using the potential temperature. This resulted in a mesh with cell-size lengths of approx. 111.1 m, 37.04 m, and 12.34 m. The resulting mesh can be seen in Fig. 5.




Fig. 3. Our approximation (solid lines) of the lid-driven cavity flow vs. reference solution (crosses) of [9]. The respective other coordinate is held constant at a value of 0.




Fig. 4. Colliding bubbles with MUSCL-Hancock. Contour values for potential temperature perturbation are $-0.05, 0.05, 0.1, \dots, 0.45$.

We observe that the L_2 difference between the potential temperature of the AMR run, which uses 1953 cells, and the one of the fully refined run with 6561 cells, is only 1.87. The relative error is 2.6×10^{-6} . We further emphasize that our AMR-criterion accurately tracks the position of the edges of the cloud instead of only its position. This is the main advantage of our gradient-based method in contrast to methods working directly with the value of the solution, as for example [12]. Overall, our result for this benchmark shows an excellent agreement to the previous solutions of [12]. In addition, we ran a simulation with the same settings for a polynomial order of 3. The lower resolution leads to spurious waves (Fig. 5) and does not capture the behavior of the cloud.

Furthermore, we simulated the same scenario with our MUSCL-Hancock method, using 7^2 patches with 90^2 finite volume cells each. As we use limiting, we do not need any viscosity. The results of this method (Fig. 4) also agree with the reference but contain fewer details. Note that the numerical dissipativity of the finite volume scheme has a smoothing effect that is similar to the smoothing caused by viscosity.

For our second scenario, the cosine bubble, we use a perturbation of the form

$$\theta' = \begin{cases} A/2 [1 + \cos(\pi r)] & r \leq a, \\ 0 & r > a, \end{cases} \quad (21)$$

where A denotes the maximal perturbation and a is the size of the bubble. We use the constants

$$A = 0.5 \text{ K}, \quad a = 250 \text{ m}, \quad x_c = 500 \text{ m}, \quad z_c = 350 \text{ m}. \quad (22)$$




Fig. 5. Left: Colliding bubbles with ADER-DG. Contour values for potential temperature perturbation are $-0.05, 0.05, 0.1, \dots, 0.45$. Right: Comparison of small scale structure between order 3 (top) and order 6 (bottom).

For the three-dimensional bubble, we set $y_c = x_c = 500\text{m}$. This corresponds to the parameters used in [11]¹. For the 2D case, we use a constant viscosity of $\mu = 0.001$ and an ADER-DG-method of order 6 with two levels of dynamic

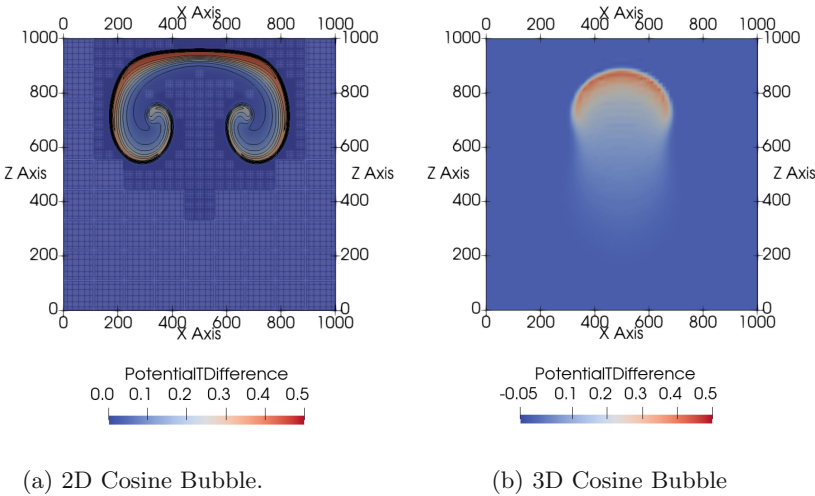


Fig. 6. Cosine bubble scenario.

¹ We found that the parameters presented in the manuscript of [11] only agree with the results, if we use the same parameters as for 2D simulations.

AMR, resulting again in cell sizes of roughly 111.1 m, 37.04 m, 12.34 m. We use slightly different AMR parameters of $T_{\text{refine}} = 1.5$ and $T_{\text{coarsen}} = -0.5$ and let the simulation run for 700 s. Note that, as seen in Fig. 6a, our AMR-criterion tracks the wavefront of the cloud accurately. This result shows an excellent agreement to the ones achieved in [10, 12].

For the 3D case, we use an ADER-DG-scheme of order 3 with a static mesh with cell sizes of 40 m and a shorter simulation duration of 400 s. Due to the relatively coarse resolution and the hence increased aliasing errors, we need to increase the viscosity to $\mu = 0.005$. This corresponds to a larger amount of smoothing. Our results (Fig. 6b) capture the dynamics of the scenario well and agree with the reference solution of [11].

4.3 Scalability

All two-dimensional scenarios presented in this paper can be run on a single workstation in less than two days. Parallel scalability was thus not the primary goal of this paper. Nevertheless, our implementation allows us to scale to small to medium scale setups using a combined MPI + Thread building blocks (TBB) parallelization strategy, which works as follows: We typically first choose a number of MPI ranks that ensure an equal load balancing. ExaHyPE achieves best scalability for 1, 10, 83, ... ranks, as our underlying framework uses three-way splittings for each level and per dimension and an additional communication rank per level. For the desired number of compute nodes, we then determine the number of TBB threads per rank to match the number of total available cores.

We ran the two bubble scenario for a uniform grid with a mesh size of 729×729 with order 6, resulting in roughly 104 million degrees of freedom (DOF), for 20 timesteps and for multiple combinations of MPI ranks and TBB threads. This simulation was performed on the SuperMUC-NG system using computational kernels that are optimized for its Skylake architecture. Using a single MPI rank, we get roughly 4.9 millions DOF updates (MDOF/s) using two TBB threads and 20.2MDOF/s using 24 threads (i.e. a half node). For a full node with 48 threads, we get a performance of 12MDOF/s. When using 5 nodes with 10 MPI ranks, we achieve 29.3MDOF/s for two threads and 137.3MDOF/s for 24 threads.

We further note that for our scenarios weak scaling is more important than strong scaling, as we currently cover only a small area containing a single cloud, where in practical applications one would like to simulate more complex scenarios.

5 Conclusion

We presented an implementation of a MUSCL-Hancock-scheme and an ADER-DG-method with AMR for the Navier-Stokes equations, based on the ExaHyPE Engine. Our implementation is capable of simulating different scenarios: We show that our method has high order convergence and we successfully evaluated our method for standard CFD scenarios: We have competitive results for both

two-dimensional scenarios (Taylor-Green vortex and lid-driven cavity) and for the three-dimensional ABC-flow.

Furthermore, our method allows us to simulate flows in hydrostatic equilibrium correctly, as our results for the cosine and colliding bubble scenarios showed. We showed that our AMR-criterion is able to vastly reduce the number of grid cells while preserving the quality of the results.

Future work should be directed towards improving the scalability. With an improved AMR scaling and some fine tuning of the parallelization strategy, the numerical method presented here might be a good candidate for the simulation of small scale convection processes that lead to cloud formation processes.

Acknowledgments. This work was funded by the European Union's Horizon 2020 Research and Innovation Programme under grant agreements No 671698 (project ExaHyPE, www.exahype.eu) and No 823844 (ChEESE centre of excellence, www.cheese-coe.eu). Computing resources were provided by the Leibniz Supercomputing Centre (project pr83no). Special thanks go to Dominic E. Charrier for his support with the implementation in the ExaHyPE Engine.

References

1. Bauer, P., Thorpe, A., Brunet, G.: The quiet revolution of numerical weather prediction. *Nature* **525**(7567), 47 (2015). <https://doi.org/10.1038/nature14956>
2. Chan, T.F., Golub, G.H., LeVeque, R.J.: Updating formulae and a pairwise algorithm for computing sample variances. In: Caussinus, H., Ettinger, P., Tomassone, R. (eds.) *COMPSTAT 1982*, pp. 30–41. Physica, Heidelberg (1982). https://doi.org/10.1007/978-3-642-51461-6_3
3. Charrier, D.E., Weinzierl, T.: Stop talking to me—a communication-avoiding ADER-DG realisation. *arXiv preprint arXiv:1801.08682* (2018)
4. Dumbser, M.: Arbitrary high order PNPM schemes on unstructured meshes for the compressible Navier-Stokes equations. *Comput. Fluids* **39**(1), 60–76 (2010). <https://doi.org/10.1016/j.compfluid.2009.07.003>
5. Dumbser, M., Balsara, D.S., Toro, E.F., Munz, C.D.: A unified framework for the construction of one-step finite volume and discontinuous Galerkin schemes on unstructured meshes. *J. Comput. Phys.* **227**(18), 8209–8253 (2008). <https://doi.org/10.1016/j.jcp.2008.05.025>
6. Dumbser, M., Peshkov, I., Romenski, E., Zanotti, O.: High order ADER schemes for a unified first order hyperbolic formulation of continuum mechanics: viscous heat-conducting fluids and elastic solids. *J. Comput. Phys.* **314**, 824–862 (2016). <https://doi.org/10.1016/j.jcp.2016.02.015>
7. Fambri, F., Dumbser, M., Zanotti, O.: Space-time adaptive ADER-DG schemes for dissipative flows: compressible navier-stokes and resistive MHD equations. *Comput. Phys. Commun.* **220**, 297–318 (2017). <https://doi.org/10.1016/j.cpc.2017.08.001>
8. Gassner, G., Lörcher, F., Munz, C.D.: A discontinuous Galerkin scheme based on a space-time expansion II. Viscous flow equations in multi dimensions. *J. Sci. Comput.* **34**(3), 260–286 (2008). <https://doi.org/10.1007/s10915-007-9169-1>
9. Ghia, U., Ghia, K., Shin, C.: High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method. *J. Comput. Phys.* **48**(3), 387–411 (1982). [https://doi.org/10.1016/0021-9991\(82\)90058-4](https://doi.org/10.1016/0021-9991(82)90058-4)

10. Giraldo, F., Restelli, M.: A study of spectral element and discontinuous Galerkin methods for the Navier-Stokes equations in nonhydrostatic mesoscale atmospheric modeling: equation sets and test cases. *J. Comput. Phys.* **227**(8), 3849–3877 (2008). <https://doi.org/10.1016/j.jcp.2007.12.009>
11. Kelly, J.F., Giraldo, F.X.: Continuous and discontinuous Galerkin methods for a scalable three-dimensional nonhydrostatic atmospheric model: limited-area mode. *J. Comput. Phys.* **231**(24), 7988–8008 (2012). <https://doi.org/10.1016/j.jcp.2012.04.042>
12. Müller, A., Behrens, J., Giraldo, F.X., Wirth, V.: An adaptive discontinuous Galerkin method for modeling cumulus clouds. In: Fifth European Conference on Computational Fluid Dynamics, ECCOMAS CFD (2010)
13. Müller, A., Kopera, M.A., Marras, S., Wilcox, L.C., Isaac, T., Giraldo, F.X.: Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA. *Int. J. High Perform. Comput. Appl.* **33**(2), 411–426 (2018). <https://doi.org/10.1177/1094342018763966>
14. Tavelli, M., Dumbser, M.: A staggered space-time discontinuous Galerkin method for the three-dimensional incompressible Navier-Stokes equations on unstructured tetrahedral meshes. *J. Comput. Phys.* **319**, 294–323 (2016). <https://doi.org/10.1016/j.jcp.2016.05.009>
15. Van Albada, G., Van Leer, B., Roberts, W.: A comparative study of computational methods in cosmic gas dynamics. In: Hussaini, M.Y., van Leer, B., Van Rosendale, J. (eds.) *Upwind and High-Resolution Schemes*, pp. 95–103. Springer, Heidelberg (1997). https://doi.org/10.1007/978-3-642-60543-7_6
16. Wasserman, L.: *All of Statistics*. Springer, New York (2004). <https://doi.org/10.1007/978-0-387-21736-9>

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 license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license 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.

