# A massively parallel Eulerian-Lagrangian method for advection-dominated transport in viscous fluids

Nils Kohl[*]    Marcus Mohr[†]    Sebastian Eibl[*]    Ulrich Rüde[*‡]

### Abstract

Motivated by challenges in Earth mantle convection, we present a massively parallel implementation of an Eulerian-Lagrangian method for the advection-diffusion equation in the advection-dominated regime. The advection term is treated by a particle-based, characteristics method coupled to a block-structured finite-element framework. Its numerical and computational performance is evaluated in multiple, two- and three-dimensional benchmarks, including curved geometries, discontinuous solutions, pure advection, and it is applied to a coupled non-linear system modeling buoyancy-driven convection in Stokes flow. We demonstrate the parallel performance in a strong and weak scaling experiment, with scalability to up to $147,456$ parallel processes, solving for more than $5.2 \times 10^{10}$ (52 billion) degrees of freedom per time-step.

**Key words**    Eulerian–Lagrangian methods, advection-diffusion, parallel algorithms

**AMS subject classifications**    65M25, 65Y05, 65M60

## 1    Introduction

While to us as human beings the ground on which we walk may appear 'rock-solid' the surface of our planet is actually in constant albeit very slow motion. Continental plates move at a rate of centimetres per year. The reason for this movement are enormous forces acting deep below our feet. Convective processes in the Earth's mantle help the planet rid itself of excess energy that is either left from the time of its formation or generated by continued radioactive decay. The mantle is a layer of Earth starting from below the crust at roughly 60 km and extending down to the core-mantle-boundary at a depth of about 3,000 km. On geologic time-scales the rocks inside the mantle behave like a highly viscous fluid. A single overturn of the material in the mantle takes about 100 mio. years.

A detailed understanding of these processes is of fundamental interest to geophysics, as they are the driving force behind phenomena such as plate tectonics, mountain and ocean building, volcanism, and finally earthquakes. As the mantle is not accessible for direct measurements studies of its convection rely mostly on simulation and form an active research topic in computational fluid dynamics (CFD). The requirements on spatial and temporal resolution render the solution of the underlying system of partial differential equations (PDEs) a grand challenge in computational science [2, 11].

[*]Computer Science 10, Friedrich-Alexander-Universität Erlangen-Nürnberg (`nils.kohl@fau.de`, `ulrich.ruede@fau.de`, `sebastian.eibl@fau.de`).

[†]Dept. of Earth and Environmental Sciences, LMU Munich (`marcus.mohr@lmu.de`).

[‡]Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France

The combination of extremely viscous material, characteristic length scale, and creeping flow of the Earth's mantle result in a Reynolds number on the order of $10^{-15}$, [40] and the Stokes equations are suitable to model momentum and mass balance. Conservation of energy can be described by an equation of advection–diffusion type for the temperature. In a buoyancy-driven flow the dimensionless Rayleigh number Ra describes the vigor of convection. For the Earth's mantle Ra is estimated to lie between $10^7$ and $10^8$ [40]. In that range temperature transport is mainly driven by fluid flow (advection) and much less by diffusive effects.

In this paper we are interested in the numerical treatment of this kind of equation in the advection-dominated regime. While the temperature equation of mantle convection forms our focus point, such kind of transport problems appear, of course, also in many other applications in CFD [15, 36]. Although the quantity of interest varies, the main characteristics of the underlying equation remain the same. Typical transported variables include for example chemical species concentration, material markers, or isotope ratios.

The solution of the advection-diffusion equation is known to be challenging in the advection-dominated regime, for instance due to stability issues at high gradients or even discontinuities in the solution [23, 38]. Well-known and established methods for the numerical treatment of advection-diffusion equations include the streamline upwind Petrov-Galerkin (SUPG) method [10], where for stability reasons, artificial diffusion is introduced into the solution. A more recent approach in the same direction is the entropy viscosity method, see e.g. [31] and references therein. Algebraic flux correction (AFC) approaches the problem by modification of the equations at the algebraic level [33]. A comparison of SUPG, AFC and other finite-element based methods for advection-dominated transport is presented in [28]. High-order, discontinuous Galerkin discretizations [16,39] are attractive as they are naturally well-suited to represent discontinuous solutions. However, the selection of adequate slope-limiters and the large number of unknowns that are introduced may be problematic.

A fundamentally different approach to the discretization of advection-diffusion equations are so-called *Lagrangian* or *characteristic* methods. Instead of employing a fixed, *Eulerian* grid, the advected property is captured by particles or volumes that move along the characteristics of the velocity field. Usually, both, Eulerian and Lagrangian discretization approaches are combined by means of a splitting-technique, where the advective term is treated by a Lagrangian, and the diffusive term by an Eulerian discretization. Solutions need to be interpolated between these two domains. These approaches are also called *Eulerian-Lagrangian methods* (ELMs). Two prominent implementations of this category are the *modified method of characteristics* (MMOC) [1, 19, 22, 35] (also referred to as characteristic Galerkin method or Lagrange-Galerkin method) and the Eulerian-Lagrangian localized adjoint method (ELLAM) [14, 42].

The MMOC is based on the backtracking of particles along the characteristics, where the transported quantity for the next time step is evaluated. This method permits large time steps, is free from parameterization and conceptually easy to understand. The particle-based method requires frequent evaluation (or interpolation) of the solution function away from the grid nodes. In general, the MMOC is not perfectly energy-conserving. A scheme to enforce global energy conservation is developed in [18]. Numerical analysis on accuracy and stability of the MMOC is found in [7,17]. Note that by following characteristics backwards in time, MMOC is conceptually different from the particle/marker-in-cell techniques often employed in geodynamical flow simulations for advecting quantities like chemical composition or water content, [24]. It also avoids some of their pitfalls such as e.g. the question of particle concentration per cell. The only investigation of MMOC-based methods for geodynamical flows seems to be [35].

ELLAM may provide local energy conservation by propagation of volumes instead of particles. This class of methods has similar advantages as the MMOC, but the inte-

gration over elements that are not aligned with the grid may be difficult, in particular in parallel implementations, and thus it can be computationally expensive.

In this article, our focus is on an ELM based on the MMOC suited for massively parallel simulations on state-of-the-art supercomputers. The parallel algorithms and data structures used in our implementation build upon the concept of hierarchical hybrid grids (HHG) [2, 6], addressing extreme-scalable, matrix-free geometric multi-grid solvers on block-structured grids. With mantle convection models as a target application, a prototype application has demonstrated scalability of Stokes solvers for systems with more than $10^{13}$ unknowns [26]. New matrix-free methods [3, 5], performance and scalability [25, 26, 29], and application to geophysical problems [4] have been studied, mainly focusing on the solution of the Stokes system. The ELM proposed in this article is developed to exploit and extend the excellent scalability of the HHG-based solvers for time-dependent mantle-convection problems.

Parallel implementations of ELMs have been designed for various applications, including research on sea-ice [43], Navier-Stokes [37, 45], and also natural convection in [12]. In the latter a target application similar to this work is considered on un-structured meshes, and an ELM is used for both the advection terms in the energy equation and also the discretization of the Navier-Stokes system itself. However, only moderate scalability with up to 1,000 parallel processes was demonstrated. To quantitatively and accurately predict the convection patterns of Earth's mantle, however, extreme-scale parallel simulations are necessary, as for instance a global spatial resolution of $\sim 1.7$km results in linear systems with more than a trillion ($10^{12}$) degrees of freedom (DoFs) [4]. Such problems require methods that can efficiently exploit the resources of today's peta- and future exascale supercomputers. With the proposed method, we demonstrate the scalability of ELM-based time-dependent simulations for up to a hundred of thousand parallel processors.

**Contribution** In this paper we will

(a) present a particle-based, massively parallel method for the advection-diffusion equation based on the MMOC that is applicable to curved geometries and largely independent of the underlying grid data structures and spatial discretization,

(b) embed the method into to a block-structured finite-element framework based on HHG,

(c) quantify the accuracy and energy conservation of our approach through multiple, two- and three-dimensional benchmarks with different spatial finite-element discretizations, discontinuous solutions, pure advection, curved domains, large time steps, Courant-Friedrichs-Lewy (CFL) number $> 1$, and coupled buoyancy-driven flow, and

(d) demonstrate the extreme-scalability of the approach on to up to 147,456 parallel processes and more than $5.2 \cdot 10^{10}$ particles, and an application to a simplified mantle convection setup.

**Reproducibility** All presented algorithms and benchmarks are implemented in the open-source software framework Hybrid Tetrahedral Grids (HyTeG)[1] [29, 30, 46], assuring reproducibility of the results.

## Governing equations

We consider the numerical approximation of the advection-diffusion equation on a bounded domain $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$, and time interval $[0, T], T \in \mathbb{R}^+$

$$\frac{\partial}{\partial t} c + \mathbf{u} \cdot \nabla c - \kappa \Delta c = q, \quad (\mathbf{x}, t) \in \Omega \times [0, T] \tag{1}$$

---

[1] `https://i10git.cs.fau.de/hyteg/hyteg`

where $c = c(\mathbf{x}, t)$ represents the advected, scalar quantity (temperature in case of our target application), $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ a given divergence-free velocity field, i.e. satisfying

$$\nabla \cdot \mathbf{u} = 0, \quad (\mathbf{x}, t) \in \Omega \times [0, T] \ ,$$

$q = q(\mathbf{x}, t)$ the given rate of internal heat production, and $\kappa \geq 0$ a diffusivity parameter. Initial, Dirichlet, and (homogeneous) Neumann boundary conditions for the temperature $c$ are given by

$$c(\mathbf{x}, 0) = c_0(\mathbf{x}), \ \mathbf{x} \in \Omega, \quad c(\mathbf{x}, t) = c_\Gamma(\mathbf{x}, t), \ \mathbf{x} \in \partial\Omega_D, \quad \frac{\partial c}{\partial \mathbf{n}}(\mathbf{x}, t) = 0, \ \mathbf{x} \in \partial\Omega_N$$

for $t \in [0, T]$, boundary $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N$, and outward normal $\mathbf{n}$. We require for the sake of simplicity that the velocity field has no inflow into the domain.

In typical applications, the advective term $\mathbf{u} \cdot \nabla c$ strongly dominates over the diffusive term $\kappa\Delta c$. Depending on the formulation and non-dimensionalization of the model, this translates to either $\kappa \ll 1$, or large velocity magnitudes.

The advection-diffusion equation can be coupled to the Stokes equation for viscous flows using the Boussinesq-approximation for natural convection, as will be described in section 5.

# 2    Eulerian-Lagrangian method

In this section we describe the parallel algorithms and data structures of the MMOC-based method for the advection-diffusion equation (1).

## 2.1    Hierarchical hybrid grids

We base the construction of the computational mesh on the concept of HHG [2, 6]. Therefore, we define a coarse unstructured mesh $\mathcal{T}_0$ of tetrahedral (or triangular) elements that partitions the domain $\Omega$. In a second step, each coarse grid element is uniformly refined according to [8]. This results in hierarchy of block-structured meshes $\mathcal{T} = \{\mathcal{T}_\ell, \ell = 0, ..., L\}$ and offers crucial performance advantages for matrix-free multigrid methods as demonstrated especially for the Stokes system [2, 3, 5, 29].

If the problem domain $\Omega$ is polyhedral, we can define a set of coarse grid elements, whose union equals $\Omega$. However, in this article we also consider a more general case, which is that $\Omega$ coincides with a polyhedral domain after a *blending function* $\Phi$ is applied to the latter. In particular, we are interested in domains with curved boundaries, such as the thick spherical shell, that is used to represent Earth's mantle in geophysical models [4, 41]. We require $\Phi$ to be a homeomorphism and its inverse to be known explicitly.

To construct the grid hierarchy for this second case, we start from an approximation of the *physical domain* $\Omega_{\mathrm{phy}} := \Omega$ by a polyhedral, *computational domain* $\Omega_{\mathrm{comp}}$ (i.e. $\Phi(\Omega_{\mathrm{comp}}) = \Omega_{\mathrm{phy}}$). This polyhedral domain is then refined as outlined above, yielding a mesh hierarchy $\mathcal{T} = \{\mathcal{T}_\ell, \ell = 0, ..., L\}$. Finally, by applying our blending function to each mesh $\mathcal{T}_\ell$ we obtain a hierarchy $\widetilde{\mathcal{T}} := \{\Phi(\mathcal{T}_\ell), \ell = 0, \ldots, L\}$ for $\Omega_{\mathrm{phy}}$. Obviously, application of this algorithm to a polyhedral physical domain $\Omega_{\mathrm{phy}}$ corresponds to the special case $\Phi = \mathrm{Id}$ as $\Omega_{\mathrm{phy}} = \Omega_{\mathrm{comp}}$. Figure 1 shows an example, where the computational domain is projected onto an annulus. The left figure shows an initial, unrefined, unstructured computational mesh $\mathcal{T}_0$, the right figure the corresponding physical mesh $\Phi(\mathcal{T}_3)$ after three refinement iterations.

Efficient and scalable, matrix-free solvers for scalar elliptic PDE problems and Stokes flow on curved domains in conjunction with HHG have been presented in [3, 5].

(a) $\mathcal{T}_0$            (b) $\Phi(\mathcal{T}_3)$

Figure 1: *Partitioning of an annular domain: (a) unstructured, initial computational mesh before refinement, (b) refined mesh mapped to physical domain.*

## 2.2 Discretization of the advection-diffusion equation

The essence of the MMOC is the elimination of the advective term $\mathbf{u} \cdot \nabla c$ from (1). For this, we define the so-called *characteristics* $\mathbf{X} : \Omega \times [0,T]^2 \to \mathbb{R}^3$ of the velocity field $\mathbf{u}$ as the solutions of

$$\frac{d}{dt}\mathbf{X}(\mathbf{x},s,t) = \mathbf{u}(\mathbf{X}(\mathbf{x},s,t),t), \quad t \in (0,T)$$
$$\mathbf{X}(\mathbf{x},s,s) = \mathbf{x} \tag{2}$$

for fixed $(\mathbf{x},s) \in \Omega \times [0,T]$. Specifying two points in time $t_0, t_1 \in [0,T]$, $t_0 < t_1$, $\mathbf{X}(\mathbf{x},t_1,t_0)$ can be interpreted as the *departure point* at time $t_0$ of a particle, that reaches point $\mathbf{x}$ at time $t_1$. Such a departure point is, thus, given by

$$\mathbf{X}(\mathbf{x},t_1,t_0) = \mathbf{x} - \int_{t_0}^{t_1} \mathbf{u}(\mathbf{X}(\mathbf{x},t_1,t),t)\,dt. \tag{3}$$

We now define, for a fixed time $s \in [0,T]$

$$\hat{c}(\mathbf{x},t) := c(\mathbf{X}(\mathbf{x},s,t),t)$$

and calculate, using the chain rule and (2)

$$\frac{\partial}{\partial t}\hat{c}(\mathbf{x},t) = \left(\frac{\partial}{\partial t}c + \mathbf{u} \cdot \nabla c\right)(\mathbf{X}(\mathbf{x},s,t),t). \tag{4}$$

At time $t = s$ we can replace the advective term in (1), since

$$\frac{\partial}{\partial t}\hat{c}(\mathbf{x},s) = \left(\frac{\partial}{\partial t}c + \mathbf{u} \cdot \nabla c\right)(\mathbf{x},s),$$

and reformulate the PDE as

$$\frac{\partial}{\partial t}\hat{c} - \kappa \Delta c = q. \tag{5}$$

Next we semi-discretize (5) in time. To this end, we divide $[0,T]$ into $N$ intervals $[t_n, t_{n+1}]$, $n \in \{0, \ldots, N-1\}$ with step size $\tau_n = t_{n+1} - t_n$. We then set $\mathbf{x} = \mathbf{X}(\mathbf{x},t_{n+1},t_{n+1})$ (or $s = t_{n+1}$ in (4)) and approximate the time derivative via a difference quotient

$$\frac{\partial}{\partial t}\hat{c}(\mathbf{x},t_{n+1}) \approx \frac{1}{\tau_n}\left[\hat{c}(\mathbf{x},t_{n+1}) - \hat{c}(\mathbf{x},t_n)\right] = \frac{1}{\tau_n}\left[c(\mathbf{x},t_{n+1}) - c(\mathbf{X}(\mathbf{x},t_{n+1},t_n),t_n)\right].$$

We perform the spatial discretization of the temperature and velocity fields using the standard Galerkin finite element method subject to the HHG grid hierarchy described in section 2.1. We therefore introduce the spaces of piecewise polynomial functions

$$\mathcal{S}_\ell^m := \{v \in \mathcal{C}^0(\Omega) : v|_T \in \mathcal{P}_m(T), \; \forall \, T \in \mathcal{T}_\ell\}, \quad l \in \{0, ..., L\}, \; m \geq 1.$$

Here, $\mathcal{P}_m(T)$ denotes the space of polynomials of degree $m$ on the element $T$. Let $V_h := \mathcal{S}_L^m \cap \mathcal{H}_0^1(\Omega)$ be a finite dimensional subspace of $\mathcal{H}_0^1(\Omega)$ with piecewise polynomial basis functions that vanish on the boundary. In particular, we employ the standard sets of Lagrange basis functions $P_m$ for polynomial degree $m$ [23]. Furthermore, given a function $c_\Gamma := c_\Gamma(\mathbf{x}, t)$ that defines suitable Dirichlet boundary conditions, let $V_h^D := \mathcal{S}_L^m \cap \mathcal{H}_D^1(\Omega)$ with $\mathcal{H}_D^1 := \{v_h \in \mathcal{H}^1(\Omega) : v_h = c_\Gamma \text{ on } \partial\Omega_D\}$.

We apply the $\Theta$-method to the time-discretization of the diffusive term [38]. The finite dimensional version of the weak formulation of (5) then reads: given $\hat{c}_h^n = \hat{c}_h^n(\mathbf{x}) \in V_h^D$, find $c_h^{n+1} = c_h^{n+1}(\mathbf{x}) \in V_h^D$ so that

$$\frac{1}{\tau_n}(c_h^{n+1} - \hat{c}_h^n, v_h) + \Theta\kappa(\nabla c_h^{n+1}, \nabla v_h) + (1 - \Theta)\kappa(\nabla\hat{c}_h^n, \nabla v_h)$$
$$= (\Theta q(t_{n+1}) + (1 - \Theta)q(t_n), v_h), \quad \text{for all } v_h \in V_h \tag{6}$$

and $c_h^0 = c_{0,h}$. $(\cdot, \cdot)$ denotes the inner product in $L^2(\Omega)$ and $\Theta \in [0, 1]$. This corresponds to an implicit Euler or Crank-Nicolson scheme for the diffusive term, for $\Theta = 1$ or $\Theta = 0.5$, respectively. For the formulation of the bilinear and linear forms in the case of a blended domain, i.e. $\Phi \neq \text{Id}$, we refer to [3, 27].

Associating $c_h^{n+1}$, $\hat{c}_h^n$, $(q(t_n), v_h)$, and $(q(t_{n+1}), v_h)$ with coefficient vectors $\underline{\mathbf{c}}^{n+1}$, $\underline{\hat{\mathbf{c}}}^n$, $\underline{\mathbf{q}}^n$, and $\underline{\mathbf{q}}^{n+1}$ we formulate (6) as the linear system

$$(M + \tau_n\Theta\kappa A)\underline{\mathbf{c}}^{n+1} = (M - \tau_n(1 - \Theta)\kappa A)\underline{\hat{\mathbf{c}}}^n + \tau_n(\Theta\underline{\mathbf{q}}^{n+1} + (1 - \Theta)\underline{\mathbf{q}}^n) \tag{7}$$

that has to be solved in each time step. $M$ represents the finite element mass matrix, and $A$ the stiffness matrix. The matrix $E := (M + \tau_n\Theta\kappa A)$ is symmetric and positive definite. This allows for efficient inversion. Especially for small time steps, $E$ tends to be more diagonally dominant than the stiffness matrix $A$ and is therefore well suited for treatment with conjugate gradient and multigrid solvers [47].

It remains to determine an approximation for $\underline{\hat{\mathbf{c}}}^n$, which requires the evaluation of $\hat{c}_h^n(\mathbf{x}) = c_h^n(\mathbf{X}(\mathbf{x}, t_{n+1}, t_n))$. The advected temperature is obtained by calculation of the departure point $\mathbf{X}(\mathbf{x}, t_{n+1}, t_n)$ via the integral in (3). Due to the initial condition and the continuous Galerkin discretization, $c_h^n(\mathbf{x})$ can be evaluated for all $\mathbf{x} \in \Omega$.

In general, the integral in (3) cannot be evaluated analytically but has to be approximated numerically. Here, we apply standard, explicit Runge-Kutta (RK) schemes that repeatedly evaluate the velocity field $\mathbf{u}$. For the general case of time-dependent and time-discrete velocity fields, evaluation at time $t^* \in (t_n, t_{n+1})$ requires interpolation. In this case, we employ linear interpolation in time. Spatially, we represent the velocity field $\mathbf{u}$ also in one of the continuous finite element spaces $\mathcal{S}_\ell^m$ resulting in a well-defined approximation $\mathbf{u}_h$. Details on the numerical integration and evaluation are presented in section 2.3.

Algorithm 1 summarizes the time-stepping scheme for the advection-diffusion equation. To determine a suitable time-step size, we employ a CFL condition via a constant $\text{CFL}_{\max}$, the length of the shortest edge of the mesh $h_{\min}$, and the maximum velocity magnitude at time-step $n$, i.e. $\max_{\mathbf{x}\in\Omega} |\mathbf{u}_h(\mathbf{x}, t_n)|$.

## 2.3 Parallel implementation

In this section, we describe the parallel implementation of the MMOC on HHG. In particular, we discuss the execution of the Lagrangian step, i.e. the calculation of $\hat{c}_h^n$, and the implementation in the HyTeG finite element framework. This corresponds to lines 3, and 4 in algorithm 1.

---
**Algorithm 1** Time-stepping scheme, advection-diffusion.
---
1: **procedure** $\mathrm{AD}(c_h^n, \mathbf{u}_h)$
2:      $\tau_n = \mathrm{CFL}_{\max} \cdot h_{\min} / \max_{\mathbf{x} \in \Omega} |\mathbf{u}_h(\mathbf{x}, t_n)|$          ▷ determine time-step size
3:      $\hat{\mathbf{x}} = \mathbf{X}(\mathbf{x}, t_{n+1}, t_n)$          ▷ calculate departure points (see section 2.3)
4:      $\hat{c}_h^n(\mathbf{x}) = c_h^n(\hat{\mathbf{x}})$          ▷ advection
5:      solve (7) to advance from $\hat{c}_h^n$ to $c_h^{n+1}$          ▷ diffusion
6:      **return** $c_h^{n+1}$
---

### 2.3.1 Particle tracing

We employ *tracer particles* that are created at the DoFs of $c_h$ at time $t_{n+1}$ and are transported backwards along the velocity trajectories, until they reach the departure points at time $t_n$. Usually, for standard Lagrange finite element discretizations, the DoFs are set to coincide with the grid vertices for a $\mathbb{P}_1$ discretization, and with the vertices and edge-midpoints for a quadratic $\mathbb{P}_2$ discretization. However, the method is not restricted to such a choice, and discretizations with a different DoF-layout such as finite-volumes may also be realized. The values of $\hat{c}_h^n$ at the DoFs are then determined by evaluation of $c_h^n$ at the departure points.

Given the continuous Galerkin approximation $c_h^n$ of $c$ on the HHG structure, we split the approximation of $\hat{c}_h^n$ into three steps: (i) particle creation, (ii) particle integration and (iii) temperature evaluation. It follows a discussion of the grid and particle data structures, and steps (i) – (iii).

**Grid data structure** For each element of the unstructured coarse grid, a *macro-primitive* (macro-faces in 2D, macro-cells in 3D) data structure is created. The macro-primitives are then uniformly refined. The HHG concept introduces *interface primitives* for each interface between two coarse grid elements. The interface primitives are also refined uniformly. As an example, in 2D, two neighboring *macro-face* primitives are interfaced by a *macro-edge* primitive, and two adjacent macro-edges are interfaced by a *macro-vertex*. This allows for a unique assignment of each individual DoF to a single primitive data structure. Each primitive is assigned a globally unique ID, and in a parallel setting, assigned to one of the parallel processes. For distributed memory architectures, communication is implemented via MPI. The coarse grid and all mesh-related metadata are distributed without global data structures, allowing for parallel runs on hundreds of thousands of parallel processes [26, 29]. More details on the HHG data structures can be found in [6, 29, 30].

**Particle data structure and synchronization** The tracer particles are realized by the Modular and Extensible Software Architecture for Particle Dynamics (MESA-PD) [20, 21], which implements particle data structures for massively parallel particle simulations. It allows to equip each particle with arbitrary properties, that are transported together with the particle through a distributed domain. The individual subdomains correspond to the volume primitives defined by the unstructured coarse grid. Particles that leave the subdomain of a process are communicated via MPI. Similar to the HHG structure, the parallel particle data structures are distributed to allow for massively parallel simulations by design.

After the position of a particle is updated, a synchronization step follows, that assigns each particle uniquely to a single neighboring volume primitive. The target primitive is determined only by the previous owner process of the particle, and therefore prevents race conditions. Detailed information on the parallel data structures and communication are found in [20].

**Step (i): particle creation** For each DoF of the Eulerian grid, a particle is created. The particles are initialized with the corresponding macro-primitive ID, DoF-index,

*Figure 2: Illustration of the parallel particle integration and temperature evaluation (steps (ii) and (iii)) on two neighboring volume primitives. In this example, a 2-stage RK method is employed: (a) initial particle position, evaluation of $\tilde{\mathbf{u}}_h^1(\mathbf{y}_i^1)$ (tangent to velocity field at that point, illustrated by dotted blue line), (b) setting particle position to $\mathbf{y}_i^2$, (c) evaluation of $\tilde{\mathbf{u}}_h^2(\mathbf{y}_i^2)$, (d) calculation of $\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n)$ using $\tilde{\mathbf{u}}_h^1(\mathbf{y}_i^1)$ and $\tilde{\mathbf{u}}_h^2(\mathbf{y}_i^2)$ according to the RK method, and particle communication to the neighboring volume primitive, (e) evaluation of $c_h^n(\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n))$, (f) communication of $c_h^n(\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n))$ back to initial DoF.*

and the process ID, so that their corresponding DoF can be backtracked in a distributed setting. Particles are also initialized on interface primitives, since they are responsible for DoFs at the interfaces of the volume primitives. The initial position of a particle corresponds to $\mathbf{x}_i = \mathbf{X}(\mathbf{x}_i, t_{n+1}, t_{n+1})$, where $\mathbf{x}_i$ is the location of a DoF with index $i$ on $\Omega_{\text{phy}}$. A following synchronization step assigns all particles that were created on an interface primitive to a single volume primitive. It is of no particular importance which volume primitive is chosen.

**Step (ii): particle integration** This step performs the backward transport along the velocity field using an explicit RK integrator with $S$ stages. This corresponds to the computation of $\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n)$ according to (3) using numerical integration. The RK integration requires the evaluation of the velocity field at a time $\tilde{t}^s \in [t_n, t_{n+1}]$ and position $\mathbf{y}_i^s$, with $\mathbf{y}_i^1 = \mathbf{x}_i$ in each stage $s \in [1, ..., S]$.

Before each RK stage, the position of a particle is set to the position $\mathbf{y}_i^s$ where the velocity field needs to be evaluated (see fig. 2 step (b)). Immediately after that, a synchronization step follows so that all particles are available on the process that owns the volume-primitive that contains $\mathbf{y}_i^s$.

We assume that the velocity field is known for the discrete time steps $t_n$ and $t_{n+1}$. Both fields $\mathbf{u}_h^n$ and $\mathbf{u}_h^{n+1}$ are evaluated and we perform linear interpolation. This means we approximate

$$\mathbf{u}(\mathbf{y}_i^s, \tilde{t}^s) \approx \tilde{\mathbf{u}}_h^s(\mathbf{y}_i^s) := \left( \frac{t_{n+1} - \tilde{t}^s}{t_{n+1} - t_n} \right) \mathbf{u}_h^n(\mathbf{y}_i^s) + \left( \frac{\tilde{t}^s - t_n}{t_{n+1} - t_n} \right) \mathbf{u}_h^{n+1}(\mathbf{y}_i^s) \qquad (8)$$

(see fig. 2 steps (a) and (c)). For scenarios where the velocity depends on the temperature field, we refer to section 5 where we discuss buoyancy-driven flows.

The intermediate result $\tilde{\mathbf{u}}_h^s(\mathbf{y}_i^s)$ is stored in the particle data structure before the next stage is executed. After the last stage, all intermediate results and weights of the RK method are combined to calculate the actual final position $\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n)$ (see fig. 2 step (d)).

**Step (iii): temperature evaluation** In this last step, the temperature field $c_h^n$ is evaluated at $\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n)$ (see fig. 2 step (e)). This gives $\hat{c}_h^n(\mathbf{x}_i)$ at the initial position $\mathbf{x}_i$ of the particle. Since the initial position was a DoF, we set the corresponding

coefficient $\hat{\underline{\mathbf{c}}}_i^n = c_h^n(\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n))$. If $\mathbf{x}_i$ is located on a different volume primitive than $\mathbf{X}(\mathbf{x}_i, t_{n+1}, t_n)$, $\hat{\underline{\mathbf{c}}}_i^n$ is communicated (see fig. 2 step (f)).

### 2.3.2 Field evaluation

The evaluation of $\mathbf{u}_h^n(\mathbf{z})$ or $c_h^n(\mathbf{z})$, $\mathbf{z} \in \Omega$ involves localization of the underlying geometric element, and computing a sum of the shape functions evaluated at $\mathbf{z}$ weighted by the corresponding DoF values. In general, as described in section 2.1, $\Omega = \Omega_{\mathrm{phy}}$ may be non-polyhedral, i. e. $\Phi \neq \mathrm{Id}$. We therefore map $\mathbf{z}$ to the computational domain $\Omega_{\mathrm{comp}}$ and set $\mathbf{z}_{\mathrm{comp}} := \Phi^{-1}(\mathbf{z})$. Since we require $\Phi$ to be a homeomorphism, we know that $\mathbf{z}_{\mathrm{comp}} \in T \subset \Omega_{\mathrm{comp}} \Leftrightarrow \mathbf{z} \in \Phi(T) \subset \Omega_{\mathrm{phy}}$. We split the search-locate algorithm on the computational domain into two steps. In a first step, the enclosing volume-primitive that contains $\mathbf{z}_{\mathrm{comp}}$ is determined by searching in the direct neighborhood of the volume-primitive that previously contained the corresponding particle. Then, we search for the containing element $T \subset \Omega_{\mathrm{comp}}$ of $\mathbf{z}_{\mathrm{comp}}$ in the uniformly refined volume primitive. Since we employ block-structured HHG, the element $T$ is found in $\mathcal{O}(1)$ cost. Finally, the value of the finite element function is computed as is standard, by application of a pull-back mapping of $T$ to the reference element.

### 2.3.3 Look-back Distance

The field evaluation in step (iii) implicitly corresponds to an interpolation of the advected temperature field $\hat{c}_h^n$ into the space $\mathcal{S}_\ell^m$. While the discretized original field at time $t_n$ satisfies $c_h^n \in \mathcal{S}_\ell^m$ this will typically not be the case for $\hat{c}_h^n$. Consequently this step introduces an interpolation error. If the field $\hat{c}_h^n$ used to update $c_h^{n+1}$ is computed from $c_h^n$, the latter already involves $n$ previous interpolations, whose errors might accumulate.

However, in the purely advective case ($\kappa = 0$ and $q = 0$) this issue can be diminished or even completely removed. To do so, one can simple follow the particle trajectory back in time over more than only one temperature time step $\tau$, i.e. instead of integrating from $t_{n+1}$ back to $t_n$ we select an earlier time $t_{n+1-b}$. We will refer to the integer $b$ as *look-back distance* as $t_{n+1-b}$ will be the time when temperature is evaluated. Of course, this approach requires that the temperature field is still known at $t_{n+1-b}$, as must be the intermediate velocity fields required by the ODE solver.

By selecting $b = n + 1$ one can derive $c_h^{n+1}$ from the initial temperature $c_{0,h}$ itself. However, the look-back distance then grows with the simulation, a fact that we will mark by using the notation $b = \infty$. This extreme approach preserves the accurate representation of the initial temperature in the Lagrangian domain, and leads to very accurate solutions, as we will see in the following benchmarks. There we will employ different look-back distances $b$ to demonstrate that the interpolation between the Lagrangian and Eulerian representation is the primary source of approximation error. A similar discussion of the accumulation of the interpolation error is found in [35].

## 3 Numerical verification

In the following subsections we assess the accuracy of our implementation through numerical benchmarks.

### 3.1 Test setup

In all benchmarks, we employ either linear ($\mathbb{P}_1$) or quadratic ($\mathbb{P}_2$) Lagrangian finite element discretizations for the temperature and velocity, block-structured triangular and tetrahedral meshes for two- and three-dimensional domains respectively. For the

particle integration we use the standard fourth-order RK integrator (often referred to as RK4). We note that the implementation supports any explicit RK integrator.

To asses the quality of our scheme, we employ the following norms and metrics: let $\tilde{c}_h$ the computed solution, $c_h$ the point-wise interpolated exact solution and $e_h = c_h - \tilde{c}_h$. The corresponding coefficient vectors are denoted as $\underline{\mathbf{c}}$, $\underline{\tilde{\mathbf{c}}}$, and $\underline{\mathbf{e}}$. A discrete version of the $\mathcal{H}^0$-norm of the error is then defined as

$$\|e_h\|_{\mathcal{H}^0} := \left(\underline{\mathbf{e}}^\top M \underline{\mathbf{e}}\right)^{\frac{1}{2}}$$

where $M$ is the finite element mass matrix. We define $\mathrm{var}(t_n)$ as in [28], and $E_{\mathrm{peak}}(t_n)$ as

$$\mathrm{var}(t_n) := \max_j \underline{\tilde{\mathbf{c}}}_j^n - \min_j \underline{\tilde{\mathbf{c}}}_j^n, \quad E_{\mathrm{peak}}(t_n) := \frac{\max_j \underline{\tilde{\mathbf{c}}}_j^n}{\max_j \underline{\mathbf{c}}_j^n} - 1$$

to indicate the amount of spurious oscillations, and to detect if peaks of the solution are preserved. To quantify the energy conservation of our implementation, we indicate a relative energy difference $\Delta m(t_n)$ compared to the initial solution by

$$\Delta m(t_n) := \frac{m(t_n)}{m(t_0)} - 1, \quad m(t_n) := \underline{\mathbf{1}}^\top M \underline{\tilde{\mathbf{c}}}^n, \quad \underline{\mathbf{1}} := (1, \dots, 1)^\top.$$

## 3.2 Circular advection

First we consider a two dimensional body rotation problem as employed in [28,34,51]. In particular, the setup is the same as in [28] to compare the numerical results.

Let $\Omega = (0,1)^2$ be the domain where the initial temperature is imposed by three bodies as shown in fig. 3a. All bodies are defined on circles with radius $r_0 = 0.15$, the initial condition is zero outside of these circles. We define $\mathbf{x} = (x_1, x_2)$, $\bar{\mathbf{x}} = (\bar{x}_1, \bar{x}_2)$, $r(\mathbf{x}) := \|\mathbf{x} - \bar{\mathbf{x}}\|_2 / r_0$, and the initial condition $c_0 = c_0^{\mathrm{slotted}} + c_0^{\mathrm{cone}} + c_0^{\mathrm{hill}}$ by

$$c_0^{\mathrm{slotted}}(\mathbf{x}) = \begin{cases} 1 & \text{if } r(\mathbf{x}) \le 1, \ |\mathbf{x} - \bar{\mathbf{x}}| \ge 0.025, \ y \ge 0.85 \\ 0 & \text{otherwise} \end{cases} \qquad \bar{\mathbf{x}} = (0.5, 0.75),$$

$$c_0^{\mathrm{cone}}(\mathbf{x}) = \begin{cases} 1 - r(\mathbf{x}) & \text{if } r(\mathbf{x}) \le 1 \\ 0 & \text{otherwise} \end{cases} \qquad \bar{\mathbf{x}} = (0.5, 0.25),$$

$$c_0^{\mathrm{hill}}(\mathbf{x}) = \begin{cases} \frac{1}{4}(1 + \cos(\pi r(\mathbf{x}))) & \text{if } r(\mathbf{x}) \le 1 \\ 0 & \text{otherwise} \end{cases} \qquad \bar{\mathbf{x}} = (0.25, 0.5).$$

The bodies are rotating counter-clockwise along the constant velocity field $\mathbf{u} = (0.5 - x_2, x_1 - 0.5)^\top$. Since we consider pure advection ($\kappa = 0$, $q = 0$), at $t = 2\pi$, the bodies have finished a full revolution and the resulting temperature field should be equal to the initial condition. The time step size $\tau$ is constant.

In table 1, the different versions of the MMOC are compared to the linear and non-linear flux-corrected transport (FCT) methods that performed best in [28]. We observe the strong influence of the look-back distance $b$ on the solution, as visualized in the plots of the computed solutions in fig. 3.

Optimal results are achieved with infinite look-back distance ($b = \infty$). This suggests that the interpolation between Eulerian and Lagrangian representation is the primary source of errors, and energy difference. We note, that in a massively parallel setting, occasional interpolation to the Eulerian domain may be desired to reduce the communication overhead during the temperature evaluation (step (iii) in section 2.3.1).

**Remark 3.2.1 (Choice of space-discretization, oscillations)** *The amount of spurious oscillations denoted by* $\mathrm{var}(2\pi)$ *in table 1 with* $\mathbb{P}_2$ *elements and* $b = \infty$ *is partly*

Table 1: Comparison of different parameterizations of the MMOC to the best performing methods from the study in [28]. The FEM-FCT employs $\mathbb{P}_1$ finite elements and a grid spacing of $h = 1/128$. We run our MMOC implementation both, with $\mathbb{P}_1$ and $\mathbb{P}_2$ elements and grid spacings of $h = 1/128$ and $h = 1/64$ respectively. The mesh size $h$ refers to quadratic elements, which are divided into two triangles each. Therefore, all grids result in 16641 DoFs, including boundary. The time-step size is $\tau \approx 1 \cdot 10^{-3}$ for all settings (which corresponds to 6283 time steps for one revolution).

| method | $\|e_h\|_{\mathcal{H}^0}$ | var$(2\pi)$ | $\Delta m(2\pi)$ |
|---|---|---|---|
| MMOC $\mathbb{P}_1$ ($b=1$) | $1.74 \cdot 10^{-1}$ | 0.5913 | $-4.73 \cdot 10^{-2}$ |
| MMOC $\mathbb{P}_1$ ($b=10$) | $1.65 \cdot 10^{-1}$ | 0.6296 | $4.76 \cdot 10^{-3}$ |
| MMOC $\mathbb{P}_1$ ($b=100$) | $8.60 \cdot 10^{-2}$ | 0.9847 | $2.67 \cdot 10^{-4}$ |
| MMOC $\mathbb{P}_1$ ($b=1000$) | $3.85 \cdot 10^{-2}$ | 1.0000 | $-5.52 \cdot 10^{-4}$ |
| MMOC $\mathbb{P}_1$ ($b=\infty$) | $1.38 \cdot 10^{-13}$ | 1.0000 | $2.22 \cdot 10^{-16}$ |
| FEM-FCT n.-l. [28] | $1.44 \cdot 10^{-1}$ | 1.0010 | no data |
| FEM-FCT [28] | $1.92 \cdot 10^{-1}$ | 1.0069 | no data |

| method | $\|e_h\|_{\mathcal{H}^0}$ | var$(2\pi)$ | $\Delta m(2\pi)$ |
|---|---|---|---|
| MMOC $\mathbb{P}_2$ ($b=1$) | $1.09 \cdot 10^{-1}$ | 1.2773 | $-2.19 \cdot 10^{-2}$ |
| MMOC $\mathbb{P}_2$ ($b=10$) | $9.71 \cdot 10^{-2}$ | 1.2943 | $-1.39 \cdot 10^{-2}$ |
| MMOC $\mathbb{P}_2$ ($b=100$) | $5.29 \cdot 10^{-2}$ | 1.3049 | $-5.56 \cdot 10^{-3}$ |
| MMOC $\mathbb{P}_2$ ($b=1000$) | $3.03 \cdot 10^{-2}$ | 1.3185 | $-1.09 \cdot 10^{-3}$ |
| MMOC $\mathbb{P}_2$ ($b=\infty$) | $1.68 \cdot 10^{-13}$ | 1.0000 | $-6.79 \cdot 10^{-14}$ |



(a) interpolated solution     (b) computed solution, $b = 1000$     (c) computed solution, $b = \infty$
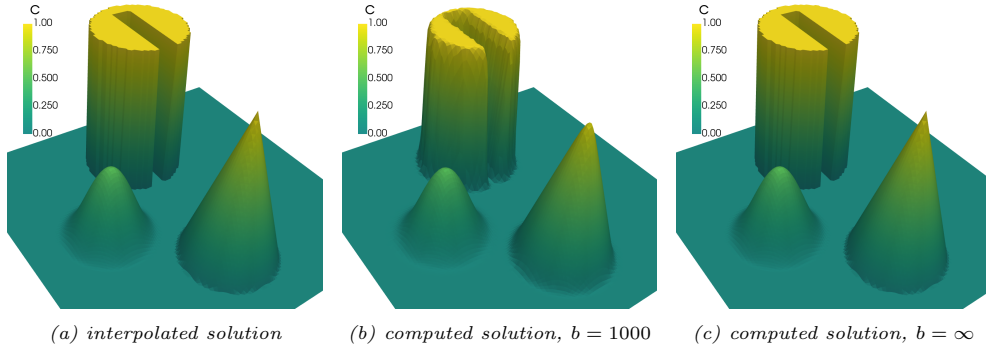
Figure 3: Interpolated and computed solutions of the body rotation problem for the setup as in table 1 with $\mathbb{P}_1$ discretization for different look-back distances $b$. The plots show the influence of the look-back distance on the quality of the computed solution.

misleading. While there are no oscillations at time $t = 2\pi$, some oscillations appear at the discontinuity around the slotted cylinder for $0 < t < 2\pi$. For the linear ($\mathbb{P}_1$) space discretization, there are no oscillations over the entire time interval.

Typically, continuous Lagrange finite elements of higher order tend to produce over- and undershoots at discontinuities. However, we note, that this is owed to the space-discretization and not to the presented time-discretization, i.e. the MMOC. An advantage of the MMOC is, that it can be applied to any space-discretization as long as the solution can be locally evaluated. In presence of discontinuities in the solution, discontinuous Galerkin space-discretizations could be considered in combination with the MMOC.

To demonstrate, the contribution of the discontinuity around the slotted cylinder to the error, we show in table 2 results for the smooth initial condition and solution $c_0 = c_0^{\text{hill}}$. In this run, the error and energy discrepancy is much smaller than for the results with a discontinuous solution in table 1, especially, for $b < \infty$.

Table 2 additionally lists results of simulations with time-steps that are increased by a factor of 10 and 100. The measured errors demonstrate that the Lagrangian approach yields promising stability and accuracy also for comparatively large time-steps.

For $b < \infty$ the resulting errors are mostly caused by the interpolation between the Eulerian and the Lagrangian representation. If we compare runs where $\tau \cdot b = \text{const}$, we obtain almost identical errors (e.g. $\tau = 1.01 \cdot 10^{-1}$ and $b = 1$ compared to $\tau = 1.00 \cdot 10^{-2}$ and $b = 10$ in table 2). In those runs, the number of time-steps in which the solution is interpolated is equal. Despite a significant time-step size reduction, the interpolation error dominates. In the case of $b = \infty$ no temperature interpolation is performed throughout the simulation. Therefore the increased accuracy of the RK integrator directly affects the error in the solution when the time-step size is reduced.

| $\tau$ | $b$ | $\|e_h\|_{\mathcal{H}^0}$ | $\mathrm{var}(2\pi)$ | $\Delta m(2\pi)$ |
|---|---|---|---|---|
| $1.01 \cdot 10^{-1}$ | 1 | $3.36 \cdot 10^{-4}$ | 0.5019 | $-6.63 \cdot 10^{-5}$ |
| $1.01 \cdot 10^{-1}$ | 10 | $5.32 \cdot 10^{-5}$ | 0.5008 | $3.83 \cdot 10^{-5}$ |
| $1.01 \cdot 10^{-1}$ | $\infty$ | $9.62 \cdot 10^{-7}$ | 0.5000 | $9.35 \cdot 10^{-7}$ |
| $1.00 \cdot 10^{-2}$ | 1 | $4.33 \cdot 10^{-3}$ | 0.5054 | $-4.46 \cdot 10^{-6}$ |
| $1.00 \cdot 10^{-2}$ | 10 | $3.43 \cdot 10^{-4}$ | 0.5020 | $-5.33 \cdot 10^{-5}$ |
| $1.00 \cdot 10^{-2}$ | 100 | $4.87 \cdot 10^{-5}$ | 0.5008 | $2.28 \cdot 10^{-5}$ |
| $1.00 \cdot 10^{-2}$ | $\infty$ | $9.13 \cdot 10^{-11}$ | 0.5000 | $9.12 \cdot 10^{-12}$ |
| $1.00 \cdot 10^{-3}$ | 1 | $6.33 \cdot 10^{-3}$ | 0.5066 | $-2.05 \cdot 10^{-5}$ |
| $1.00 \cdot 10^{-3}$ | 10 | $4.33 \cdot 10^{-3}$ | 0.5054 | $-4.38 \cdot 10^{-6}$ |
| $1.00 \cdot 10^{-3}$ | 100 | $3.40 \cdot 10^{-4}$ | 0.5020 | $-4.80 \cdot 10^{-5}$ |
| $1.00 \cdot 10^{-3}$ | 1000 | $4.96 \cdot 10^{-5}$ | 0.5009 | $3.27 \cdot 10^{-5}$ |
| $1.00 \cdot 10^{-3}$ | $\infty$ | $9.87 \cdot 10^{-15}$ | 0.5000 | $4.88 \cdot 10^{-15}$ |



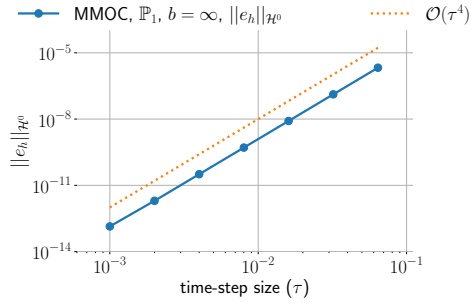Figure 4: Time-step size study for the circular advection benchmark with $c_0 = c_0^{slotted} + c_0^{cone} + c_0^{hill}$ and $b = \infty$.

Table 2: Results for the circular advection benchmark with $c_0 = c_0^{hill}$ and a $\mathbb{P}_2$ space discretization.

In fig. 4 we plot the $\mathcal{H}^0$ error of the solution of the original benchmark problem ($c_0 = c_0^{slotted} + c_0^{cone} + c_0^{hill}$, i.e. with discontinuous solution) discretized with $\mathbb{P}_1$ finite elements for different time-step sizes and $b = \infty$. For the largest time-step size in fig. 4 ($\tau \approx .065$) and a maximum absolute velocity of $\approx .7$ this results in a CFL number of roughly 3.

## 3.3 Swirling advection

Next we move to a three-dimensional setting with a time-dependent velocity field. The benchmark is taken from [34]. Let $\Omega = (0,1)^3$ and $t \in [0,T]$ with $T = 1.5$. The the initial condition $c_0$ and the velocity field $\mathbf{u}(\mathbf{x}, t)$ are defined by

$$c_0(\mathbf{x}) := \begin{cases} 1 & \text{if } x_1 < 0.5 \\ 0 & \text{otherwise} \end{cases} \quad , \quad \mathbf{u}(\mathbf{x}, t) := \begin{pmatrix} 2\sin^2(\pi x_1)\sin(2\pi x_2)\sin(2\pi x_3)g(t) \\ -\sin(2\pi x_1)\sin^2(\pi x_2)\sin(2\pi x_3)g(t) \\ -\sin(2\pi x_1)\sin(2\pi x_2)\sin^2(\pi x_3)g(t) \end{pmatrix},$$

with $g(t) := \cos(\pi t/T)$. The temperature field undergoes a deformation which reverses at $t = T/2$ and should return to the initial solution at $t = T$. Again, we consider pure advection ($\kappa = 0$, $q = 0$).

The results for the MMOC with $b = \infty$ at $T = 1.5$ are listed in table 3. The $\mathcal{H}^0$-errors are small and no spurious oscillations are detected for all chosen time step and grid sizes. Figure 5 shows the computed solution at $x_3 = 0.425$ for $t = T/2$, and

Table 3: Results for application of the MMOC with $b = \infty$ and $\mathbb{P}_1$ finite elements in space to the swirling flow benchmark in 3D.

| DoFs | $h_{\min}$ | $\tau$ | $\|e_h\|_{\mathcal{H}^0}$ | $\mathrm{var}(1.5)$ | $\Delta m(1.5)$ |
|---|---|---|---|---|---|
| 35,937 | $3.12 \cdot 10^{-2}$ | $1.00 \cdot 10^{-1}$ | $8.67 \cdot 10^{-4}$ | 1.0000 | $-9.75 \cdot 10^{-5}$ |
| 35,937 | $3.12 \cdot 10^{-2}$ | $5.00 \cdot 10^{-2}$ | $5.48 \cdot 10^{-5}$ | 1.0000 | $-2.37 \cdot 10^{-6}$ |
| 35,937 | $3.12 \cdot 10^{-2}$ | $2.50 \cdot 10^{-2}$ | $5.11 \cdot 10^{-6}$ | 1.0000 | $-1.37 \cdot 10^{-8}$ |
| 2,146,689 | $7.81 \cdot 10^{-3}$ | $1.00 \cdot 10^{-1}$ | $1.26 \cdot 10^{-3}$ | 1.0000 | $-2.77 \cdot 10^{-5}$ |
| 2,146,689 | $7.81 \cdot 10^{-3}$ | $5.00 \cdot 10^{-2}$ | $5.01 \cdot 10^{-5}$ | 1.0000 | $-1.11 \cdot 10^{-6}$ |
| 2,146,689 | $7.81 \cdot 10^{-3}$ | $2.50 \cdot 10^{-2}$ | $2.57 \cdot 10^{-6}$ | 1.0000 | $-2.18 \cdot 10^{-8}$ |

$t = T$ with $h_{\min} = 7.81 \cdot 10^{-3}$ (refinement level 7), and $\tau = 5.00 \cdot 10^{-2}$. At $t = T$, the initial temperature field is restored without visible artifacts or numerical diffusion. The slice is chosen to coincide with the slice shown in [34, figure 11.2].

## 3.4 Advection-diffusion on blended geometry

Finally we apply the MMOC to a problem with a diffusion coefficient $\kappa > 0$, without internal heating ($q = 0$), on a blended geometry. Since $\kappa > 0$, we need to solve the

*(a) interpolated initial condition*   *(b) $t = T/2$ (computed solution)*   *(c) $t = T$ (computed solution)*

Figure 5: *Elevated slice of the computed solution at $x_3 = 0.425$, as in [34, figure 11.2]. Parameters: $\mathbb{P}_1$ discretization, $b = \infty$, $h_{min} = 7.81 \cdot 10^{-3}$ (refinement level 7), $\tau = 5.00 \cdot 10^{-2}$. The discontinuities are preserved without any oscillations or numerical diffusion.*

linear system (7) in each time-step. Thus the solution must be interpolated to the Eulerian grid in each time-step and we must limit the look-back distance to $b = 1$.

The physical domain $\Omega_{\text{phy}}$ is an annulus defined by $\Omega = \{\mathbf{x} \in \mathbb{R}^2 : r_{\min} \leq \|\mathbf{x}\|_2 \leq r_{\max}\}$ with $(r_{\min}, r_{\max}) = (0.5, 1.5)$. The computational domain $\Omega_{\text{comp}}$ approximates the annulus with a coarse triangular mesh that is uniformly refined and projected onto $\Omega_{\text{phy}}$, see fig. 1.

The benchmark is inspired by the unsteady advection-diffusion benchmark in [32]. A circular velocity field $\mathbf{u}(\mathbf{x}) = (-x_2, x_1)$ transports a gradually smeared Gaussian hill around the annulus. The time-dependent position of the hill is given by an initial position $(\bar{x}_1, \bar{x}_2) = (0, 1)$ and $\hat{\mathbf{x}}(t) = (\hat{x}_1(t), \hat{x}_2(t))$ with $\hat{x}_1(t) = \bar{x}_1 \cos(t) - \bar{x}_2 \sin(t)$, and $\hat{x}_2(t) = -\bar{x}_1 \sin(t) + \bar{x}_2 \cos(t)$. The analytical solution $c$ is defined by

$$c(\mathbf{x}, t) := \frac{1}{4\pi t \kappa} \exp\left( -\frac{r(\mathbf{x}, t)^2}{4t\kappa} \right), \quad \kappa > 0, \ t > 0 \ ,$$

where $r(\mathbf{x}, t) := \|\mathbf{x} - \hat{\mathbf{x}}\|_2$. At time $t = 0$, $c$ becomes a Dirac delta function, which is why we start the simulation at a later time. To better compare the proposed method for different $\kappa$, we parameterize the initial time $t_0$ via $t_0(\kappa) = \frac{2\pi \times 10^{-3}}{\kappa}$. The moving hill has therefore a different initial position when applying a different diffusion coefficient, but the initial shape is identical for all choices of $\kappa$. The computed solution is evaluated after a full revolution, at time $T = t_0(\kappa) + 2\pi$. We employ a $\mathbb{P}_2$ space-discretization in all runs. Results after one revolution are listed in table 4. For a

Table 4: *Results after one revolution of the Gaussian hill, with time-step size $\tau \approx 1e - 1$, implicit Euler time-integration for the parabolic part ($\Theta = 1$ in (6)), three different diffusion coefficients $\kappa$, and different refinement levels.*

| | | $\kappa = 1 \cdot 10^{-3}$ | | $\kappa = 1 \cdot 10^{-5}$ | | $\kappa = 1 \cdot 10^{-7}$ | |
|---|---|---|---|---|---|---|---|
| DoFs | $h_{\min}$ | $\|e^h\|_{\mathcal{H}^0}$ | $E_{\text{peak}}(2\pi)$ | $\|e^h\|_{\mathcal{H}^0}$ | $E_{\text{peak}}(2\pi)$ | $\|e^h\|_{\mathcal{H}^0}$ | $E_{\text{peak}}(2\pi)$ |
| 12 480 | $3.12 \cdot 10^{-2}$ | $1.48 \cdot 10^{-2}$ | $-1.30 \cdot 10^{-3}$ | $8.32 \cdot 10^{-2}$ | $-2.02 \cdot 10^{-2}$ | $8.51 \cdot 10^{-2}$ | $-1.98 \cdot 10^{-2}$ |
| 49 536 | $1.56 \cdot 10^{-2}$ | $3.86 \cdot 10^{-3}$ | $3.45 \cdot 10^{-3}$ | $7.38 \cdot 10^{-3}$ | $-1.90 \cdot 10^{-3}$ | $7.84 \cdot 10^{-3}$ | $-1.56 \cdot 10^{-3}$ |
| 197 376 | $7.81 \cdot 10^{-3}$ | $4.32 \cdot 10^{-3}$ | $3.98 \cdot 10^{-3}$ | $6.30 \cdot 10^{-4}$ | $-1.25 \cdot 10^{-4}$ | $6.92 \cdot 10^{-4}$ | $-1.01 \cdot 10^{-4}$ |

comparably large time step size (CFL between 4 and 20) and varying diffusivity, we observe satisfying results for sufficiently small mesh sizes.

# 4   Parallel performance

Scalability of the HHG data structures in HyTeG, and of the particle dynamics framework MESA-PD has separately been demonstrated on some of the worlds largest

supercomputers, [21, 29]. It remains to assess the parallel performance of our MMOC implementation, in which both software architectures are coupled.

For the scalability benchmark, we set up an elongated, three-dimensional cuboid domain, where a smooth initial temperature field is initialized, and transported along a constant velocity field, resembling flow through a pipe. This setup allows for straightforward parameterization of domain size, and number of coarse grid primitives.

We employ $\mathbb{P}_2$ finite elements for the space-discretization and perform a single time-step, consisting of particle creation, particle integration, and temperature evaluation, including synchronization (steps (i)–(iii) in section 2.3.1).

All runs in this section were performed on SuperMUC-NG, ranked 15[th] in the Top500[2] list (Nov 2020). The system is composed of 6,336 so-called thin-nodes, 3,072 of which we had access to at the time of writing. Two Intel® Skylake Xeon® Platinum 8174 CPUs are installed on each node, which sums up to 48 cores per node and 147,456 cores in total on the accessible 3,072 nodes. Per node 96GB of main memory are available.

**Strong scaling**   We conduct a strong scaling experiment on a grid that consists of 3,072 tetrahedral coarse grid elements, each of which is refined 4 times, resulting in $\approx 1.7 \cdot 10^7$ DoFs in total. Leaving the grid fixed, we increase the number of processes, so that the number of particles per process decreases. As a baseline for the parallel performance we consider a single-node run. We plot the parallel performance and the number of updated particles per second in fig. 6a. For the largest setting with 64 nodes (3,072 processes) we obtain a parallel efficiency of roughly 36 % for one macro-cell and $\approx 5,500$ particles per process.

**Weak scaling**   Additionally, we perform a weak-scaling experiment, where the number of DoFs per process is kept constant. In this setting, each process is assigned a single tetrahedral macro-cell and we refine the initial grid 5 times. This results in about $3.55 \cdot 10^5$ DoFs per process. Starting from a single node, again used as baseline for parallel efficiency, we scale up to the available 147,456 processes of SuperMUC-NG. In the largest scenario, this amounts to more than $5.2 \cdot 10^{10}$ DoFs in total for the discretization of the solution of the advection-diffusion equation (1). All runs maintain an excellent parallel efficiency of more than 92 %.
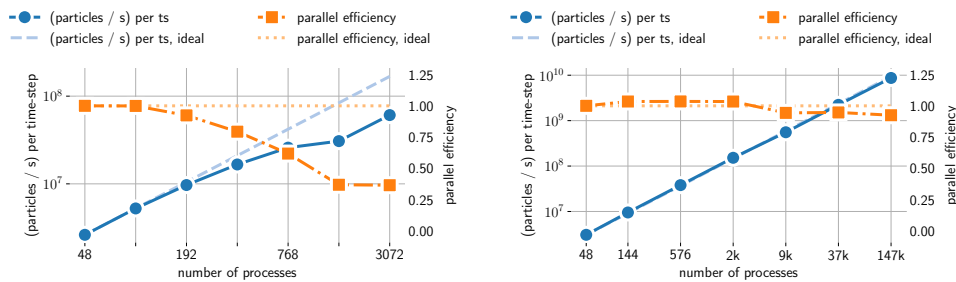


(a) strong scaling, $\approx 1.7 \times 10^7$ temperature DoFs (i. e. particles) in total

(b) weak scaling, $\approx 3.55 \times 10^5$ temperature DoFs (i. e. particles) per process

Figure 6: Weak and strong scaling results (ts = time-step)

Overall, we observe a run time per time step of about 5 or less seconds in all tested scenarios, and less than a second in the strong-scaling limit. We note, however, that it is not sufficient to consider the number of updated DoFs per second alone as a measure to quantify the efficiency of the method. The results of section 3 show that the stability

---

and accuracy of the MMOC allows for large time-steps even in the strongly advection-dominated problems. This may be an advantage in coupled convection simulations as they appear in Earth mantle convection, where the majority of the run time is spent for the solution of the Stokes system [25, 31]. Given sufficiently accurate coupling schemes, the MMOC does not only allow for a reduction of simulation time by itself, but also permits to advance faster in time, due to less restrictive CFL limitations.

# 5   Coupled flow

Finally, we apply our scheme to a buoyancy-driven, coupled flow problem. Due to the negligible Reynolds number in mantle convection models, the Stokes equations are used to model the creeping flow of the medium. We consider the incompressible formulation for the Boussinesq approximation [40]

$$-\nabla \cdot \sigma = \mathbf{F}(c), \quad \nabla \cdot \mathbf{u} = 0$$
$$\sigma(\mathbf{u}, p) = 2\mu\epsilon(\mathbf{u}) - pI, \quad \epsilon(\mathbf{u}) = \frac{1}{2}\left(\nabla\mathbf{u} + (\nabla\mathbf{u})^\top\right). \tag{9}$$

where $\mu$ is a viscosity field, $p$ the pressure, $\mathbf{F}(c) := \operatorname{Ra} c \mathbf{g}$ a temperature dependent forcing term, Ra the Rayleigh number, $\mathbf{g}$ the normalized gravitation, and $\sigma$ the Cauchy stress tensor associated with an incompressible, highly viscous Newtonian fluid.

The PDEs (1) and (9) are coupled through both, the velocity which is the solution of the Stokes system and drives the advection of the temperature, and the temperature which enters the Stokes equation through the forcing term. Different from the benchmarks in section 3, the convectivity is in the following setups steered solely through the Rayleigh number Ra, i.e. we set $\kappa = 1$ in (1). When Ra is large, so is the right-hand side (RHS) of (9) and the velocity that enters (1) has a large magnitude, resulting in advection-dominated transport. Note that since $\kappa > 0$, the look-back distance is set to $b = 1$ in the following benchmarks.

The advection-diffusion equation (1) is constrained by the Stokes equation (9) at all times, and a non-linear system must be solved at each time-step, which is at least a computationally expensive challenge for large-scale simulations. In practice, (1) is thus usually *decoupled* from the constraints to the velocity $\mathbf{u}$, so that the systems can be solved in an alternating fashion [31, 50].

The solution of the Stokes system in each time-step dominates the computational cost of this scheme and is therefore crucial to performance. We employ an efficient, monolithic matrix-free geometric multigrid solver as described in [29], and large CFL-numbers to reduce the number of required solves. The Stokes system is discretized with a mixed $\mathbb{P}_2$-$\mathbb{P}_1$ finite element approximation. For more in-depth discussion of efficient matrix-free geometric multigrid solvers on HHG, we refer to [2, 4, 5, 25, 26, 29].

In section 5.2, we outline a predictor-corrector scheme (see algorithm 3) to approximate the solution of the non-linear, coupled system, and apply the method to two benchmark problems.

## 5.1   Strang-splitting

We employ a tighter coupling of the advection- and diffusion-step via a Strang-splitting approach [44]. Instead of an alternating application of the advection and diffusion step, the diffusion step is split, and the advection step is framed by two fractional diffusion steps with reduced time-step size, giving a scheme with three stages. The algorithm is listed in algorithm 2.

The splitting procedure noticeably increases the accuracy of the method in the benchmarks of this section, however, we did not observe relevant differences when applying it to the advection-diffusion benchmark in section 3.4.

---

**Algorithm 2** Time-stepping scheme, advection-diffusion, with Strang-splitting

---
1: **procedure** ADS($c_h^n, \mathbf{u}_h$)
2:     $\tau_n = \text{CFL}_{\max} \cdot h_{\min} / \max_{\mathbf{x} \in \Omega} |\mathbf{u}_h(\mathbf{x}, t_n)|$          ▷ determine time-step size
3:     solve (7) with $\tau_n^* = \tau_n/2$ to advance from $c_h^n$ to $c_h^{n+(1/3)}$        ▷ diffusion
4:     $\hat{\mathbf{x}} = \mathbf{X}(\mathbf{x}, t_{n+1}, t_n)$          ▷ calculate departure points (see section 2.3)
5:     $c_h^{n+(2/3)}(\mathbf{x}) = c_h^{n+(1/3)}(\hat{\mathbf{x}})$          ▷ advection
6:     solve (7) with $\tau_n^* = \tau_n/2$ to advance from $c_h^{n+(2/3)}$ to $c_h^{n+1}$      ▷ diffusion
7:     **return** $c_h^{n+1}$

---

## 5.2 A predictor-corrector scheme

To resolve the non-linear coupling of the advection-diffusion equation (1) and the Stokes problem (9), we apply a predictor-corrector method [48], as outlined in algorithm 3.

---

**Algorithm 3** Predictor-corrector scheme to couple (1) and (9). In each time-step both PDEs are solved twice. For the advection-diffusion step, algorithm 2 is employed.

---
1: **procedure** PC
2:     solve (9) for $\mathbf{u}_h^0$          ▷ initial velocity field
3:     **for** $n \in \{0, 1, \dots\}$ **do**
4:         $\tilde{\mathbf{u}}_h(\mathbf{x}, t) \leftarrow \mathbf{u}_h^n(\mathbf{x})$      ▷ time-invariant velocity field at time-step $n$ for temp. predictor
5:         $c_h^{\text{pr}} \leftarrow \text{ADS}(c_h^n, \tilde{\mathbf{u}}_h)$          ▷ predict temperature
6:         solve (9) for $\mathbf{u}_h^{\text{pr}}$ with RHS $\mathbf{F}(c_h^{\text{pr}})$          ▷ predict velocity
7:         $\tilde{\mathbf{u}}_h(\mathbf{x}, t) \leftarrow \text{lerp}(\mathbf{u}_h^n, \mathbf{u}_h^{\text{pr}}, t)(\mathbf{x})$      ▷ linear interpolation in $t$ between $\mathbf{u}_h^n$ and $\mathbf{u}_h^{\text{pr}}$
8:         $c_h^{n+1} \leftarrow \text{ADS}(c_h^n, \tilde{\mathbf{u}}_h)$          ▷ correct temperature
9:         solve (9) for $\mathbf{u}_h^{n+1}$ with RHS $\mathbf{F}(c_h^{n+1})$          ▷ correct velocity

---

For the temperature prediction step, we approximate the velocity field with the time-invariant state at $t = t_n$, i.e. the interpolation (8) yields $\mathbf{u}_h^n$ for all $t \in [t_n, t_{n+1}]$. A prediction $\mathbf{u}_h^{\text{pr}}$ for the velocity is then computed using the predicted temperature field for the RHS force term of (9). The correction step is then executed, employing the interpolation in (8) between $\mathbf{u}_h^n$ at $t = t_n$, and $\mathbf{u}_h^{\text{pr}}$ at $t = t_{n+1}$. Finally a new velocity solution is computed using the corrected temperature field.

## 5.3 Time-dependent convection benchmark

To verify our implementation, we consider a classical benchmark from Blankenbach et al. [9] (case 3) that was also investigated e.g. in [49]. The test considers time-dependent convection with constant viscosity ($\mu = 1$ in (9)) and internal heating ($q = 1$ in (1)) in a two-dimensional, rectangular domain $\Omega = [0, L] \times [0, H]$, $L = 1.5, H = 1$. The top, bottom, and side boundaries are denoted as $\Gamma_t$, $\Gamma_b$, and $\Gamma_s$. $\mathbf{n}$ and $\mathbf{t}$ represent the outward normal, and tangential vectors respectively. For the velocity free-slip conditions are prescribed at the vertical boundaries ($\mathbf{u} \cdot \mathbf{n} = \sigma \mathbf{n} \cdot \mathbf{t} = 0$ for $\mathbf{x} \in \Gamma_s$), and no-slip conditions at the horizontal boundaries ($\mathbf{u} = 0$ for $\mathbf{x} \in \Gamma_t \cup \Gamma_b$). For the temperature zero Dirichlet boundary conditions are prescribed at the top boundary ($c = 0$ for $\mathbf{x} \in \Gamma_t$), and Neumann boundaries otherwise ($\partial_{\mathbf{n}} c = 0$ for $\mathbf{x} \in \Gamma_b \cup \Gamma_s$). We employ the initial condition $c_0(\mathbf{x}) = 0.5 \left(1 - x_2^2\right) + 0.01 \cos(\pi x_1/L) \sin(\pi x_2/H)$ given in [49].

The benchmark solution is expected to exhibit a characteristic, periodic development of downwelling plumes and is quantified via the local extrema of the root-mean-square velocity $\mathbf{u}_{\text{rms}}$ and the Nusselt number Nu, defined as

$$\mathbf{u}_{\text{rms}} = \left( \frac{1}{|\Omega|} \int_\Omega \|\mathbf{u}\|^2 \, dx \right)^{1/2}, \quad \text{Nu} = -\frac{\int_0^L \partial_{x_2} c(\mathbf{x}, x_2 = H) dx}{\int_0^L c(\mathbf{x}, x_2 = 0) dx} \quad .$$

At low Rayleigh numbers, every plume shows the same behavior. With increasing Ra the periodicity is characterized by every $n$-th plume behaving identically, resulting

in a *Pn*-cycle. In particular, the benchmark suggests that the convective motion transitions from a *P2*- to a *P4*-cycle between Ra = 216,000 and Ra = 218,000. We partition a Pn-cycle into $n$ time intervals, denoted as stages $S0, \ldots, S(n-1)$. Each stage of a cycle comprises a local maximum of $\mathbf{u}_{\mathrm{rms}}$ and Nu, followed by a local minimum.

We apply the time-stepping scheme algorithm 3 on two meshes with sizes $24 \times 16$ and $48 \times 32$ squares (each divided into 2 triangles) and for two CFL-numbers (0.5 and 1) running the simulation from $t = 0$ to $t = 3$. For $t \in [2.5, 3]$, the described repetitive cyclic motion of the plumes is observed. The computed solution is compared to the reference values in [9, table 9] for Ra = 216,000, and [9, table 8a, Code Ha, $96 \times 64$] for Ra = 218,000. We selected the latter reference from the various codes compared in [9] as the presumably most accurate implementation, and note that no analytical solution is known.

The relative errors (compared to the reference) of the minima and maxima of Nu and $\mathbf{u}_{\mathrm{rms}}$ are calculated. For Ra = 216,000, all extrema coincide with the reference up to a relative error of less than $0.4\%$ for both meshes and CFL numbers. For $Ra = 218,000$, a maximum relative error of less than $1\%$ for all extrema is reached for the finer mesh with $48 \times 32$ squares. We conclude that the computed results agree well with those reported in [9, 49]. The characteristic trends of $\mathbf{u}_{\mathrm{rms}}$ and Nu for both scenarios, with mesh size $48 \times 32$ and CFL = 1.0 are plotted in fig. 7.


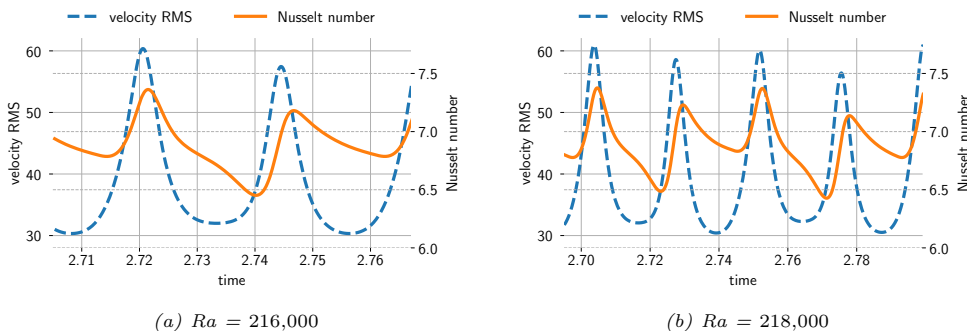
*(a) Ra = 216,000*                         *(b) Ra = 218,000*

*Figure 7:* $\mathbf{u}_{rms}$ *and Nu plotted over a P2-cycle for Ra = 216,000, and a P4-cycle for Ra = 218,000 (mesh: $48 \times 32$, CFL = 1).*

## 5.4 Mantle convection on a spherical shell

As a demonstrator for the applicability to large scale applications, we employ the coupled method to simulate isoviscous convection with Ra = $10^8$ and no internal heating ($q = 0$). The domain approximates Earth's mantle by the spherical shell $\Omega = \{\mathbf{x} \in \mathbb{R}^3 : r_{\min} \leq \|\mathbf{x}\|_2 \leq r_{\max}\}$ with $r_{\min} = 0.5$ and $r_{\max} = 1$. The computational grid is composed of 19,200 tetrahedral macro-cells, which are refined 4 times and projected onto the sphere, resulting in more than $3.2 \cdot 10^8$ unknowns for the Stokes equation, and $1.0 \cdot 10^8$ DoFs (and therefore particles) for the advection-diffusion equation, solved for in every time step.

The initial and Dirichlet boundary conditions for the temperature are prescribed by $c_0(r) = \exp\left(-10 \frac{r - r_{\min}}{r_{\max} - r_{\min}}\right)$ where $r$ is the distance to the origin. For the velocity, we set no-slip boundary conditions at all boundaries.

We apply the predictor-corrector scheme in algorithm 3 with Strang-splitting, and simulate 3,000 time-steps with a CFL-number of 1. The Stokes system is solved with a monolithic geometric multigrid solver that employs an inexact Uzawa smoother with weighted Jacobi relaxation [29]. Its excellent performance and scalability to linear systems more than a trillion $(10^{12})$ unknowns is discussed in [25, 26, 29]. For

the diffusive term, i.e. the solution of the linear system (7), we employ a standard conjugate gradient iteration, which turns out to be sufficient.

The simulation is performed on 400 nodes (19,200 processes) of SuperMUC-NG in roughly 16 hours. In fig. 8a we list a summary of the benchmark parameters. Figure 8b shows a stacked bar chart of the fractional run time of the relevant components of the predictor-corrector scheme. On average, the computation of a single time-step takes about 19 seconds. About 85 % of the total run time is spent for the solution of the Stokes system. Almost half of that time (roughly 41 % of total run time) accounts for communication during the Jacobi-relaxation. Especially during iterations on the coarser grids, communication time strongly dominates time spent in the compute kernels. Strategies to further improve the performance of the coarse grid solver are presented in [13].



| Mantle convection benchmark parameters | |
| --- | --- |
| machine | SuperMUC-NG |
| nodes | 400 |
| cores | 19,200 |
| numerical scheme | predictor-corrector (see algorithm 3) |
| Stokes | |
| discretization | $\mathbb{P}_2$-$\mathbb{P}_1$ (Taylor-Hood) |
| solver | monolithic geometric multigrid (GMG) |
| DoFs | $3.2 \cdot 10^8$ |
| temperature | |
| discretization | $\mathbb{P}_2$ finite elements + MMOC |
| advection scheme | MMOC |
| solver diffusion | conjugate gradient (CG) |
| DoFs | $1.0 \cdot 10^8$ |
| Rayleigh number | $10^8$ |
| CFL | 1 |
| avg. run time / ts | $\approx$ 19s (incl. pred. + corr., I/O) |

(a) parameter summary

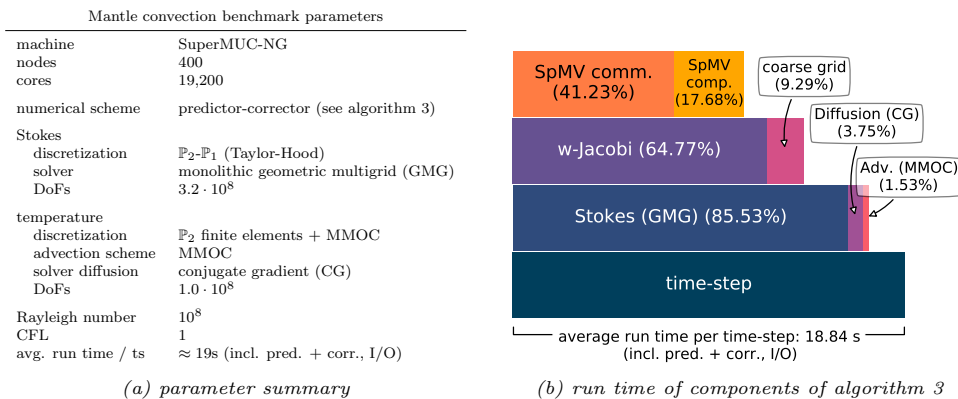(b) run time of components of algorithm 3

Figure 8: Mantle convection benchmark: (a) Summary of parameters. (b) Stacked bar chart of average fractional run time of components of algorithm 3. The percentage in parentheses indicates the fractional run time with respect to the overall run time of a predictor-corrector step.

In fig. 9, the contour surfaces of the temperature at $c_{\text{cont}} = 0.15$ are shown at time steps # 200, and 3,000. Thin, chaotically rising plumes are observed as expected at
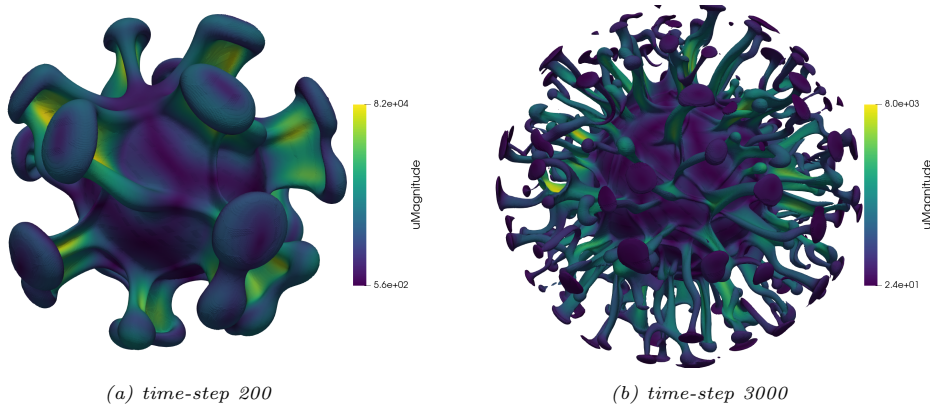


(a) time-step 200

(b) time-step 3000

Figure 9: Contour plot at $c = 0.15$ of the temperature solution on the spherical shell for $Ra = 10^8$, colored by velocity magnitude.

such large Rayleigh numbers.

# Conclusion

In this article, we presented an implementation of an Eulerian-Lagrangian discretization based on the method of characteristics to treat the advection-diffusion equation

in the advection-dominated regime. Its numerical performance was demonstrated on multiple two- and three-dimensional benchmarks, including cases with pure advection, curved geometries, and discontinuous solutions.

Motivated by the demand of extreme spatial resolution in mantle convection simulations, the parallel scalability of our implementation was assessed in a weak and strong scaling benchmark for the advection-diffusion equation. We demonstrate a parallel efficiency of more than $92\,\%$, solving for more than $5.2 \cdot 10^{10}$ DoFs per time-step on 147,456 parallel processes. Finally, we applied the method to buoyancy-driven Stokes flow, embedding it into a non-linear scheme based on a predictor-corrector method. The scheme was verified through a classical benchmark for time-dependent convection, and its practical applicability to large scale problems demonstrated in a mantle convection benchmark on the spherical shell, with combined more than $4.0 \cdot 10^8$ unknowns solved for in each time step for 3,000 time steps.

# Acknowledgements

# References

[1] Alejandro Allievi and Rodolfo Bermejo. Finite element modified method of characteristics for the Navier-Stokes equations. *Int. J. Numer. Meth. Fluids*, 32(4):439–463, 2000.

[2] S. Bauer, H.-P. Bunge, D. Drzisga, S. Ghelichkhan, M. Huber, N. Kohl, M. Mohr, U. Rüde, D. Thönnes, and B. Wohlmuth. TerraNeo — Mantle Convection Beyond a Trillion Degrees of Freedom. In H.-J. Bungartz, S. Reiz, B. Uekermann, P. Neumann, and W. Nagel, editors, *Software for Exascale Computing - SPPEXA 2016-2019*, volume 136 of *Lecture Notes in Computational Science and Engineering*, pages 569–610. Springer, 2020.

[3] S. Bauer, D. Drzisga, M. Mohr, U. Rüde, C. Waluga, and B. Wohlmuth. A stencil scaling approach for accelerating matrix-free finite element implementations. *SIAM J. Sci. Comp.*, 40(6):C748–C778, 2018.

[4] S. Bauer, M. Huber, S. Ghelichkhan, M. Mohr, U. Rüde, and B. Wohlmuth. Large-scale Simulation of Mantle Convection Based on a New Matrix-Free Approach. *J. Comput. Sci.*, 31:60–76, 2019.

[5] S. Bauer, M. Mohr, U. Rüde, J. Weismüller, M. Wittmann, and B. Wohlmuth. A two-scale approach for efficient on-the-fly operator assembly in massively parallel high performance multigrid codes. *Appl. Numer. Math.*, 122:14–38, 2017.

[6] B. Bergen and F. Hülsemann. Hierarchical hybrid grids: data structures and core algorithms for multigrid. *Numer. Linear Algebra Appl.*, 11:279–291, 2004.

[7] Alfredo Bermúdez, Maria R Nogueiras, and Carlos Vázquez. Numerical analysis of convection-diffusion-reaction problems with higher order characteristics/finite

elements. part i: time discretization. *SIAM J. Numer. Anal.*, 44(5):1829–1853, 2006.

[8] Jürgen Bey. Tetrahedral grid refinement. *Computing*, 55(4):355–378, 1995.

[9] B. Blankenbach, F. Busse, U. Christensen, L. Cserepes, D. Gunkel, U. Hansen, H. Harder, G. Jarvis, M. Koch, G. Marquart, D. Moore, P. Olson, H. Schmeling, and T. Schnaubelt. A benchmark comparison for mantle convection codes. *Geophys. J. Int.*, 98(1):23–38, 1989.

[10] Alexander N. Brooks and Thomas J. R. Hughes. Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Comp. Meth. Appl. Mech. Engrg.*, 32(1-3):199–259, 1982.

[11] Carsten Burstedde, Georg Stadler, Laura Alisic, Lucas C. Wilcox, Eh Tan, Michael Gurnis, and Omar Ghattas. Large-scale adaptive mantle convection simulation. *Geophys. J. Int.*, 192(3):889–906, 2013.

[12] Saray Busto, Maurizio Tavelli, Walter Boscheri, and Michael Dumbser. Efficient high order accurate staggered semi-implicit discontinuous Galerkin methods for natural convection problems. *Computers & Fluids*, 198, 2020.

[13] Alfredo Buttari, Markus Huber, Philippe Leleux, Théo Mary, Ulrich Ruede, and Barbara Wohlmuth. Block Low Rank Single Precision Coarse Grid Solvers for Extreme Scale Multigrid Methods. Apr 2020. Submitted.

[14] Michael A. Celia, Thomas F. Russell, Ismael Herrera, and Richard E. Ewing. An Eulerian-Lagrangian localized adjoint method for the advection-diffusion equation. *Adv. Water Resources*, 13(4):187–206, 1990.

[15] Zhangxin Chen, Guanren Huan, and Yuanle Ma. *Computational Methods for Multiphase Flows in Porous Media*. SIAM, 2006.

[16] Bernardo Cockburn and Chi-Wang Shu. The Local Discontinuous Galerkin Method for Time-Dependent Convection-Diffusion Systems. *SIAM J. Numer. Anal.*, 35(6):2440–2463, 1998.

[17] CN Dawson, TF Russell, and MF Wheeler. Some improved error estimates for the modified method of characteristics. *SIAM J. Numer. Anal.*, 26(6):1487–1512, 1989.

[18] Jim Douglas, Jr., Chieh-Sen Huang, and Felipe Pereira. The modified method of characteristics with adjusted advection. *Numer. Math.*, 83(3):353–369, 1999.

[19] Jim Douglas, Jr. and Thomas F. Russell. Numerical Methods for Convection-Dominated Diffusion Problems Based on Combining the Method of Characteristics with Finite Element or Finite Difference Procedures. *SIAM J. Numer. Anal.*, 19(5):871–885, 1982.

[20] Sebastian Eibl and Ulrich Rüde. A Modular and Extensible Software Architecture for Particle Dynamics. Submitted.

[21] Sebastian Eibl and Ulrich Rüde. A local parallel communication algorithm for polydisperse rigid body dynamics. *Parallel Comput.*, 80:36–48, 2018.

[22] Mofdi El-Amrani and Mohammed Seaïd. Eulerian-Lagrangian time-stepping methods for convection-dominated problems. *Int. J. Comput. Math.*, 85(3-4):421–439, 2008.

[23] Howard C Elman, David J Silvester, and Andrew J Wathen. *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics.* Oxford University Press, 2nd edition, 2014.

[24] R. Gassmöller, H. Lokavarapu, W. Bangerth, and E. G. Puckett. Evaluating the accuracy of hybrid finite element/particle-in-cell methods for modelling incompressible Stokes flow. *Geophys. J. Int.*, 219(3):1915–1938, 2019.

[25] B. Gmeiner, U. Rüde, H. Stengel, C. Waluga, and B. Wohlmuth. Performance and scalability of hierarchical hybrid multigrid solvers for stokes systems. *SIAM J. Sci. Comput.*, 37(2):C143–C168, 2015.

[26] Björn Gmeiner, Markus Huber, Lorenz John, Ulrich Rüde, and Barbara Wohlmuth. A quantitative performance study for Stokes solvers at the extreme scale. *J. Comput. Sci.*, 17(3):509–521, 2016.

[27] William J Gordon and Charles A Hall. Transfinite Element Methods: Blending-Function Interpolation over Arbitary Curved Element Domains. *Numer. Math.*, 21(2):109–129, 1973.

[28] Volker John and Ellen Schmeyer. Finite element methods for time-dependent convection-diffusion-reaction equations with small diffusion. *Comp. Meth. Appl. Mech. Engrg.*, 198(3-4):475–494, 2008.

[29] Nils Kohl and Ulrich Rüde. Textbook efficiency: massively parallel matrix-free multigrid for the Stokes system, 2020. Submitted.

[30] Nils Kohl, Dominik Thönnes, Daniel Drzisga, Dominik Bartuschat, and Ulrich Rüde. The HyTeG finite-element software framework for scalable multigrid solvers. *Int. J. Par., Emerg. Distrib. Sys.*, 34(5):477–496, 2019.

[31] M. Kronbichler, T. Heister, and W. Bangerth. High Accuracy Mantle Convection Simulation through Modern Numerical Methods. *Geophys. J. Int.*, 191(1):12–29, 2012.

[32] Dimitri Kuzmin, Matthias Möller, and Stefan Turek. High-resolution FEM–FCT schemes for multidimensional conservation laws. *Comput. Methods Appl. Mech. Engrg.*, 193(45-47):4915–4946, 2004.

[33] Dmitri Kuzmin. Algebraic Flux Correction I - Conservation Laws. In *Flux-Corrected Transport - Principles, Algorithms, and Applications*. Springer, 2012.

[34] Randall J. LeVeque. High-Resolution Conservative Algorithms for Advection in Incompressible Flow. *SIAM J. Numer. Anal.*, 33(2):627–665, 1996.

[35] A. V. Malevsky and D. A. Yuen. Characteristics-based methods applied to infinite Prandtl number thermal convection in the hard turbulent regime. *Phys. Fluids A*, 3(9):2105–2115, 1991.

[36] Keith W. Morton. *Numerical Solution of Convection-Diffusion Problems.* CRC Press, 2019.

[37] Pablo Ouro, Bru no Fraga, and Unai Lopez-Novoac amd Thorsten Stoesser. Scalability of an Eulerian-Lagrangian large-eddy simulation solver with hybrid MPI/OpenMP parallelisation. *Computers & Fluids*, 179:123–136, 2019.

[38] Alfioc Quarteroni and Alberto Valli. *Numerical Approximation of Partial Differential Equations.* Springer, 2008.

[39] Anne Reinarz, Dominic E Charrier, Michael Bader, Luke Bovard, Michael Dumbser, Kenneth Duru, Francesco Fambri, Alice-Agnes Gabriel, Jean-Matthieu Gallard, Sven Köppel, et al. ExaHyPE: an engine for parallel dynamically adaptive simulations of wave problems. *Computer Physics Communications*, 254:107251, 2020.

[40] Y. Ricard. Physics of Mantle Convection. In David Bercovici, editor, *Mantle Dynamics*, volume 7 of *Treatise on Geophysics*, pages 31–89. Elsevier, 2007.

[41] Johann Rudi, A. Cristiano I. Malossi, Tobin Isaac, Georg Stadler, Michael Gurnis, Peter W. J. Staar, Yves Ineichen, Costas Bekas, Alessandro Curioni, and Omar Ghattas. An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '15, pages 5:1–5:12. ACM, 2015.

[42] Thomas F. Russell and Michael A. Celia. An overview of research on Eulerian–Lagrangian localized adjoint methods (ELLAM). *Adv. Water Resources*, 25(8-12):1215–1231, 2002.

[43] Abdoulaye Samaké, Pierre Rampal, Sylvain Bouillon, and Einar Ólason. Parallel implementation of a Lagrangian-based model on an adaptive mesh in C++: Application to sea-ice. *J, Comput. Phys.*, 350:84–96, 2017.

[44] Gilbert Strang. On the construction and comparison of difference schemes. *SIAM J. Numer. Anal.*, 5(3):506–517, 1968.

[45] Maurizio Tavelli and Walter Boscheri. A high-order parallel Eulerian-Lagrangian algorithm for advection-diffusion problems on unstructured meshes. *Int. J. Numer. Meth. Fluids*, 91:332–347, 2019.

[46] TerraNeo Team. HyTeG, 2021. SHA: 3e5a93130cc9b86143cfa5249351914ee99cd6e3.

[47] Ulrich Trottenberg, Cornelius Oosterlee, and Anton Schüller. *Multigrid*. Academic Press, 2001.

[48] Arie P. van den Berg, Peter E. van Keken, and David A. Yuen. The effects of a composite non-Newtonian and Newtonian rheology on mantle convection. *Geophys. J. Int.*, 115(1):62–78, 1993.

[49] L. Vynnytska, M. E. Rognes, and S. R. Clark. Benchmarking FEniCS for mantle convection simulations. *Comp. & Geosci.*, 50:95–105, 2013.

[50] Christian Waluga, Barbara Wohlmuth, and Ulrich Rüde. Mass-corrections for the conservative coupling of flow and transport on collocated meshes. *J. Comput. Phys.*, 305:319–332, 2016.

[51] Steven T. Zalesak. Fully multidimensional flux-corrected transport algorithms for fluids. *J. Comput. Phys.*, 31(3):335–362, 1979.