Slope limiting the velocity field in a discontinuous Galerkin divergence free two-phase flow solver

Tormod Landet[*], Kent-Andre Mardal and Mikael Mortensen
Matematisk Institutt, Moltke Moes vei 35, 0851 Oslo, Norway

March 20, 2018

Abstract

Solving the Navier-Stokes equations when the density field contains a large sharp discontinuity—such as a water/air free surface—is numerically challenging. Convective instabilities cause Gibbs oscillations which quickly destroy the solution. We investigate the use of slope limiters for the velocity field to overcome this problem in a way that does not compromise on the mass conservation properties. The equations are discretised using the interior penalty discontinuous Galerkin finite element method that is divergence free to machine precision.

A slope limiter made specifically for exactly divergence free (solenoidal) fields is presented and used to illustrated the difficulties in obtaining convectively stable fields that are also exactly solenoidal. The lessons learned from this are applied in constructing a simpler method based on the use of an existing scalar slope limiter applied to each velocity component.

We show by numerical examples how both presented slope limiting methods are vastly superior to the naive non-limited method. The methods can solve difficult two-phase problems with high density ratios and high Reynolds numbers—typical for marine and offshore water/air simulations—in a way that conserves mass and stops unbounded energy growth caused by the Gibbs phenomenon.

Keywords: DG FEM; divergence free; solenoidal; Navier-Stokes; two-phase; slope limiter; Gibbs oscillations; density jump

1 Introduction

The incompressible and variable density Navier-Stokes equations for the unknown velocity $\boldsymbol{u}$ and pressure p, with gravity $\boldsymbol{g}$ and spatially varying density ρ and viscosity μ,

$$\rho\left(\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u}\cdot\nabla)\boldsymbol{u}\right) = \nabla\cdot\mu\left(\nabla\boldsymbol{u} + (\nabla\boldsymbol{u})^T\right) - \nabla p + \rho\boldsymbol{g}, \tag{1}$$

$$\nabla\cdot\boldsymbol{u} = 0, \tag{2}$$

$$\frac{\partial \rho}{\partial t} + \boldsymbol{u}\cdot\nabla\rho = 0, \tag{3}$$

---

[*]Corresponding author. Email: tormodla@math.uio.no

are used in applications where the fluid velocity is much smaller than the Mach number and where gravity effects are important such as for internal and surface gravity waves. Unfortunately, numerical problems will immediately occur when using equations (1) to (3) to study air/water free surface physics with a higher order solution method. These problems are due to the sharp factor 1000 jump in momentum across the liquid/gas interface which causes Gibbs oscillations in the velocity field, even though the true velocity has no discontinuity at the interface. The energy in the velocity field will eventually blow up to destroy the solution if this non-linear convective instability is not handled with care.

The convergence of numerical approximations to the Navier-Stokes equations with variable and potentially discontinuous density and viscosity was studied by Liu and Walkington (2007). They show that if the continuous problem has an unique solution then a stable discontinuous Galerkin (DG) discretisation with a piecewise constant density field will converge to that solution. This work presents such a stable DG discretisation. Stable fractional step methods for solving the equations have been studied by Guermond and Quartapelle (2000); Pyo and Shen (2007); Guermond and Salgado (2009). Our results are computed using a direct solver on the coupled velocity and pressure system in order to eliminate any fractional step splitting errors from influencing the conclusions.

In this work we will use a piecewise constant density field and the volume of fluid, VOF, method by Hirt and Nichols (1981) for evolving this density field in time while maintaining a sharp interface. In the VOF method a transport equation for the density is solved for a normalised so called *colour function*, $C \in [0, 1]$, which is linearly related to the density. Among the most commonly used variations are the algebraic VOF schemes CICSAM (Ubbink, 1997) and HRIC (Muzaferija et al., 1998). The level set method (Osher and Sethian, 1988) is another alternative for tracking free surfaces. Modern versions exist that significantly improve on the historical mass conservation problems of the level set method—see, e.g., Olsson and Kreiss (2005); Touré, Fahsi, and Soulaïmani (2016).

For most industrial applications the above methods are implemented in solvers based on the finite volume method where it is possible to ensure convective stability by appropriate use of flux limiters, see, e.g., the illustrative discussions and diagrams of Sweby (1984) and Leonard (1988). Flux limiters based on TVD or ENO/WENO schemes are applied in the solution of the momentum equation, and some sort of stable interface sharpening scheme such as CICSAM or HRIC is used for the flux of density. Finite volume methods can be partially extended to higher order by use of larger geometrical stencils, but this approach is not trivial to implement when using irregular grids, so in practice it is common that only immediate neighbours are used and most schemes are hence low order.

For irregular geometries there are two major research directions that look to enable higher order methods: immersed boundary methods and finite element methods. Immersed boundary methods use regular background grids and perform special treatment of grid cells near or inside objects embedded in the computational domain (Peskin, 2002). Finite element methods use basis functions with local support to enable high order approximating functions on irregular meshes. This work is based on a discontinuous Galerkin (DG) finite element method (FEM). DG-FEM has two main advantages over continuous Galerkin methods, the ease of using an upwind flux limiter for stabilising the linear convective instabilities and the option to create exactly divergence free and mass

conserving numerical schemes (Cockburn, Kanschat, and Schötzau, 2004, 2005). There is one important drawback, and that is the significantly increased number of degrees of freedom. A variation of the scheme from Cockburn, Kanschat, and Schötzau (2005) is used, where instead of using a local discontinuous Galerkin, LDG, treatment of the elliptic term, a symmetric interior penalty, SIP, treatment (Arnold, 1982) is employed.

When using higher order approximating polynomials it is no longer sufficient to use only a flux limiter to obtain convective stability—a slope limiter must also be included in the method (Cockburn and Shu, 1998, 2001; Kuzmin, 2010). The flux limiter ensures that the cell average values are bounded. When using higher order basis functions, the solution can go out of bounds in localized regions due to steep gradients inside each cell. A slope limiter prevents this by flattening steep slopes near discontinuities while leaving the solution untouched near smooth maxima so that the method's high order accuracy is retained. While a flux limiter is an implicit part of the equation system and has its stabilising effect included in the results from the linear equation solver, a slope limiter is typically applied to the resulting function as an explicit post-processing operator. Other options could be to use non-linear diffusion to combat the non-linear instability, or to apply spectral filtering, see, e.g., Michoski et al. (2016); Zingan et al. (2013).

This paper starts with a description of the Gibbs instability and the discontinuous Galerkin method employed for solving the variable density Navier-Stokes equations in sections 2.1 to 2.4. Slope limiting is then presented; first the hierarchical Taylor based slope limiter by Kuzmin (2010, 2013) in section 3.1, and then the possibility of constructing a vector field slope limiter that leaves the resulting velocity field both solenoidal and free from local maxima is explored in section 3.2. After showing that it is likely not possible to obtain a single field that is both solenoidal and stable, separate limiting of the convecting and the convected velocity fields is introduced. Two alternatives are presented in section 3.3; both ensure solenoidal convecting velocities while keeping convective stability. Readers familiar with DG methods and slope limiting can start at section 3.2, though the presentation builds directly on the preceding methods, so some referring back may be needed. Results from numerical tests are shown in section 4 and discussion and concluding remarks can be found in section 5.

2 The numerical method

2.1 Instabilities

The most common numerical instabilities related to handling of large density jumps are illustrated in figures 1a to 1c. A 'block' of water starts at rest in a box filled with air. Already in the very first time step (figure 1a) the solution will start to blow up if one does not either apply smoothing to the density field or stabilise the numerical scheme. Smoothing the initial field is not sufficient as the interface may pinch and become non smooth (figure 1b), so continuous smoothing of the density field is necessary if this approach is selected. A level set method is a natural way to implement a smoothed density field, see, e.g., Unverdi and Tryggvason (1992); Sussman, Smereka, and Osher (1994). In this work we will not apply stabilisation through smoothing. Our aim is to decouple the stability of the method from the treatment of the density field.

Figure 1c illustrates a situation where the solenoidal properties of $\boldsymbol{u}$ are important.

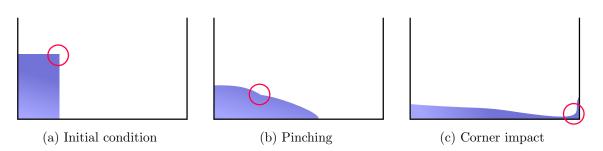| (a) Initial condition | (b) Pinching | (c) Corner impact |

Figure 1: Illustration of typical numerical problems

A divergence free convecting velocity is required for the stability of the density transport equation (3) and any divergence quickly becomes a problem in difficult situations such as the corner impact where in our experience either mass loss or unbounded densities will occur if care is not taken to ensure that the convecting velocity used for density transport is solenoidal.

2.2 Preliminaries

The notation used in this paper is relatively standard. When looking at a facet between two cells (finite elements) we will denote one of the cells K^+ and the other K^- in an arbitrary, but repeatable manner. Function values in each cell will be given the same superscripts to distinguish between the values on opposite sides of the facet in the discontinuous space. The average and jump operators across an internal facet are defined as

$$\{\{u\}\} = \frac{1}{2}(u^+ + u^-), \tag{4}$$

$$[\![u]\!] = u^+ - u^-, \tag{5}$$

$$[\![\boldsymbol{u}]\!]_{\boldsymbol{n}} = \boldsymbol{u}^+ \cdot \boldsymbol{n}^+ + \boldsymbol{u}^- \cdot \boldsymbol{n}^-. \tag{6}$$

Vector terms such as the velocity $\boldsymbol{u}$ and gravity $\boldsymbol{g}$ will be denoted with a bold font while scalars such as the pressure p, the density ρ, and the dynamic viscosity μ are shown in italics. Projection operators are written in blackboard bold. An example is the projection $\mathbb{D}$ into a solenoidal vector space. Sets are written in a calligraphic typeface. Fluxes of a quantity, used in DG facet integrals, are marked with a circumflex accent over the quantity, e.g., $\hat{\boldsymbol{u}}$, $\hat{p}$. An additional superscript is added when different definitions of the flux are used in different parts of the weak form, e.g., $\hat{\boldsymbol{u}}^{\boldsymbol{w}}$ and $\hat{\boldsymbol{u}}^{p}$, the flux of velocity related to convection and the flux of velocity related to the incompressibility constraint.

Let ∂K denote the boundary of element K. Γ_I is the inlet portion of the domain boundary $\partial \Omega$ where $\boldsymbol{u} \cdot \boldsymbol{n} < 0$ for an outwards pointing normal, $\boldsymbol{n}$. The set of all grid cells is $\mathcal{T}$ (the tessellation), and the set of all facets is $\mathcal{S}$ (the mesh skeleton). The set of outside facets is $\mathcal{S}_O = \mathcal{S} \cap \partial \Omega$ and the set of inside facets shared between two cells is $\mathcal{S}_I = \mathcal{S} \setminus \mathcal{S}_O$.

For a facet on the boundary, $F_b \in \mathcal{S}_O$, let the boundary cell be denoted K^+ with regards to the facet and hence $\boldsymbol{n}^+ \cdot [\![\boldsymbol{v}]\!] = \boldsymbol{n}^+ \cdot \boldsymbol{v}^+ = \boldsymbol{n} \cdot \boldsymbol{v}$ on F_b. We set terms related to K^- to zero and also set the average to be the K^+ value $\{\{\boldsymbol{v}\}\} = \boldsymbol{v}$ on $\mathcal{S}_O$.

Let $P_k(K)$ denote the space of polynomials of order k on an element and $P_k(F)$ denote the space of polynomials of order k on a facet. The basis functions in $P_k(K)$ are discontinuous across facets and functions in $P_k(F)$ are discontinuous across edges (vertices in 2D). Let the superscript $\bullet^{n+1}$ denote a value at time step $t = \Delta t(n + 1)$. For nabla the conventions $(\nabla \boldsymbol{u})_{ij} = \partial_j \boldsymbol{u}_i$ and $(\nabla \cdot \boldsymbol{\sigma})_i = \partial_j \sigma_{ij}$ are used.

2.3 Discretisation

We will approximate the unknown functions in space by using discontinuous Lagrange polynomial function spaces with polynomial order $k = 2$ for the velocity in $d = 2$ spatial dimensions. Let the Galerkin test functions for $\{\boldsymbol{u}, p, \rho\}$ be denoted respectively $\{\boldsymbol{v}, q, r\}$. The discontinuous function spaces are not restricted at the boundaries—all boundary conditions will be imposed weakly—so the trial and test functions share spaces,

$$
\begin{aligned}
\boldsymbol{u}, \boldsymbol{v} &\in [P_k(K)]^d, \\
p, q &\in P_{k-1}(K), \\
\rho, r &\in P_0(K).
\end{aligned}
\tag{7}
$$

2.3.1 Two-phase density transport

The transport equation (3) for the density is modified by the introduction of a solenoidal convecting velocity field $\boldsymbol{w}$ which is close to $\boldsymbol{u}$, but may not be identical as will be explained in sections 2.4 and 3.3. Boundary conditions are needed only on the inlet. The strong form can then be written

$$
\frac{\partial \rho}{\partial t} + \boldsymbol{w} \cdot \nabla \rho = 0 \qquad \text{in } \Omega,
\tag{8}
$$

$$
\rho = \rho_I \qquad \text{on } \Gamma_I.
$$

In VOF methods the fluid properties are expressed in terms of an indicator function, $C \in [0, 1]$. The true density and viscosity fields can easily be recovered from C when the density and kinematic viscosity properties of the two fluids are known,

$$
\rho = C\rho_{\text{water}} + (1 - C)\rho_{\text{air}},
\tag{9}
$$

$$
\mu = [C\nu_{\text{water}} + (1 - C)\nu_{\text{air}}]\,\rho.
\tag{10}
$$

In order to compute C the transport equation (8) for the density is modified by inserting equation (9). Now the solution $C^{n+1} \in P_0(K)$ can be found by expressing the resulting equation on weak form and integrating by parts,

$$
\int_{\mathcal{T}} \frac{1}{\Delta t}(\gamma_1 C^{n+1} + \gamma_2 C^n + \gamma_3 C^{n-1})r \, \mathrm{d}\boldsymbol{x}
\tag{11}
$$

$$
- \int_{\mathcal{T}} C^{n+1}\boldsymbol{w} \cdot \nabla r \, \mathrm{d}\boldsymbol{x} + \int_{\mathcal{S}} \hat{C}^{n+1}\,\boldsymbol{w} \cdot \boldsymbol{n}^+ \, [\![\, r \,]\!] \, \mathrm{d}s = 0,
$$

where we have assumed that the convecting velocity $\boldsymbol{w}$ is H^{div}-conforming such that the flux is continuous across facets, $[\![\, \boldsymbol{w} \,]\!]_{\boldsymbol{n}} = 0$.

A second order backwards differencing formulation, BDF2, is used for time integration. The parameters are $\{\gamma_1, \gamma_2, \gamma_3\} = \{3/2, -2, 1/2\}$. The BDF2 method is monotonicity preserving when started using a backward Euler step (Hundsdorfer, Ruuth, and Spiteri, 2003), though the time step required to preserve monotonicity is half that of backward Euler. The benefit is that the method is second order—and small time steps are anyhow required to keep the interface sharp (Muzaferija et al., 1998). Second order extrapolation is used for the convective velocity,

$$\boldsymbol{w} = 2\boldsymbol{w}^n - \boldsymbol{w}^{n-1}, \tag{12}$$

which makes ρ^{n+1} independent of the computed velocity at time $t = (n+1)\Delta t$. The density transport equation is hence uncoupled from the momentum equation.

For the density flux, $\hat{C}^{n+1}$, the most stable choice is to take the upwind value, which means that the boundary condition $\hat{\rho} = \rho_I$ is used to determine $\hat{C}^{n+1}$ on inlet facets, $\mathcal{S} \cap \Gamma_I$. For the internal facets the term related to the flux can be written on upwind form as

$$\hat{C}^{n+1} \boldsymbol{w} \cdot \boldsymbol{n}^+ = \left[\!\left[C^{n+1} \frac{1}{2} (\boldsymbol{w} \cdot \boldsymbol{n} + |\boldsymbol{w} \cdot \boldsymbol{n}|) \right]\!\right], \tag{13}$$

and by replacing the plus by a minus in equation (13) the downwind flux can similarly be computed. Using the upwind and downwind fluxes on each facet a downwind-blended compressive interface flux can be applied to the density transport such as CICSAM (Ubbink, 1997) or HRIC (Muzaferija et al., 1998). Such blended fluxes sharpen the interface between the two fluid layers, but remain convectively stable unlike downwind or central fluxes. Both CICSAM and HRIC are algebraic VOF methods which define facet-wise blending factors to combine the upwind and downwind fluxes into one linearly stable flux. The results in section 4 are computed using the HRIC method.

One important note is that standard VOF flux limiters from finite volume methods ensure that the C field remains sharp and bounded based on a convecting velocity field that is piecewise constant on each facet. Such a field can easily be computed from $\boldsymbol{w}$ and as long as $\boldsymbol{w}$ is solenoidal, so is the piecewise constant field. It is this field that is used when solving for C^{n+1} since $\boldsymbol{w}$ is only needed on the facets—for $r \in P_0(K)$ the volume integral term containing $\boldsymbol{w}$ in equation (11) is identically zero.

2.3.2 The variable density Navier-Stokes equations

The strong form of the variable density Navier-Stokes equations, where the convecting velocity is replaced by $\boldsymbol{w}$, can be written

$$\rho \left(\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{w} \cdot \nabla)\boldsymbol{u} \right) = \nabla \cdot \mu(\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^T) - \nabla p + \rho \boldsymbol{g} \qquad \text{in } \Omega, \tag{14}$$

$$\nabla \cdot \boldsymbol{u} = 0 \qquad \text{in } \Omega,$$

$$\boldsymbol{u} = \boldsymbol{u}_D \qquad \text{on } \Gamma_D,$$

$$\frac{\partial \boldsymbol{u}}{\partial n} = \boldsymbol{a} \qquad \text{on } \Gamma_N,$$

where Γ_D and Γ_N are the parts of the boundary where Dirichlet and Neumann boundary conditions are applied respectively; $\Gamma_N = \partial\Omega \setminus \Gamma_D$. Dirichlet boundary conditions can

be enforced on external facets due to the elliptic viscosity term, $\nabla \cdot \mu(\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^T)$. The Navier-Stokes equations will be written on weak form, see equation (19), but first we will briefly discuss the symmetric interior penalty (SIP) method. First, the boundary condition, $\boldsymbol{u} = \boldsymbol{u}_D$, is written on weak form on an external facet F,

$$\int_F \boldsymbol{u} \cdot \boldsymbol{v} \, \mathrm{d}s = \int_F \boldsymbol{u}_D \cdot \boldsymbol{v} \, \mathrm{d}s, \tag{15}$$

and then—using the stabilisation scheme proposed by Nitsche (1971)—the test function $\boldsymbol{v}$ is replaced by a Petrov-Galerkin test function $\tilde{\boldsymbol{v}} = \kappa_\mu \boldsymbol{v} - \mu \left(\nabla \boldsymbol{v} + (\nabla \boldsymbol{v})^T\right) \cdot \boldsymbol{n}$. This method is extended to enforce continuity across internal facets which is necessary for stability (Arnold, 1982). This gives

$$\int_F \kappa_\mu \, [\![\boldsymbol{u}]\!] \cdot [\![\boldsymbol{v}]\!] \, \mathrm{d}s - \int_F (\{\!\{ \mu \left(\nabla \boldsymbol{v} + (\nabla \boldsymbol{v})^T\right) \}\!\} \cdot \boldsymbol{n}^+) \cdot [\![\boldsymbol{u}]\!] \, \mathrm{d}s = 0. \tag{16}$$

where κ_μ is a penalty parameter which must be sufficiently large to ensure stability. The analyses in Epshteyn and Rivière (2007) and Shahbazi, Fischer, and Ethier (2007) guide us in defining the penalty parameter as a function of the minimum and maximum dynamic viscosities, $\mu_{\min}$ and $\mu_{\max}$, the order k of the approximating polynomials, and the surface area S_K and volume V_K of each cell K,

$$\kappa_\mu = 3 \frac{\mu_{\max}^2}{\mu_{\min}} k(k+1) \max_K \left(\frac{S_K}{V_K}\right). \tag{17}$$

The same scheme is used for the left-hand side of the momentum equation as for the density transport equation, but a pure upwind flux is used without any downwind blending for the convective term. To avoid overloading the notation we now drop the $\bullet^{n+1}$ superscript on the unknown quantities. The upwind flux related to convection can then be written

$$\hat{\boldsymbol{u}}^{\boldsymbol{w}} \, \boldsymbol{w} \cdot \boldsymbol{n}^+ = \left[\!\left[\boldsymbol{u} \frac{1}{2} (\boldsymbol{w} \cdot \boldsymbol{n} + |\boldsymbol{w} \cdot \boldsymbol{n}|) \right]\!\right]. \tag{18}$$

Both the pressure gradient and the viscosity on the right-hand side of the momentum equation are integrated by parts. The resulting weak form is a direct combination of the LDG Navier-Stokes method by Cockburn, Kanschat, and Schötzau (2005) and the SIP diffusion method by Arnold (1982). Both these references contain more details and proofs of stability. The stability of the convective and the diffusive terms are not interconnected, so replacing the LDG elliptic operator with the SIP version is unproblematic. The treatment of the momentum transport and the inter-cell continuity, both described above, are easily recognisable in the complete weak form,

$$\int_{\mathcal{T}} \frac{\rho}{\Delta t} (\gamma_1 \boldsymbol{u} + \gamma_2 \boldsymbol{u}^n + \gamma_3 \boldsymbol{u}^{n-1}) \boldsymbol{v} \, \mathrm{d}\boldsymbol{x} \tag{19}$$

$$- \int_{\mathcal{T}} \boldsymbol{u} \cdot \nabla \cdot (\rho \boldsymbol{v} \otimes \boldsymbol{w}) \, \mathrm{d}\boldsymbol{x} \; + \; \int_{\mathcal{S}} \boldsymbol{w} \cdot \boldsymbol{n}^+ \, \hat{\boldsymbol{u}}^{\boldsymbol{w}} \cdot [\![\rho \boldsymbol{v}]\!] \, \mathrm{d}s$$

$$+ \int_{\mathcal{T}} \mu \left(\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^T\right) : \nabla \boldsymbol{v} \, \mathrm{d}\boldsymbol{x} \; + \; \int_{\mathcal{S}_I} \kappa_\mu \, [\![\boldsymbol{u}]\!] \cdot [\![\boldsymbol{v}]\!] \, \mathrm{d}s$$

$$- \int_{\mathcal{S}} (\{\{\mu\left(\nabla\boldsymbol{u} + (\nabla\boldsymbol{u})^T\right)\}\} \cdot \boldsymbol{n}^+) \cdot [\![\boldsymbol{v}]\!] \, \mathrm{d}s$$

$$- \int_{\mathcal{S}_I} (\{\{\mu\left(\nabla\boldsymbol{v} + (\nabla\boldsymbol{v})^T\right)\}\} \cdot \boldsymbol{n}^+) \cdot [\![\boldsymbol{u}]\!] \, \mathrm{d}s$$

$$- \int_{\mathcal{T}} p\,\nabla\cdot\boldsymbol{v}\,\mathrm{d}\boldsymbol{x} \; + \; \int_{\mathcal{S}} \hat{p}\,\boldsymbol{n}^+\cdot[\![\boldsymbol{v}]\!]\,\mathrm{d}s \;\; = \;\; \int_{\mathcal{T}} \rho\,\boldsymbol{g}\,\mathrm{d}\boldsymbol{x}.$$

where the flux of pressure is taken as $\hat{p} = \{\{p\}\}$.

The continuity equation (2) is also integrated by parts using $\hat{\boldsymbol{u}}^p = \{\{\boldsymbol{u}\}\}$ as the flux related to the incompressibility constraint,

$$\int_{\mathcal{S}} \hat{\boldsymbol{u}}^p \cdot \boldsymbol{n}^+ \, [\![q]\!] \, \mathrm{d}s \; - \; \int_{\mathcal{T}} \boldsymbol{u}\cdot\nabla q \, \mathrm{d}\boldsymbol{x} \;\; = \;\; 0. \tag{20}$$

Dirichlet boundaries On the inflow part of the Dirichlet boundary take $\hat{\boldsymbol{u}}^w = \boldsymbol{u}_D$, and on the outflow part take $\hat{\boldsymbol{u}}^w = \boldsymbol{u}$, i.e., the upwind values are used. On the whole Dirichlet boundary let $\hat{\boldsymbol{u}}^p = \boldsymbol{u}_D$ and $\hat{p} = p$. The viscous penalty and symmetrisation terms from equation (16) can be used also on the domain boundary by replacing $\boldsymbol{u}^-$ by $\boldsymbol{u}_D$. As noted in the references used to define κ_μ above, the best choice is to use twice the amount of penalisation on external facets compared to the interior facets. This leads to external boundary integrals

$$\int_{\Gamma_D} 2\kappa_\mu(\boldsymbol{u} - \boldsymbol{u}_D)\cdot\boldsymbol{v}\,\mathrm{d}s - \int_{\Gamma_D} (\mu\nabla\boldsymbol{v}\cdot\boldsymbol{n})\cdot(\boldsymbol{u} - \boldsymbol{u}_D)\,\mathrm{d}s = 0. \tag{21}$$

Neumann boundaries For the Neumann boundaries we take $\hat{\boldsymbol{u}}^w = \boldsymbol{u}$, $\hat{\boldsymbol{u}}^p = \boldsymbol{u}$ and $\hat{p} = p$. The extra viscous terms for symmetrisation and penalty are removed and only the normal integration by parts terms are left. The surface term becomes

$$- \int_{\Gamma_N} \mu\,\boldsymbol{a}\cdot\boldsymbol{v}\,\mathrm{d}s. \tag{22}$$

Pure Neumann boundary conditions will not be used, but we will use Dirichlet for one velocity component and Neumann for the other to implement free-slip boundary conditions on planes that are parallel to the axes. The definitions above can easily be split into component-wise treatment of boundary conditions.

Solution algorithm

The Navier-Stokes equations and the density transport equation are solved in a decoupled manner for each time step in the following order:

1. Compute an explicit convecting velocity $\boldsymbol{w}^{n+1}$ by use of equation (12).

2. Find $C^{n+1} \in P_0(K)$ such that equation (11) is satisfied.

3. Compute the density ρ^{n+1} and viscosity μ^{n+1} by use of equations (9) and (10).

4. Use the computed coefficient fields $\boldsymbol{w}^{n+1}$, μ^{n+1} and ρ^{n+1} to find $\boldsymbol{u}^{n+1} \in [P_k(K)]^d$ and $p^{n+1} \in P_{k-1}(K)$ from equations (19) and (20).

2.4 H^{div} projection of the velocity field

In finite element methods for solving the Navier-Stokes equations it is common to impose the incompressibility criterion, $\nabla \cdot \boldsymbol{u} = 0$, weakly by multiplying with a scalar test function, q, and integrating over the domain. This term, $\int_\Omega \nabla \cdot \boldsymbol{u}\, q\,\mathrm{d}\boldsymbol{x}$, will appear directly in a coupled solver and as the right-hand side in the Poisson equation for the pressure in a pressure correction fractional step scheme such as the commonly used incremental pressure correction scheme.

Imposing the incompressibility criterion weakly in the space of the pressure is sufficient for stability, but one must require exact incompressibility to locally conserve mass and momentum. Below we explain the core ideas of the method presented in Cockburn, Kanschat, and Schötzau (2005). By using this method the resulting divergence will be zero almost to machine precision, approximately 10^{-13} on each cell in our tests. We calculate the cell-wise error by computing the integrated absolute value of the divergence internally in each cell, $\int_K |\nabla \cdot \boldsymbol{u}|\,\mathrm{d}\boldsymbol{x}$, and add to it the error in flux continuity between cells on each connected facet, $\int_{\partial K} |[\![\,\boldsymbol{u} \cdot \boldsymbol{n}\,]\!]|\,\mathrm{d}s$.

Let us start by noting that H^{div}-conforming finite elements exist, and that such elements of a given polynomial order will be subspaces of the fully discontinuous elements of the same order with the same cell geometry. Such elements impose continuity of normal fluxes across facets and hence have fewer global degrees of freedom. We follow Cockburn, Kanschat, and Schötzau (2005) and define a projection from our fully discontinuous velocity $\boldsymbol{u}$ into a velocity $\boldsymbol{w} = \mathbb{P}\boldsymbol{u}$ that exist in a space of polynomials that are consistent with the H^{div}-conforming elements.

The projection operator $\boldsymbol{w} = \mathbb{P}\boldsymbol{u}$ is defined in a cell-wise manner. This local projection is hence very fast and consists of finding $\boldsymbol{w} \in [P_k(K)]^d$ such that

$$\int_F \boldsymbol{w} \cdot \boldsymbol{n}\, v_1\,\mathrm{d}s = \int_F \hat{\boldsymbol{u}}^p \cdot \boldsymbol{n}\, v_1\,\mathrm{d}s \qquad \forall v_1 \in P_k(F), F \in \partial K, \qquad (23)$$

$$\int_K \boldsymbol{w} \cdot \boldsymbol{v}_2\,\mathrm{d}x = \int_K \boldsymbol{u} \cdot \boldsymbol{v}_2\,\mathrm{d}x \qquad \forall \boldsymbol{v}_2 \in \boldsymbol{N}_{k-1}(K). \qquad (24)$$

The first equation ensures the continuity of the normal velocities across each facet. The continuity stems from using a single valued flux which is here $\hat{\boldsymbol{u}}^p = \{\!\{\boldsymbol{u}\}\!\}$ on internal facets and $\hat{\boldsymbol{u}}^p = \boldsymbol{u}_D$ on external facets ($\hat{\boldsymbol{u}}^p = \boldsymbol{u}$ on Γ_N). This flux is consistent for continuous velocity fields.

In equation (24) the space $\boldsymbol{N}_{k-1}(K)$ is the Nédélec $H(\mathrm{curl})$ element of the first kind of order $k-1$, see, e.g., Kirby et al. (2012); Nédélec (1986). The dimension of the Brezzi-Douglas-Marini (BDM) element $P_k(F) \times \boldsymbol{N}_{k-1}(K)$ is the same as that of the Discontinuous Lagrange DGk element on each cell (Brezzi and Fortin, 1991). It is hence possible to form square projection matrices between the two spaces in each cell. The projection $\mathbb{P}\boldsymbol{u}$ is not square globally since the degrees of freedom related to v_1 are shared between exactly two cells on all internal facets. This leads to the test space having fewer global degrees of freedom than the fully discontinuous trial space.

The properties of this BDM-like projection used as a velocity post-processing step in a discontinuous Galerkin method is given by Cockburn, Kanschat, and Schötzau (2005). The projection gives both a continuous flux $\boldsymbol{w} \cdot \boldsymbol{n}$ on all inter-element facets and also ensures that the total flux across each individual cell's facets is zero. Fulfilling these two criteria is what we mean by exact incompressibility.

3 Slope limiting

The numerical diffusion due to upwinding in the spatial DG scheme is sufficient to stabilise convective operators and avoid spurious oscillations when piecewise constant approximating functions are employed. One way to stabilise convective operators when using higher order approximating functions is to employ slope limiters (Cockburn and Shu, 1998).

3.1 The hierarchical Taylor polynomial based slope limiter

The hierarchical vertex based slope limiter by Kuzmin (2010, 2013) is used to remove high frequency oscillations near discontinuities when solving an equation containing a convective term for a scalar quantity ϕ. This slope limiter is based on using discontinuous Taylor function spaces. We are using Lagrange polynomials, so the first step is to project the unknown function ϕ from the discontinuous Lagrange function space to a function ϕ_t in the discontinuous Taylor function space. This projection $\phi_t = \mathbb{T}\phi$ is local to each cell and can be applied and inverted exactly by a single matrix vector product in each cell, so the cost of converting back and forth is negligible.

The discontinuous Taylor function space is slightly altered from the standard definition by use of cell averages. The expansion

$$
\phi_t(x,y) = \bar{\phi} + \left.\frac{\partial \phi}{\partial x}\right|_c \alpha_1 (x - x_c) + \left.\frac{\partial \phi}{\partial y}\right|_c \alpha_1 (y - y_c) + \tag{25}
$$
$$
\left.\frac{\partial^2 \phi}{\partial x^2}\right|_c \alpha_2 \left[\frac{(x - x_c)^2}{2} - \frac{\overline{(x - x_c)^2}}{2}\right] + \left.\frac{\partial^2 \phi}{\partial y^2}\right|_c \alpha_2 \left[\frac{(y - y_c)^2}{2} - \frac{\overline{(y - y_c)^2}}{2}\right] +
$$
$$
\left.\frac{\partial^2 \phi}{\partial x \partial y}\right|_c \alpha_2 \left[\frac{(x - x_c)(y - y_c)}{2} - \frac{\overline{(x - x_c)(y - y_c)}}{2}\right]
$$

is used to describe a second order polynomial function on a triangle where the six coefficients are $\{\bar{\phi}, \partial \phi/\partial x|_c, \cdots, \partial^2 \phi/\partial x \partial y|_c\}$. The number of coefficients is the same as the number of nodes in a second order Lagrangian function space on a triangle which ensures that the transformation matrix resulting from $\mathbb{T}$ is square on each cell. The same is also true for tetrahedra and for higher and lower polynomial orders.

Over-lined terms such as $\bar{\phi}$ and $\overline{(x - x_c)^2/2}$ denote cell averages and the restriction $\cdot|_c$ signifies evaluation in the cell centre. Slope limiters $\alpha_i \in [0,1]$ will be applied to the i'th derivative terms as shown in equation (25) using the same α-factors for all derivatives of the same order. How to compute α_i is described below. As long as no slope limiting is performed ($\alpha_i = 1$), we have $\phi = \mathbb{T}^{-1}\mathbb{T}\phi$.

To find α_i we follow Kuzmin (2013) and define three bi-linear functions on each cell,

$$
\tilde{\phi}_0(x,y) = \bar{\phi} + \tilde{\alpha}_1 \left[\left.\frac{\partial \phi}{\partial x}\right|_c (x - x_c) + \left.\frac{\partial \phi}{\partial y}\right|_c (y - y_c)\right], \tag{26}
$$

$$
\tilde{\phi}_x(x,y) = \left.\frac{\partial \phi}{\partial x}\right|_c + \tilde{\alpha}_{2x} \left[\left.\frac{\partial^2 \phi}{\partial^2 x}\right|_c (x - x_c) + \left.\frac{\partial^2 \phi}{\partial x \partial y}\right|_c (y - y_c)\right], \tag{27}
$$

$$
\tilde{\phi}_y(x,y) = \left.\frac{\partial \phi}{\partial y}\right|_c + \tilde{\alpha}_{2y} \left[\left.\frac{\partial^2 \phi}{\partial^2 y}\right|_c (y - y_c) + \left.\frac{\partial^2 \phi}{\partial x \partial y}\right|_c (x - x_c)\right]. \tag{28}
$$

The parameters $\tilde{\alpha}_i$ are intermediate slope limiters for the approximate function $\tilde{\phi}_i$. Equations (26) to (28) are used to linearly approximate ϕ and its first derivatives at each of the three vertices of the triangular cells. These approximate vertex values should not form local maxima or minima when compared to the cell centre values of the linear representations in the surrounding cells.

Pick any one of the above linear approximations in equations (26) to (28) and call it $\tilde{\phi}(x, y)$. The corresponding value in the cell centre is denoted $\tilde{\phi}_c = \tilde{\phi}(x_c, y_c)$. Iterate through all cells in the mesh and for each cell consider each vertex and record the extremal $\tilde{\phi}_c$ in the cells that share the vertex. This gives allowable bounds $\tilde{\phi}_i^{\min}$ and $\tilde{\phi}_i^{\max}$ at vertex i for the selected linear approximation in the given cell. For each vertex i located at (x_i, y_i) one must ensure that $\tilde{\phi}_i = \tilde{\phi}(x_i, y_i)$ is bounded by the surrounding cell values, $\tilde{\phi}_i^{\min} \leq \tilde{\phi}_i \leq \tilde{\phi}_i^{\max}$. To find the maximum admissible value of $\tilde{\alpha}$ that ensures this, first for each vertex i set

$$\tilde{\alpha}_i = \begin{cases} \min\{1, \frac{\tilde{\phi}_i^{\max} - \tilde{\phi}_c}{\tilde{\phi}_i - \tilde{\phi}_c}\} & \text{if } \tilde{\phi}_i - \tilde{\phi}_c > 0, \\ 1 & \text{if } \tilde{\phi}_i - \tilde{\phi}_c = 0, \\ \min\{1, \frac{\tilde{\phi}_i^{\min} - \tilde{\phi}_c}{\tilde{\phi}_i - \tilde{\phi}_c}\} & \text{if } \tilde{\phi}_i - \tilde{\phi}_c < 0, \end{cases} \tag{29}$$

and then take the minimum of the calculated $\tilde{\alpha}_i$-values in each vertex to find $\tilde{\alpha}_1$, $\tilde{\alpha}_{2x}$ and $\tilde{\alpha}_{2y}$ for each cell by performing the above calculations for each of the three linear approximations in equations (26) to (28). The last step is to calculate the final slope limiter coefficients for the second order derivatives,

$$\alpha_2 = \min\{\tilde{\alpha}_{2x}, \tilde{\alpha}_{2y}\}, \tag{30}$$

and, since one can expect higher regularity of the first order derivatives than the second order derivatives, take

$$\alpha_1 = \max\{\tilde{\alpha}_1, \alpha_2\}, \tag{31}$$

which at a smooth extremal point will stop any limiting from happening since $\alpha_2 = 1$ here, even if it is likely that $\tilde{\alpha}_1 < 1$.

Having found the slope limiter coefficients α_1 and α_2 for a given cell one can go back to the discontinuous Lagrange function space with the help of the projection $\mathbb{T}^{-1}$. If there were no spurious oscillations in the cell then the slope limiters coefficients should end up as $\alpha_i = 1$ and the slope limiter projection—which we call $\mathbb{S}$—will hence be an identity transform, and we get $\phi^{\text{lim}} = \mathbb{T}^{-1} \mathbb{S} \mathbb{T} \phi = \mathbb{T}^{-1} \mathbb{T} \phi = \phi$. This ensures that the order of spatial convergence is kept the same as the underlying DG scheme.

3.2 On slope limiting of solenoidal fields

Slope limiting of solenoidal vector fields is more complex than limiting scalar fields due to an increased number of invariants. For scalar fields the only invariant is that the average value in each cell must be unchanged after limiting. For vector fields the (i) divergence inside each element and the (ii) flux between neighbouring elements are new invariants in addition to the (iii) average value of each of the velocity component in each cell.

Breaking these invariants means that the result is not solenoidal (i), the slope limiting post processor is not a non-local operation (ii) or that momentum is not preserved (iii).

In the following sections we will describe how these invariants can be used to reduce the number of unknowns in the cell-wise limiting problem from 12 to only 4 for a DG2 vector field, $\boldsymbol{u} \in [P_2(K)]^2$, on a triangle. We will then describe an optimisation method which can be used to decide the value of the remaining four unknowns in a way that minimises the tendency to produce local extrema. Removing all local extrema would render the method stable. We do not claim that the below method is optimal, but it does give strong indications that the four remaining unknowns in the cell-wise limiting problem are not sufficient to fully control the local extrema. Creating a single velocity field that is both solenoidal and stable is hence likely not possible by use of a cell-wise slope limiting process.

3.2.1 A reduced basis for the solenoidal slope limiting problem

Let the cell volume be denoted V_K and the area of facet i be denoted L_i. Preserving the facet average of the flux between neighbouring cells gives a result which is H$^{\mathrm{div}}$-compatible in a DG0 sense, which is what is needed for mass conservation. The cell-wise invariants we require to be left unchanged by the slope limiting procedure are then

$$R_i = \frac{1}{L_i} \int_{F_i} \boldsymbol{u} \cdot \boldsymbol{n} \, \mathrm{d}s \ \ \forall i \in \{1, 2, 3\}, \qquad U_j = \frac{1}{V_K} \int_K \boldsymbol{u}_j \, \mathrm{d}x \ \ \forall j \in \{1, 2\}. \tag{32}$$

When applying this to a DG2 vector field on a triangle there are five invariants, three average fluxes, R_i, and two component averages, U_j. By not changing the average flux on each facet the total flux for each cell will still sum to zero, and it is still possible to have a solenoidal description of the vector field.

The 12 degrees of freedom in a DG2 vector field on a triangle can be reduced. The divergence free solution will never span the full Lagrange space, but be restricted to the solenoidal subspace which has only 9 degrees of freedom, see e.g. Baker, Jureidini, and Karakashian (1990). This space is spanned by

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}; \quad \begin{pmatrix} y \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ x \end{pmatrix}, \begin{pmatrix} x \\ -y \end{pmatrix}; \quad \begin{pmatrix} y^2 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ x^2 \end{pmatrix}, \begin{pmatrix} x^2 \\ -2xy \end{pmatrix}, \begin{pmatrix} -2xy \\ y^2 \end{pmatrix}; \tag{33}$$

which is here shown grouped into the two zeroth order, three first order and four second order vector valued polynomials.

A projection from the Lagrangian description of $\boldsymbol{u}$ to the solenoidal description $\boldsymbol{u}_s$ can be implemented as $\boldsymbol{u}_s = \mathbb{D}\boldsymbol{u}$ on a cell by cell level. To do this, first define local x and y coordinates with the origin in the centre of each cell and then form the cell-wise rectangular inverse projection $\boldsymbol{u} = \mathbb{D}^{-1}\boldsymbol{u}_s$ by evaluating the solenoidal polynomials from equation (33) in the location of the Lagrangian nodes. The least squares pseudo-inverse of this matrix is then $\mathbb{D}$. It is important to note that in our simulations the projection is lossless such that $\boldsymbol{u} = \mathbb{D}^{-1}\mathbb{D}\boldsymbol{u}$ since the H$^{\mathrm{div}}$ projection along with the DG weak form of the momentum and continuity equations causes $\boldsymbol{u}$ to be fully described by the solenoidal subspace. In general, we have $\mathbb{D}^{-1}\mathbb{D} \neq \mathbb{I}$, so this identity property does not hold for a generic vector field.

The velocity vector field can now be described by nine degrees of freedom in each cell, coefficients $s_1 \cdots s_9$. Each of the nine coefficients is multiplied by the corresponding vector valued polynomials in equation (33), numbered from left to right. The sum of these products form the complete vector valued polynomial field.

We will further restrict the number of degrees of freedom by making use of the five invariants from equation (32). Let the coefficients for the constant and linear polynomials be determined by the five invariants. The four quadratic terms are now the only degrees of freedom. Picking four arbitrary numbers, σ_1 to σ_4, for the coefficients s_6 to s_9, one can find corresponding constant and linear weights that make the final vector field keep the selected invariants. This is done by forming a 9x9 matrix system,

$$
\begin{bmatrix}
q_{11} & q_{12} & q_{13} & q_{14} & q_{15} & q_{16} & q_{17} & q_{18} & q_{19} \\
q_{21} & q_{22} & q_{23} & q_{24} & q_{25} & q_{26} & q_{27} & q_{28} & q_{29} \\
q_{31} & q_{32} & q_{33} & q_{34} & q_{35} & q_{36} & q_{37} & q_{38} & q_{39} \\
q_{41} & q_{42} & q_{43} & q_{44} & q_{45} & q_{46} & q_{47} & q_{48} & q_{49} \\
q_{51} & q_{52} & q_{53} & q_{54} & q_{55} & q_{56} & q_{57} & q_{58} & q_{59} \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}
\cdot
\begin{bmatrix}
s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9
\end{bmatrix}
=
\begin{bmatrix}
U_1 \\ U_2 \\ R_1 \\ R_2 \\ R_3 \\ \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4
\end{bmatrix},
\tag{34}
$$

where the first five rows contain appropriate weights q_{ki} such that the correct integral from equation (32) is computed when the row is multiplied by the vector of coefficients, s_i. These weights can easily be computed by quadrature. By solving equation (34) the full set of nine solenoidal coefficients, s_i, can be recovered from the reduced set of four degrees of freedom, σ_j. The original 12 degrees of freedom of the Lagrangian velocity field $\boldsymbol{u}$ can be recovered by use of the inverse projection $\boldsymbol{u} = \mathbb{D}^{-1} \boldsymbol{u}_s$.

We should note that we have no proof that the matrix in equation (34) is invertible for all possible cells, but in our extensive testing we have never found a cell where this was not the case. For other possible matrices where σ_j do not control the quadratic terms in the solenoidal basis—but instead let's say that σ_j are identical to the first four s_i—then it is easy to find cells where the condition number of the resulting matrix is very high, so the choice of the limited degrees of freedom σ_j is important.

3.2.2 Optimisation and cost functions

We have used optimisation to determine the remaining four degrees of freedom by minimising a cost function in $\mathbb{R}^4$. This optimisation is performed for each cell to find the coefficients resulting in the lowest cost, which should correspond to the most stable solution with an appropriate choice of cost function. For each choice of cost function we get a new method with unique properties.

An ideal cost function will keep the limited solution very close to the original non-limited solution while still damping non-physical oscillations. To achieve this we have selected to first calculate a component-wise slope limited solution $\boldsymbol{u}_{\mathrm{HT}}$ by applying the hierarchical Taylor slope limiter described in section 3.1 to each of the velocity components. This makes the result frame dependent, which is not desirable, but sadly unavoidable when using a scalar field slope limiter for a vector field.

The $\boldsymbol{u}_{\mathrm{HT}}$ field is used as the target in the optimisation since it maintains the correct convergence order. This target field does not keep all the solenoidal invariants, so it is always slightly out of reach. The minimum and maximum allowable velocity component values are computed for a set of points along each cell boundary by considering the average values in the neighbouring cells just like in the hierarchical Taylor slope limiter. The range of allowable values is extended to include the already limited field $\boldsymbol{u}_{\mathrm{HT}}$ in order to avoid limiting at smooth maxima.

The cost function is implemented in two steps. First take the coefficients s_i corresponding to the non-limited velocity field and use these in equation (34) to calculate the five invariants from equation (32). Then, for a set of points $P_i \in \partial K$ along the perimeter of a cell K, compute for each velocity component u_j ($j \in [1, 2]$) the current value $u_{ij} = u_j(P_i)$, the allowable minimum $u_{ij,\mathrm{min}}$, the preferred target $u_{ij,\mathrm{HT}}$, and the allowable maximum $u_{ij,\mathrm{max}}$. The cost at point P_i for velocity component u_j is then the squared distance from the target normalised by the range of allowed values,

$$C_{ij} = \left(\frac{u_{ij} - u_{ij,\mathrm{HT}}}{u_{ij,\mathrm{max}} - u_{ij,\mathrm{min}}} \right)^2. \tag{35}$$

As the second step, if the current value u_{ij} is higher than $u_{ij,\mathrm{max}}$, add

$$C_{ij,\kappa} = \kappa_L \left(\frac{u_{ij} - u_{ij,\mathrm{max}}}{u_{ij,\mathrm{max}} - u_{ij,\mathrm{min}}} \right)^2 + \kappa_C, \tag{36}$$

and similar if $u_{ij} < u_{ij,\mathrm{min}}$. If the current value is inside the stable bounds, then $C_{ij,\kappa} = 0$. We have used $\kappa_L = 1.0$ and $\kappa_C = 1.0$ as additional penalties for going outside the stable zone. We have used the six Lagrangian nodes as the points P_i, i.e., the vertices and facet midpoints. The total cost for the cell K to be optimised is then calculated as

$$C_K = \sum_{i=1}^{6} \sum_{j=1}^{2} \left(C_{ij} + C_{ij,\kappa} \right). \tag{37}$$

3.2.3 Optimized velocity fields

The cost function in equation (37) is unfortunately not able to steer even a perfect optimisation routine towards a solution that is completely free from spurious local maxima in all cells. There are three times more contributions to the cost function than the degrees of freedom $\sigma_1 \cdots \sigma_4$ in each cell—and that is when the cost function is looking for overshoots only in the Lagrange nodes, there might still be local maxima between the Lagrange nodes.

Detailed studies of the cost functions for some selected problematic cells show that there are cells where there does not exist any point in the four dimensional coefficient space σ_j where the cost is below $\kappa_C = 1.0$. Using a brute force optimisation algorithm that explores the entire solution space would not help for those cells. In conclusion, it is likely not possible—by cell-wise slope limiting—to reconstruct a velocity field that is both solenoidal and free from local maxima based on the presented reduced basis and the chosen criteria to detect local overshoots.

3.3 A split solenoidal slope limiting algorithm

Motivated by the above findings we introduce separate limiters for the *convected* velocity $\boldsymbol{u}$ and the *convecting* velocity $\boldsymbol{w}$. Looking at the convected and convecting velocity fields as separate, but related, is inspired by the common practice of using an explicit convecting velocity to linearise the Navier-Stokes equations—which is also what we have done in equation (14). A similar splitting is done in cell-centered finite volume schemes where the convecting velocity field is interpolated to the facets while the unknown convected velocity field consists of cell averages. The convecting velocity can be partially slope limited in a way that does not compromise on the solenoidal properties, or left entirely unlimited. The convected velocity field must be slope limited to avoid instabilities. The overall algorithm is shown in figure 2.
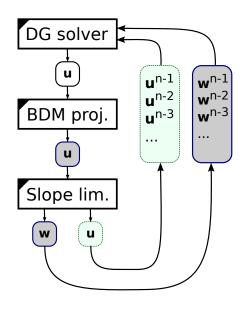

Figure 2: The slope limiting algorithm splits the velocity field into two separate time histories. Shaded velocities are solenoidal when the box boundary is a continuous line (grey background) and free from local maxima when the line is dotted (light green background).

In the following results section we have applied three different slope limiting procedures for the velocity field. The first is the naive unlimited method, the second applies the hierarchical Taylor polynomial based limiter to each velocity component of the convected velocity, and applies no limiting to the convecting velocity. The third method applies slightly different versions of the solenoidal slope limiter from section 3.2.2 to the two velocity fields. The convecting velocity field is limited without regards for the resulting cost function value in each cell; hence, for some cells the suppression of local maxima will not be successful. The convected velocity field is treated similarly, but here the cell velocity fields are replaced with the component-wise limited fields from the hierarchical Taylor polynomial based limiter, $\boldsymbol{u}_{\mathrm{HT}}$, in the cells where the local maxima suppression was unsuccessful.

4 Results

The methods described above have been implemented in Ocellaris (Landet, 2017), a two-phase solver framework which is built on top of FEniCS (Logg, Mardal, and Wells, 2012) and implemented in a mix of Python and C++. The optimiser employed is the BFGS

algorithm as implemented in SciPy (Jones et al., 2001), but for efficiency reasons we have written the optimisation algorithm along with the cost function in C++ to decrease the running time of the solver. The hierarchical Taylor based slope limiter for scalar fields has been implemented in C++ and does not contribute significantly to the running time of the solver. The rest of the Ocellaris solver is implemented in Python and depends on the C++ code generation facilities in FEniCS to utilize the computational resources optimally. The source code, input files, and scripts to reproduce the figures shown below can be found in Landet (2017).

In the following text we will compare three numerical algorithms, (i) a naive DG implementation without any slope limiters, (ii) a simple velocity slope limiting implementation using the scalar Taylor based limiter on each of the convected velocity components and no limiting of the convecting velocity (referred to as 'Hierarchical Taylor'), and (iii) the more involved cell-wise optimisation described above where both the convected and convecting velocity fields are optimised with slightly different criteria (referred to as 'Solenoidal').

To properly test the presented limiters it is important to ensure that artificially high viscosity is not a contributing factor to the method's stability. For the air/water free surface test cases we have applied $\nu = 1.0 \times 10^{-6}\,\mathrm{m^2\,s^{-1}}$ for both phases, which is artificially low for the air phase. With this choice of viscosity the non-linear convective instability impacts the results early in the simulations. This is also true for, e.g., $\nu = 1.0 \times 10^{-4}\,\mathrm{m^2\,s^{-1}}$, but not for very high kinematic viscosities, when ν approaches 1.0. Note that the dynamic viscosity μ—the parameter that goes into the weak form in equation (19)—is not the same in the two phases, it has the same factor 1000 jump as the density.

4.1   Taylor-Green vortex

The first test case is a study of the effect of slope limiting on the spatial convergence of the numerical scheme by considering a Taylor-Green vortex where the density is constant and the solution $\boldsymbol{u} = [u, v], p$ is given by,

$$
\begin{aligned}
u &= -\sin(\pi y)\cos(\pi x)\exp(-2\pi^2\nu t) \\
v &= \quad \sin(\pi x)\cos(\pi y)\exp(-2\pi^2\nu t) \\
p &= -\tfrac{1}{4}\,\rho(\cos 2\pi x + \cos 2\pi y)\exp(-4\pi^2\nu t)
\end{aligned}
\tag{38}
$$

for $t \in [0, 1]$ on a domain $\Omega = \{(x, y) \in [0, 2] \times [0, 2]\}$. This test case is often solved on a periodic domain, but we use Dirichlet boundary conditions for $\boldsymbol{u}$ to test the effect of the boundary. Initial conditions are given both at $t = 0$ and $t = -\Delta t$ to be able to use second order time stepping from the start. A constant time step of $\Delta t = 0.01$ is used and the kinematic viscosity is $\nu = 0.005$.

A slope limited solution should be identical to the non-limited solution for this test case—there are no non-smooth maxima in the true solution. Still, exact equivalence is not obtained since the DG FEM vector field is only a numerical approximation. The main discrepancy is in the handling of boundary conditions. Vertices on the boundary will be missing neighbour cells and may hence be unnecessarily slope limited. Figure 3 shows that the results converge towards the non-limited solution when slope limiting is avoided in the boundary facing cells. When boundary cells are included in the limiter the expected third order L^2 convergence rate is still obtained, but the constant is larger.
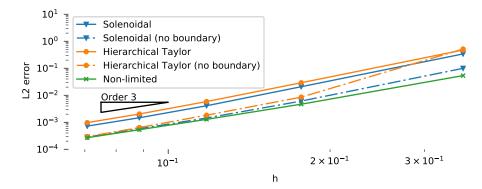
16

Figure 3: Spatial convergence rate on the Taylor-Green vortex test case. The rate of convergence is as expected, but missing neighbour cell info near boundaries in the limiters gives a larger constant in the L^2 error plot.

4.2 Dam break

The first two-phase flow example is a classic a dam break in a box simulation as illustrated in figure 4 based on the experiments by Martin and Moyce (1952). This is one of the most commonly used test cases in the high Reynolds number, low surface tension regime. This is the regime most interesting for studying marine and offshore structures, which is the ultimate goal of the method.
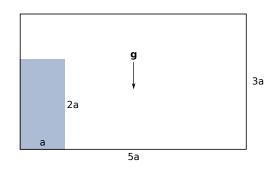

Figure 4: Dam break simulation geometry

The 2D rectangular water column is $1a$ wide and $2a$ high and starts at rest in the lower left corner of a box filled with air which is $5a$ wide and $3a$ high. The size of the column is the same as in the experiments, $a = 2.25\,\text{in} = 0.057\,15\,\text{m}$. The two phases are water and air with densities $1000\,\text{kg}\,\text{m}^{-3}$ and $1.0\,\text{kg}\,\text{m}^{-3}$ respectively. This corresponds to tables 2 and 6 in Martin and Moyce (1952). The acceleration of gravity is $g = 9.81\,\text{m}\,\text{s}^{-2}$ in the negative y-direction and the kinematic viscosity is $\nu = 1.0 \times 10^{-6}\,\text{m}^2\,\text{s}^{-1}$ for both phases.

The total kinematic energy, $E_k = \int_\Omega 1/2\,\rho\,\boldsymbol{u} \cdot \boldsymbol{u}\,\mathrm{d}\boldsymbol{x}$, is shown in figure 5 as a function of time for the slope limited and the non-limited methods. From the start, with both fluids at rest, the kinetic energy increases as the water mass starts to flow down and towards the right wall. There is a slight reduction as the water hits this wall at approximately $t = 0.19\,\text{s}$. The Gibbs oscillations start to dominate the non-limited solution after a short time and from $t = 0.04\,\text{s}$ the solution is non-physical and completely dominated by

17

numerical errors. In the rest of this paper we will remove the non-limited method from the results as it leads to non-physical solutions in all cases.
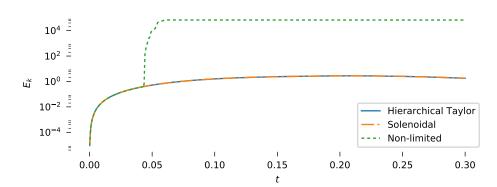

Figure 5: Kinetic energy as a function of time for three slope limiting strategies

The results of the two slope limited methods are compared to the experimental results in figure 6. The experimental data points are the ensemble averages of the results reported for the same geometry by Martin and Moyce (1952). We have employed free slip boundary conditions, and hence the surge front moves slightly faster than in the experimental results, but the qualitative behaviour is correct. The maximum height of the water column matches very well until the water hits the domain boundary and creates a jet shooting up along the right wall at $t \approx 0.19\,\mathrm{s}$ which corresponds to $T \approx 3.5$ and $\tau \approx 2.8$. The right wall was placed much further away in the experimental setup.

As can be seen in figures 5 and 6, the two investigated limiters perform very similarly on the dam break test case. The 'Solenoidal' method requires significantly more computation per time step than the component-wise Taylor based limiter and is more complex to implement. One reason why one could still consider using this method is that the optimisation may keep the difference between the convected and the convecting velocity fields smaller and hence closer to the true solution. To test how much this optimisation contributes we have calculated the difference between the velocity fields, $\Delta U = \|\boldsymbol{u} - \boldsymbol{w}\| / \|\boldsymbol{u}\|$, for all time steps and the results can be seen in figure 7.

The difference is large in the beginning since both $\boldsymbol{u}$ and $\boldsymbol{w}$ are close to zero at this time. After the initial phase the difference is low for both limiters. We can see that right before the water impacts the wall at approximately $t = 0.19\,\mathrm{s}$ the slope limiters start to perform differently for some time until the impact is over. The reason is that the 'Hierarchical Taylor' slope limiter introduces a very high divergence in the convected velocity field in the tank corner. The optimisation based limiter performs as expected and the difference between the two fields is lower, but not entirely removed.

Later in the time series there is a new event when the water is shooting up in a thin jet along the right wall at approximately $t = 0.27\,\mathrm{s}$. Here we can see that the optimised slope limiter is introducing slightly higher ΔU differences than the Taylor based slope limiter, quite contrary to the intention. We have no direct explanation of this, except that it seems related to the locally high shear in the vertical velocity field.
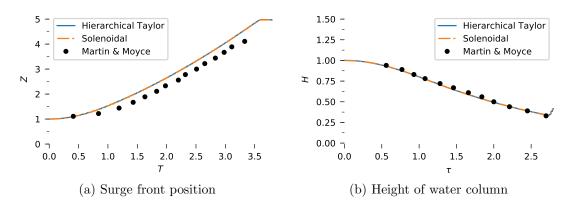
(a) Surge front position
(b) Height of water column

Figure 6: Comparison of the numerical results with experimental results by Martin and Moyce (1952). The vertical axis scaling is $Z = z/a$ and $H = \eta/2a$. Measured from the lower left corner, z is the surge front position and η is the height of the water column. The horizontal axis scaling is $T = t\sqrt{2g/a}$ and $\tau = t\sqrt{g/a}$.
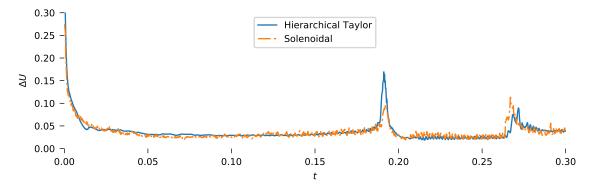

Figure 7: Dam break; comparison of the two velocity slope limiters in terms of, ΔU, the scaled L^2 distance between convected and convecting velocity fields.
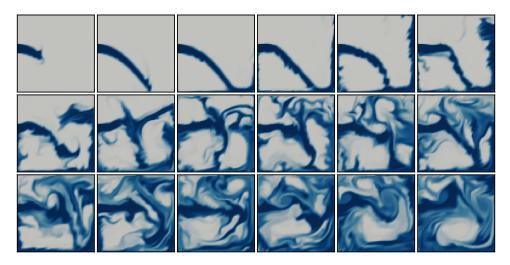
Figure 8: Tank filling results. The first row shows the density field at $t = 0.167$, 0.333, 0.500, 0.667, 0.833, and 1.000, the second and third row have the same temporal spacing, hence the lower rightmost image shows the simulation end point, $t = 3$ s. The finest resolution is shown, using the 'Hierarchical Taylor' slope limiter.

4.3   Tank filling

To provide a more challenging test of the method's stability with a large and complex free surface we have performed a set of tank filling simulations inspired by Guermond, de Luna, and Thompson (2017), but unlike in their work we have used physical properties closer to normal water and air, and the results are hence more energetic. The physical properties are the same as in the dam break test case.

The geometry of the tank can be seen along with the resulting density field in figure 8. The domain is the unit square and the inlet is located between 0.5 m and 0.625 m above the floor on the left wall, $1/8$ m wide. The outlet is located in the roof between 0.5 m and 0.625 m from the left wall and is also $1/8$ m wide. The inlet and outlet velocities are prescribed as $2\,\mathrm{m\,s^{-1}}$, constant for all time. At t=0 the tank is completely filled with air, the velocity inside is zero, and the pressure is hydrostatic.

The results are computed on a regular mesh with triangular cells and no mesh grading. There are eight cells across the inlet and 64x64x2 cells in total. The domain is first divided into squares and then these are subdivided into triangles. The time step is $\Delta t = 0.0001$ s. The solution has also been calculated on a finer mesh with 16 elements across the inlet, 128x128x2 cells in total. The time step was then $\Delta t = 5.0 \times 10^{-5}$ s and only the 'Hierarchical Taylor' slope limiter was tested. The results from all three simulations qualitatively exhibit the same physics and look similar. The density field from the simulations on the fine mesh is shown in figure 8 at regular intervals up to the maximum simulation time of 3 s.

The performance of the slope limiters is compared in terms of stability and conservation. We have calculated the total mass and energy inside the tank and the results can be seen in figures 9a and 9b. The analytical solutions in the figures are computed under the assumption that only air exits through the outlet, which is not true in the second half of the simulation. When looking at conservation of mass in the first part of the simulations, only very minor differences can be seen between the methods; both methods preserve
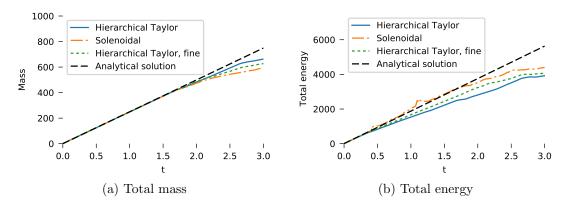
Figure 9: Tank filling: comparison of the two velocity slope limiters in terms of conservation of mass and energy
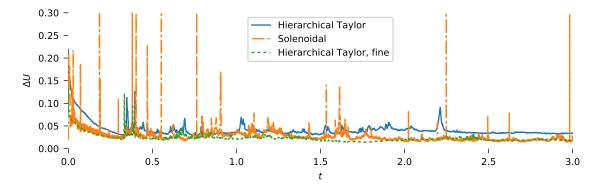

Figure 10: Tank filling: comparison of the two velocity slope limiters in terms of, ΔU, the scaled L^2 distance between convected and convecting velocity fields.

mass very well. This is exactly as expected since mass conservation is the key invariant in both slope limiters—the convecting velocity field is always solenoidal.

The loss of mass and energy through the outlet—starting at approximately 1.7 seconds—can also be seen in the time history of the total energy in figure 9b. Up to that point the optimisation based solenoidal slope limiter can be seen to preserve total energy better than the component-wise hierarchical Taylor based limiter. Refining the mesh improves the results from the hierarchical Taylor based limiter, but the optimised solenoidal limiter still outperforms it, even if it is running on a much coarser mesh.

Both methods are stable for the duration of the simulation, but some spurious increases in total energy can be seen in the 'Solenoidal' method. The 'Hierarchical Taylor' method has no noticeable spurious increases, at the expense of slowly loosing energy. The reason for the spurious increase in energy in the 'Solenoidal' method around 1.2 s in figure 9b may be the fact that local maxima are allowed to occur in between the Lagrange polynomial nodes.

The difference between the convected and the convecting velocity, ΔU, has been calculated in the same manner as in the dam break test case and the results can be seen in figure 10. Both methods preserve a low difference between the velocity fields for most of the time steps, but both have instances when the difference grows larger. This is particularly noticeable in the 'Solenoidal' method which has some large spikes. In general

21

the optimised method maintains a lower difference, on par with the refined hierarchical Taylor based method. Both methods are shown to recover from temporary spikes and return to a low difference, which is the correct solution.

5   Discussion

We have shown how convective instabilities in higher order two-phase flow simulations can be stabilised without sacrificing mass conservation. Our efforts have been focused on the use of slope limiters to achieve this stabilisation. Using optimisation on a solenoidal reduced basis to create one single velocity field that is both stable and solenoidal was investigated. The results show strong indications that this may in fact not be possible. To overcome this problem, two methods that treats the convected velocity field differently from the convecting velocity field were proposed. These methods combine exact mass and momentum conservation with convective stability for two-phase simulations containing large density jumps inside the computational domain.

The suggested methods have been tested to see if there were any negative impact on the solution of smooth problems; and optimal convergence on a Taylor-Green vortex problem was show. Further, a two-dimensional dam break with a factor 1000 sharp density jump in the computational domain was tested. Such problems are not possible to handle without some form of higher order convective stabilisation. Both slope limiting strategies handled this problem well. Finally, we have studied mass and energy conservation in a very energetic flow—a tank filling simulation. A summary of our results is that slope limiting is a possible way to stabilise a mass conserving discontinuous Galerkin method for the incompressible Navier-Stokes equation with large density jumps.

There are other methods than slope limiting for stabilising high order convective instabilities in discontinuous Galerkin methods. Instead of using an explicit post processing operator one can add specifically tailored implicit damping to the cells near the free surface with a method such as the entropy viscosity method, (Zingan et al., 2013). Another method that has been used is to selectively reduce the polynomial approximation order in cells near the free surface, a special use of p-adaptivity (Robert J. Labeur, private communication, 2017). The advantages of the explicit slope limiting method is the decoupling of the convective stabilisation from the definition of the weak form and the assembly of the equation system. The relative ease of implementation is also a large advantage of the component-by-component velocity limiting method. Some type of scalar slope limiter is likely to be already present in a DG FEM code. Both the presented methods require no extra method dependent viscosity parameters to be derived, and no selective p-adaptivity is needed in the software. There is no need to tune parameters in the methods to achieve stability. Extension to higher order is also straight forward.

Acknowledgements

References

Arnold, Douglas Norman. 1982. "An interior penalty finite element method with discontinuous elements." *SIAM journal on numerical analysis* 19 (4): 742–760.

Baker, Garth A., Wadi N. Jureidini, and Ohannes A. Karakashian. 1990. "Piecewise Solenoidal Vector Fields and the Stokes Problem." *SIAM Journal on Numerical Analysis* 27 (6): 1466–1485.

Brezzi, Franco, and Michel Fortin. 1991. *Mixed and Hybrid Finite Element Methods.* no. 15 In Springer Series in Computational Mathematics. Springer-Verlag, New York.

Cockburn, Bernardo, Guido Kanschat, and Dominik Schötzau. 2004. "The local discontinuous Galerkin method for the Oseen equations." *Mathematics of Computation* 73 (246): 569–594.

Cockburn, Bernardo, Guido Kanschat, and Dominik Schötzau. 2005. "A locally conservative LDG method for the incompressible Navier-Stokes equations." *Mathematics of Computation* 74 (251): 1067–1096.

Cockburn, Bernardo, and Chi-Wang Shu. 1998. "The Runge–Kutta Discontinuous Galerkin Method for Conservation Laws V." *Journal of Computational Physics* 141 (2): 199–224.

Cockburn, Bernardo, and Chi-Wang Shu. 2001. "Runge-Kutta discontinuous Galerkin methods for convection-dominated problems." *Journal of scientific computing* 16 (3): 173–261.

Epshteyn, Yekaterina, and Béatrice Rivière. 2007. "Estimation of penalty parameters for symmetric interior penalty Galerkin methods." *Journal of Computational and Applied Mathematics* 206 (2): 843–872.

Guermond, Jean-Luc, Manuel Quezada de Luna, and Travis Thompson. 2017. "An conservative anti-diffusion technique for the level set method." *Journal of Computational and Applied Mathematics* 321: 448–468.

Guermond, Jean-Luc, and Luigi Quartapelle. 2000. "A Projection FEM for Variable Density Incompressible Flows." *Journal of Computational Physics* 165 (1): 167–188.

Guermond, Jean-Luc, and Abner Salgado. 2009. "A splitting method for incompressible flows with variable density based on a pressure Poisson equation." *Journal of Computational Physics* 228 (8): 2834–2846.

Hirt, Cyril W., and Billy D. Nichols. 1981. "Volume of fluid (VOF) method for the dynamics of free boundaries." *Journal of computational physics* 39 (1): 201–225.

Hundsdorfer, Willem, Steven J. Ruuth, and Raymond J. Spiteri. 2003. "Monotonicity-Preserving Linear Multistep Methods." *SIAM Journal on Numerical Analysis* 41 (2): 605–623.

Jones, Eric, Travis Oliphant, Pearu Peterson, and others. 2001. "SciPy: Open source scientific tools for Python." `http://www.scipy.org/`.

Kirby, Robert C., Anders Logg, Marie E. Rognes, and Andy R. Terrel. 2012. "Common and unusual finite elements." In *Automated Solution of Differential Equations by the Finite Element Method,* Lecture Notes in Computational Science and Engineering. 95–119. Springer, Berlin, Heidelberg.

Kuzmin, Dmitri. 2010. "A vertex-based hierarchical slope limiter for -adaptive discontinuous Galerkin methods." *Journal of Computational and Applied Mathematics* 233 (12): 3077–3085.

Kuzmin, Dmitri. 2013. "Slope limiting for discontinuous Galerkin approximations with a possibly non-orthogonal Taylor basis." *International Journal for Numerical Methods in Fluids* 71 (9): 1178–1190.

Landet, Tormod. 2017. "Ocellaris DG-FEM software and input files to reproduce results." Zenodo. http://doi.org/10.5281/zenodo.845352.

Leonard, Brian Phillip. 1988. "Simple high-accuracy resolution program for convective modelling of discontinuities." *International journal for numerical methods in fluids* 8 (10): 1291–1318.

Liu, Chun, and Noel J. Walkington. 2007. "Convergence of Numerical Approximations of the Incompressible Navier–Stokes Equations with Variable Density and Viscosity." *SIAM Journal on Numerical Analysis* 45 (3): 1287–1304.

Logg, Anders, Kent-Andre Mardal, and Garth Wells. 2012. *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book.* Springer Science & Business Media.

Martin, John Christopher, and William James Moyce. 1952. "Part IV. An Experimental Study of the Collapse of Liquid Columns on a Rigid Horizontal Plane." *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 244 (882): 312–324.

Michoski, C., C. Dawson, E. J. Kubatko, D. Wirasaet, S. Brus, and J. J. Westerink. 2016. "A Comparison of Artificial Viscosity, Limiters, and Filters, for High Order Discontinuous Galerkin Solutions in Nonlinear Settings." *Journal of Scientific Computing* 66 (1): 406–434. `https://link.springer.com/article/10.1007/s10915-015-0027-2`.

Muzaferija, Samir, Milovan Peric, Pierre C. Sames, and Thomas E. Schellin. 1998. "A Two-Fluid Navier-Stokes Solver to Simulate Water Entry." In *Proceedings from the 22nd symposium on naval hydrodynamics,* 277–289. Washington, DC.

Nitsche, Joachim A. 1971. "Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind." *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 36 (1): 9–15.

Nédélec, Jean-Claude. 1986. "A new family of mixed finite elements in $\mathbb{R}^3$." *Numerische Mathematik* 50 (1): 57–81.

Olsson, Elin, and Gunilla Kreiss. 2005. "A conservative level set method for two phase flow." *Journal of Computational Physics* 210 (1): 225–246.

Osher, Stanley, and James A. Sethian. 1988. "Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations." *Journal of computational physics* 79 (1): 12–49.

Peskin, Charles S. 2002. "The immersed boundary method." *Acta Numerica* 11: 479–517.

Pyo, Jae-Hong, and Jie Shen. 2007. "Gauge–Uzawa methods for incompressible flows with variable density." *Journal of Computational Physics* 221 (1): 181–197.

Shahbazi, Khosro, Paul F. Fischer, and Christopher Ross Ethier. 2007. "A high-order discontinuous Galerkin method for the unsteady incompressible Navier-Stokes equations." *Journal of computational physics* 222 (1): 391–407.

Sussman, Mark, Peter Smereka, and Stanley Osher. 1994. "A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow." *Journal of Computational Physics* 114 (1): 146–159.

Sweby, Peter K. 1984. "High resolution schemes using flux limiters for hyperbolic conservation laws." *SIAM Journal on Numerical Analysis* 21 (5): 995–1011.

Touré, Mamadou Kabirou, Adil Fahsi, and Azzeddine Soulaïmani. 2016. "Stabilised finite-element methods for solving the level set equation with mass conservation." *International Journal of Computational Fluid Dynamics* 30 (1): 38–55.

Ubbink, Onno. 1997. "Numerical prediction of two fluid systems with sharp interfaces." Ph.D. thesis. Imperial College, University of London.

Unverdi, Salih Ozen, and Grétar Tryggvason. 1992. "A front-tracking method for viscous, incompressible, multi-fluid flows." *Journal of Computational Physics* 100 (1): 25–37.

Zingan, Valentin, Jean-Luc Guermond, Jim Morel, and Bojan Popov. 2013. "Implementation of the entropy viscosity method with the discontinuous Galerkin method." *Computer Methods in Applied Mechanics and Engineering* 253: 479–490.