On the Numerical Convergence Properties of a Flow Solver Based on the Hybridizable Discontinuous Galerkin Method

Pedro Manuel Brito Monteiro
pedro.m.b.monteiro@tecnico.ulisboa.pt

Instituto Superior Técnico, Universidade de Lisboa, Portugal

October 2018

Abstract

In this project the Hybridizable Discontinuous Galerkin (HDG) method has its performance and robustness tested. For the performance analysis, a vortex-shedding flow around a cylinder is simulated and the convergence of the HDG method is compared to the convergence of the Finite Volume Method (FVM). The unsteady incompressible Navier–Stokes equations were used to simulate the flow. For a comparison with the FVM, a mesh generated by blocks was used, with second-order quadrilateral elements. The convergence with the element size and the residual of the nonlinear solver is analysed to conclude about the performance of the HDG method. The influence of the nonlinear solver used, Picard iteration or Newton's method, is also studied. For the robustness analysis, the method is tested at high Reynolds numbers for the flows over a flat plate and around a NACA0012 airfoil. In this implementation no stabilization is added to the method. Firstly, the limits of the HDG method were found using the laminar flow over the flat plate and the airfoil. Once the limits for the method were found, a turbulence model (Spalart–Allmaras) was implemented to check its influence on the method. The effect of the linearization of the source terms is analysed as well as the influence of a user defined parameter intrinsic to the HDG method.

Keywords: Hybridizable Discontinuous Galerkin, high-order, convergence, robustness, vortex-shedding, Spalart–Allmaras.

1. Introduction

The study of the flow of fluids is indispensable for many engineering applications, such as vehicle design, aero-acoustics and turbomachinery. The simulation of fluid flows is studied in Computational Fluid Dynamics (CFD), and it is the focus of this work.

The behaviour of fluids is described by the Navier–Stokes and the energy equations. These equations are very complex due to the non-linearity and the coupling between them. Due to these problems, obtaining an exact solution directly from the system of partial differential equations is almost impossible. Usually, the solution of this system is obtained approximately through methods that divide the domain into many subdomains where the solution is approximated by simple polynomial functions.

Currently, the most used method to solve the Navier–Stokes equations is the Finite Volume Method (FVM), which ensures the conservation of the mass flux, essential in a fluid analysis. While this method is robust and flexible, additional stabilization terms are usually required to solve the Navier–Stokes equations even at low Reynolds numbers. Additionally, it is very complex to implement high-order approximations to the method.

This work is focused on an alternative. A method applicable to high-order capable of solving the Navier–Stokes equations coupled. This method is the Hybridizable Discontinuous Galerkin method (HDG). Theoretically, this method is high-order stable, has many less degrees of freedom than the Discontinuous Galerkin method and conserves the mass flow.

For the HDG method to be widely applied in engineering, its convergence and robustness must be tested. The motivation for this work is to address the numerical convergence properties of the HDG. It's intended to compare this HDG method to the classic Finite Volume Method, testing mainly its robustness and convergence. The HDG code used for this thesis was written by Paipuri (2018).

This work starts with an investigation on the methods used to date.

A second section will mention the discretization

used for the HDG method.

The third section will contain the convergence analysis of the method. In here, the unsteady case of a laminar vortex-shedding flow around a cylinder is studied. The results are compared to those obtained using the Finite Volume Method.

The fourth section is where the robustness of the method is tested. Using the flow over a flat plate and around the NACA0012, the limits of the method using the incompressible Navier–Stokes equations were found. Afterwards, the robusntess of the method using the Spalart–Allmaras turbulence model (Spalart and Allmaras (1994)) is analysed. No stabilization terms were added to the equations.

Finally, some concluding remarks on the behaviour of the method are presented, as well as future developments.

2. Background

With the development of numerical methods, the solution of many complex problems was at engineers grasp. Most methods divide the domain in non-overlapping elements, where the solution is approximated. The differences appear in the way the PDE is satisfied, either satisfying the equations conservatively or point-wise. From Ferziger and Peric (2002) it is possible to understand the main differences between the methods. The Finite Volume Method has a big advantage in its application to fluid flows, which is the conservation of mass. This method uses the average of the solution over the cell and a flux between the elements to satisfy the PDE conservatively. Although the method can be applied to complex geometries, extending the method to high-order is not a simple task. For fluid mechanics, the Finite Element Method is not commonly used, since the mass conservation is not intrinsic to the method. However, the method is very flexible and performs well with high-order approximations. Consequently, there has been research in obtaining the best of each method, so that a high-order conservative method can be implemented.

2.1. Finite Volume Method

The Finite Volume Method is the most used method for CFD. This method uses the cell average solution as a local approximation satisfying conservatively the PDE. The conservation of the mass flux is essential to obtain a good approximated solution. The method can handle unstructured meshes and hanging nodes, making it very flexible. There may be some problems in the stability of this method, mainly solving the incompressible Navier–Stokes equations (Ferziger and Peric (2002)).

The coupling between the momentum equations and the continuity equation can generate problems. For this reason, many codes use a segregated approach. Additionally, the method usually requires techniques to avoid pressure oscillations in the solution. Using unstructured grids, it is very complex to obtain a high-order approximation, since the accuracy is worse due to the non-orthogonality between the element boundary and the gradient approximation. Deferred correction schemes can be implemented to attempt to reduce the error due to this non-orthogonality.

While the method is very flexible and robust when using the correct stabilization methods, its accuracy and efficiency are not ideal.

2.2. Discontinuous Galerkin method

The classic Finite Element Method (FEM) uses the Continuous Galerkin (CG) approach. This method uses polynomial functions to approximate the solution in each cell. These polynomials are continuous between each cell, hence the name Continuous Galerkin. Due to its use of polynomials to approximate the solution locally, it is much simpler to apply this method to high-order than the FVM. Combining the conservation of flux of the FVM and the high-order polynomials approximation of the FEM is the intent of the Discontinuous Galerkin method (DG).

The Discontinuous Galerkin method uses completely discontinuous basis functions and the information is passed through a numerical flux on the element boundary. Although this change makes the method conservative, with divergence-free solutions, it doubles the number of degrees of freedom (DoF) at the element interface. However, the data structure is extremely local since the elements only communicate with immediate neighbours regardless of the order of the approximation.

2.3. Hybridizable Discontinuous Galerkin Method

To attempt to reduce the number of degrees of freedom of the Discontinuous Galerkin method, the trace is idealized. The trace is a new variable that is unique at the element's boundaries interface, therefore it is shared by both neighbours. This unique trace at inter-element boundaries is obtained by enforcing the continuity of flux. Using static condensation, it is possible to reduce the number of degrees of freedom of the system, separating the problem into global and local problems. Then, on each element the equations are solved with a Dirichlet condition applied on its boundary nodes.

The Hybridizable Discontinuous Galerkin method was introduced by Cockburn et al. (2009b) for a second-order elliptic problem, although many hybridization techniques had already been studied by the author. This HDG method has an optimal convergence rate of $k + 1$ for the solution and the derivatives, whereas with a post-processing a super-convergence of $k + 2$ in the solution is achieved, using a degree k to approximate the

unknowns.

The application of HDG to incompressible unsteady Navier—Stokes flow was first reported in Nguyen et al. (2011). Here, once again the properties of optimal convergence, super-convergence with local post-processing, unified treatment of boundary conditions and the numerical fluxes were shown for the Navier–Stokes discretization. In Paipuri (2018) a comparison between the high-order Continuous Galerkin and HDG was performed and in Cockburn et al. (2009a) it was concluded that the HDG method of degree k is as efficient as a CG method of degree $k+1$ for convection diffusion-reaction problems.

Although HDG methods are locally conservative and stable for convection-dominated flows, they can produce oscillatory solutions using turbulence models. Techniques like those applied in FVM have been applied to solve these problems.

The main advantages of the HDG method when compared to the usual FVM codes are: the coupling of the Navier–Stokes equations; high-order approximations; high-order terms integrated into the matrices; no deferred corrections in unstructured meshes.

2.4. The Spalart–Allmaras Turbulence Model

Most of the flows studied by engineers are turbulent. Most engineering simulations are performed using the Reynolds Averaged Navier–Stokes (RANS) equations with a turbulence model. Turbulence models intend to determine the Reynolds stresses necessary to close the RANS equations. The Spalart–Allmaras (SA) model is a one equation model introduced in Spalart and Allmaras (1994). To reduce the numerical problems with negative eddy viscosity, the model was changed in Allmaras et al. (2012) to accommodate a negative value of the main variable. This model would later be known as the Negative Spalart–Allmaras (SA-neg).

Implementing turbulence models in high-order methods is a challenging task. The coupling between the RANS equations and the turbulence model can lead to the divergence of the solution (Moro et al. (2011)). Usually, it is necessary to supplement the turbulence model with some sort of stabilization to converge the simulation.

In Moro et al. (2011) different approaches are studied to avoid oscillations and negative values of the eddy viscosity, using artificial dissipation and artificial viscosity. In practice, the SA model applied to the HDG method diverges if no stabilization is implemented. Here, the SA-neg introduced in Allmaras et al. (2012) is used to deal with the negative viscosity. However, no stabilization was introduced to the equation.

3. Mathematical Foundation

The applied Hybridizable Discontinuous Galerkin (HDG) method solves the unsteady incompressible Navier–Stokes equations, where the momentum equations are coupled with the continuity equation. The domain Ω has the boundary $\partial\Omega$, which is divided into Dirichlet and Neumann boundaries, $\partial\Omega_D$ and $\partial\Omega_N$. The system is shown in (1), valid in the domain Ω.

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla\cdot(\mathbf{u}\otimes\mathbf{u}) - \nabla\cdot(-p\mathbf{I} + \nu\nabla\mathbf{u}) = \bar{\mathbf{s}}\,, \tag{1a}$$

$$\nabla\cdot\mathbf{u} = 0 \tag{1b}$$

On the Dirichlet boundary condition, $\partial\Omega_D$, equation $\mathbf{u} = \bar{\mathbf{u}}$ is applied and on the Neumann boundary condition, $\partial\Omega_N$, equation $(-p\mathbf{I} + \nu\nabla\mathbf{u})\cdot\mathbf{n} = \bar{\mathbf{t}}$ is enforced. The condition for $t = 0$ is $\mathbf{u} = \mathbf{u_0}$.

Where in the equations above, $\mathbf{u}$ denotes the velocity vector, p the pressure, t the time, $\bar{\mathbf{u}}$ the imposed velocity on the Dirichlet boundary and $\bar{\mathbf{t}}$ the prescribed flux on the Neumann boundary. The initial velocity field is denoted by $\mathbf{u_0}$. ν is the kinematic viscosity. The body force $\bar{\mathbf{s}}$ is zero in all the domain. The matrix $\mathbf{I}$ represents the identity matrix.

In order to perform the discretization of the system the velocity gradient will originate a new variable, $\mathbf{L}$. The domain Ω is split into a series of elements Ω^e, in which the equations defined in (1) are valid. The elements are non-overlapping, and the union of the edges of the elements is denoted by Γ, $\Gamma = \bigcup_{e=1}^{n_{el}} \partial\Omega^e$, where n_{el} is the number of elements the domain is split into. Due to the nature of the method, the discontinuities between the elements must be taken into consideration, originating two more equations, equation (2b) and (2c).

$$\mathbf{L} - \nabla\mathbf{u} = 0 \text{ in } \Omega^e\,, \tag{2a}$$

$$[\![\mathbf{u}\otimes\mathbf{n}]\!] = 0 \text{ on } \Gamma\setminus\partial\Omega\,, \tag{2b}$$

$$[\![(-p\mathbf{I} + \nu\nabla\mathbf{u})\cdot\mathbf{n}]\!] = 0 \text{ on } \Gamma\setminus\partial\Omega \tag{2c}$$

Where the operator $[\![\cdot]\!]$ denotes the jump. The jump operator imposes the continuity of a variable over the boundary.

The advantage of a HDG is that the flux between the elements is approximated by a new variable, the trace variable $\hat{\mathbf{u}}$. This extra variable allows for the definition of two problems, a local and a global problem.

In the local problem the Navier–Stokes equations are solved for each element with the velocity imposed on its boundary ($\mathbf{u} = \hat{\mathbf{u}}$ on $\partial\Omega^e$), and the average pressure along the element boundary, $\hat{p}_e$. With these two variables, the local problem is completely defined, and it is possible to compute the velocity field and pressure distribution in the element.

The global problem only solves for the trace variable $\hat{\mathbf{u}}$ and the average pressure on the boundary $\hat{p}_e$.

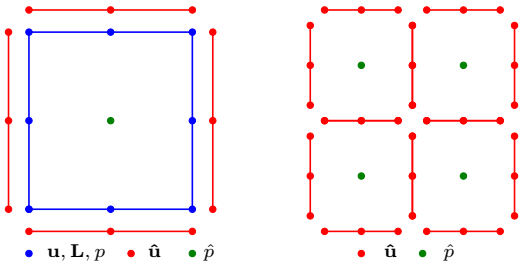

(a) Local Problem (b) Global Problem

Figure 1: Variables for Local and Global problems

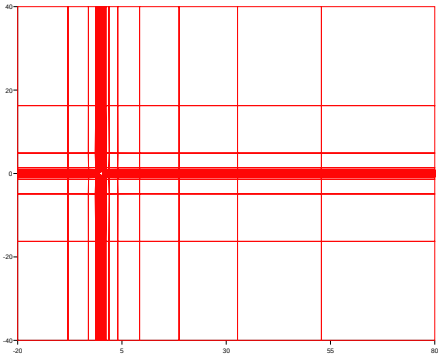

Figure 2: Mesh by blocks used for the analysis

The time derivative is approximated using second-order a Backward Differentiation Formulae (BDF) scheme.

The condition in (2b) was applied weakly through the trace $\hat{\mathbf{u}}$ and using the term τ_u. Since the trace $\hat{\mathbf{u}}$ is unitary in the border between two elements and not discontinuous, forcing the jump equal zero is the same as forcing the velocity $\mathbf{u}$ equal to the trace. In this manner, the Dirichlet boundary condition for the local problem is also weakly enforced. The parameter τ_u is an intrinsic part of the method, since the method does not work with this parameter set as zero. However, the robustness and accuracy of the method can be influenced by this parameter. In this work, the parameter τ_u is set as 1 for most simulations, except some simulations where the value is higher due to the high Reynolds number flow, in which case the value is later defined.

4. Laminar Flow Around a Cylinder

In this chapter the performance of the Hybridizable Discontinuous Galerkin method and the classic Finite Volume Methods is compared. For this comparison, the unsteady laminar vortex-shedding flow around a cylinder will the simulated.

The simulation was computed with a Reynolds of 100, which ensures the flow is laminar. The cylinder has a diameter of 1 m and a no slip condition is applied in its wall (velocity is zero). The inlet has an imposed horizontal velocity of 1 m/s and the outlet pressure is imposed as 0. In order to get the Reynolds number of 100, the kinematic viscosity is $1 \times 10^{-2}\,\mathrm{m}^2/\mathrm{s}$.

All the meshes used in this analysis have second-order quadrilateral elements, with nine nodes each. Therefore, the shape functions for the velocity and the pressure are second-order. Due to the post-processing computations, the velocity and the pressure converge with fourth and third-order respectively.

The mesh used for the following simulations is composed of 9 blocks.

4.1. Convergence with Mesh size

For this analysis, 4 second-order meshes were used in the HDG analysis.

Table 1: Square Meshes

Mesh	Ratio h_i/h_1	Elements	Time Step [s]
1	2	34 176	0.05
2	3.2	13 350	0.08
3	4	9 655	0.10
4	8	2 136	0.20

The ratio h_i/h_1 represents the ratio of a mesh element size to the most refined mesh element size. However, a more refined mesh not present in table 4.1 was used for the FVM. Consequently, the ratio was computed with the most refined mesh overall (for both FVM and HDG). Since the elements used for the FVM are first-order and those used for the HDG are second-order, the ratio was computed using the number of nodes. Two meshes have the same ratio h_i/h_1 if they have the same number of nodes. All meshes have a maximum Courant number of 15.

Following the study performed by Paipuri (2018), the HDG method's element has 13 degrees of freedom (in the global problem) if second-order quadrilateral elements are used. The finite volume only has 3 degrees of freedom per element, however, the elements are comprised of only 4 nodes instead of the 9 nodes that comprise the second-order quadrilateral elements. Since for the meshes used one second-order quadrilateral element is generated from four first-order elements, the global problem of the HDG method only has 1.08 times more degrees of freedom than the FVM.

The HDG data will be compared to data obtained using the FVM. The Finite Volume Method data was computed using ReFRESCO (MARIN (2015)), a viscous-flow CFD code that solves multiphase unsteady incompressible flows using the RANS equations.

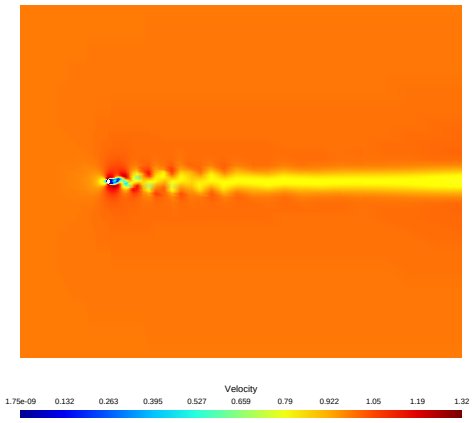

Figure 3: Velocity magnitude for square mesh with 53 400 elements

Figures 4 and 5 show the convergence of the Strouhal number and the drag coefficient for the HDG and the FVM in red and blue respectively. Due to the higher-order convergence of the HDG method, meshes with less elements were used.

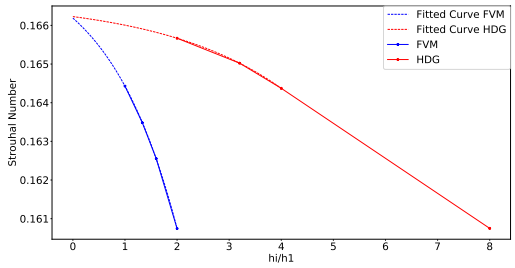

Figure 4: FVM and HDG convergence analysis of the Strouhal number.

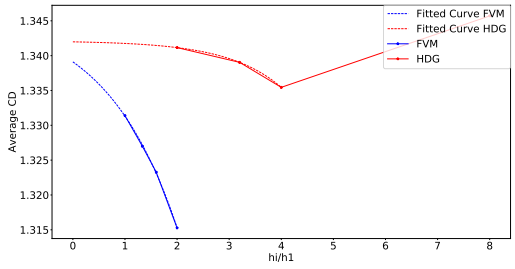

Figure 5: FVM and HDG convergence analysis of the drag coefficient.

Both methods converge to the same drag coefficient. It is clear that the HDG method manages to converge with a much coarser mesh. This advantage makes it worth the extra degrees of freedom of the HDG method. The mesh with ratio 4 is much more accurate than any of the Finite Volume Method simulations, which used meshes with up to eight times as many nodes.

In the computation of the Strouhal number the HDG method is still more efficient. Both methods converge to a Strouhal number of 0.166.

It is clear that the HDG method performs better than the classic FVM. The higher-order of the HDG method allows a better approximation of the solution with many times less nodes than the FVM. Additionally, the convergence with the element size is also quicker.

4.2. Convergence with Residual

Now that the HDG method was shown to be more accurate, it is important to compare the convergence of the method with the FVM, since it is one of its main advantages.

Figures 6 and 7 show the results for the HDG and the FVM in function of the residual of the nonlinear solver. It should be noted that the FVM used the Picard iteration and the HDG used Newton's method as nonlinear solver. However, the differences shown are due the discretization method and not to the nonlinear solver, since the nonlinear solver influences mainly the number of iterations. The residual is normalized by the largest residual in these simulations, which is in this case 1×10^{-2}. These values are plotted in the $x-$axis.

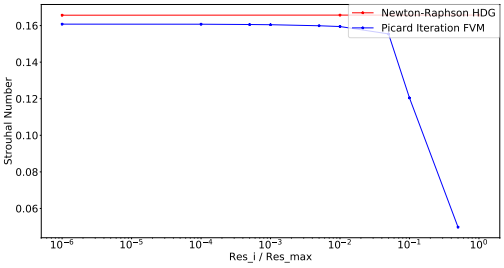

Figure 6: FVM and HDG residual convergence analysis of the Strouhal number.

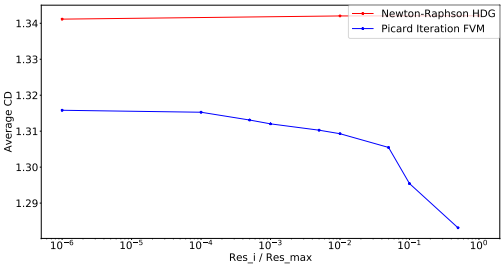

Figure 7: FVM and HDG residual convergence analysis of the drag coefficient.

The HDG method has little change in function of

the residual. The FVM requires a lower residual to converge. Mainly, when the residual is more than 1×10^{-4} the evolution in the solution is more emphasized. Additionally, the FVM method required more than 10 the number of iterations of the HDG to reach a residual of 1×10^{-8}. The HDG method only required 3 iterations on average to reach a residual of 1×10^{-8} (6 if the Picard iteration is used).

The methods converged to a different value of average drag coefficient, however the difference is only of 2%. This is due to the higher-order convergence of the HDG method. To obtain the same value, the FVM should have a thinner mesh.

The Hybridizable Discontinuous Galerkin method is a viable option to the Finite Volume Method. The solutions provided by both methods are equal, with the HDG method requiring less iterations, since the convergence is faster, due to the use of Newton's method and the coupling of the Navier–Stokes equations. However, the HDG method has more degrees of freedom than the FVM, despite both meshes having the same number of nodes. Since the HDG method is more exact with 10 times less iterations and with a less refined mesh, this method manages to be more efficient than the FVM.

5. Turbulent Flow

In this section the robustness of the HDG method will be analysed. The method's limits will be tested without any stabilization by computing the flow over a flat plate and a NACA0012 airfoil. For this simulation, the Spalart–Allmaras turbulence model will be implemented in the HDG code developed by Paipuri (2018).

5.1. Flat Plate Laminar Flow

The first step is to test the limits of the HDG method without any turbulence model or stabilization techniques. It is important to know the maximum Reynolds number the method is capable of converging without the turbulence model so that its can be correctly predicted.

The flow over a flat plate is important due to its simplicity, since only one boundary layer is being computed and there is a zero pressure gradient. There are plenty of numerical and experimental results, which make this simulation typical for CFD codes. There are also analytical approximations to the behaviour of the boundary layer, which will be used to compare the results obtained.

The plate is 1 m in length, the domain extends 0.25 m before, after and above the plate. Only in the plate is a no-slip boundary condition applied. The left boundary is the inlet and the remaining boundaries have a Neumann boundary condition applied (the normal gradient of the velocity and the pressure are zero).

A Cartesian mesh is used with a refinement near the plate so that the dimensionless wall distance y^+ is equal to one in the first element. Near the stagnation point there is also a refinement due to the higher gradients present in this area. The mesh has 64×64 elements, all second-order quadrilaterals. In figure 8 the plate is located at $y = 0$ and $0.25 \leq x \leq 1.25$.

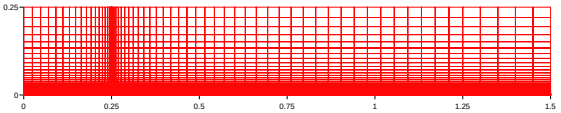

Figure 8: Flat plate mesh

Without the turbulence model, the method was capable of achieving a Reynolds number of 1×10^6. Although the method was able to converge for higher Reynolds numbers, due to the lack of a turbulence model these simulations were not accurate. The parameter τ remained 1 during these simulations.

The HDG method provided a drag coefficient of 1.344×10^{-3}. The Blasius solution provides the theoretical values of the drag coefficient for a laminar flow: $C_d = \frac{1.328}{Re^{1/2}}$, where Re is the Reynolds number (from White (1998)). The theoretical value for the drag coefficient is 1.328×10^{-3}, which is a difference of 1%. Figure 9 shows the shear stress distribution along the plate computed using the HDG method and the theoretical shear stress distribution (from White (1998)).

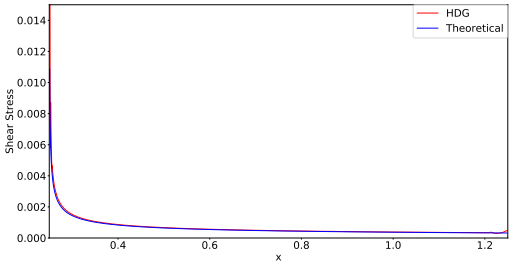

Figure 9: Shear stress distribution for flat plate at Reynolds 1×10^6

The main differences are near the leading and trailing edges of the plate. The analytical solution is not valid near the leading and trailing edges of the plate. Refinement in the trailing edge of the plate should have been used to reduce the oscillations.

The HDG method was able to correctly predict a high Reynolds number flow without the assist of a turbulence model and without the need to increase

the τ parameter. For higher Reynolds number flows it is necessary to implement a turbulence model.

5.2. Naca Laminar Flow

The flow around the NACA0012 airfoil is much more complex than the flat plate, mainly due to the merger of the upper and lower boundary layers in the wake. This merger is the main cause for instabilities due to the inflexion points of the velocity profile.

In this case, the parameter τ_u needs to be increased with the Reynolds number. Despite reported in many papers that the HDG method is stable for a wide range of τ_u (Cockburn et al. (2011)), its value was found to be limited by two problems. If the τ_u parameter is too low, convergence is not achieved due to instabilities in the flow. In opposition, if τ_u is too high, the solution obtained is not correct, due to the τ_u parameter's dominance in the matrix.

The NACA0012 airfoil is at zero angle of attack. Since the airfoil is symmetric, the lift of the airfoil should be null. For a larger τ, despite the residual being low, the lift of the airfoil is not equal zero.

The mesh used is a circular mesh, with an outer diameter of 26 m. The airfoil has a chord of 1 m and is centred in the mesh. The elements are second-order quadrilaterals with a refinement near the airfoil (the height of the first element must have a dimensionless wall distance y^+ of one). The mesh is composed of 4096 elements.

The inlet velocity is set to 1 m/s and the Reynolds number is imposed through the kinematic viscosity. A no-slip condition is applied on the airfoil nodes. The outer nodes are split into the inlet and the outlet, where the Dirichlet condition for the velocity and the Neumann condition for the pressure and velocity gradient are imposed respectively .

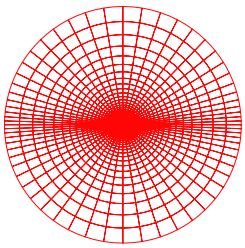

Figure 10: Mesh for the NACA0012 airfoil

Using a τ_u of 1 it was found that the method stopped converging for a Reynolds higher than 2×10^4. For this Reynolds, the HDG method was still capable of converging with a residual lower than 1×10^{-10} without stabilization. For this simulation, the lift coefficient was -1.5×10^{-5}. Since this is a numeric method, the lift coefficient is not exactly zero, however, the lift coefficient is 1000 smaller than the drag coefficient.

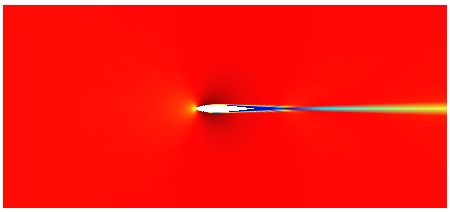

Figure 11: Velocity distribution on the NACA0012 at zero angle of attack and Reynolds 2×10^4 with a τ_u of 1

Since for Reynolds 2×10^4 the HDG method converges and provides a lift coefficient of practically zero, the τ_u parameter influence was analysed at Reynolds 4×10^4. The minimum τ_u to stabilize the solution is 16 000. Despite converging with a residual of 1×10^{-8}, the lift coefficient obtained is -0.1923, which is more than 10 times higher than the drag coefficient. This is due to the order of magnitude of the τ_u parameter, which dominates the stiffness matrix of the system. Consequently, the residual is lowered.

In conclusion, the HDG managed to converge to Reynolds numbers lower than 2×10^4 with minimal error using a τ_u of unity. Consequently, the method only has problems at high Reynolds, were the turbulence affects the stability. Using the τ_u parameter to stabilize the method is not appropriate. Consequently, the τ_u parameter will only be used to improve the convergence of already stable flows. If the simulation is unstable, the required τ_u parameter to stabilize the method would increase the error in the solution, making it useless.

5.3. Spalart–Allmaras Turbulence Model Implementation

The one equation Spalart–Allmaras turbulence model was introduced in Spalart and Allmaras (1994). In Allmaras et al. (2012) the Negative Spalart–Allmaras was created to deal with negative values of the SA working variable.

Merging the equations of both models into a single equation, it is possible to obtain equation (3), valid for both positive and negative values of the SA working variable.

$$\begin{aligned} \nabla \cdot \left((\boldsymbol{u}\tilde{\nu}) - \frac{1}{\sigma} \cdot ((\nu + \tilde{\nu} f_n) \nabla \tilde{\nu}) \right) = C_{b1} \tilde{S} \tilde{\nu} - \\ - C_{w1} f_w \left(\frac{\tilde{\nu}}{d} \right)^2 + \frac{C_{b2}}{\sigma} |\nabla \tilde{\nu}|^2 \end{aligned} \tag{3}$$

The SA working variable is $\tilde{\nu}$ (undamped eddy viscosity). The eddy viscosity ν_t is given by: $\nu_t =$

$\tilde{\nu} f_{v1}$. The remaining parameter are:

$$\chi = \frac{\tilde{\nu}}{\nu}, \tag{4a}$$

$$S = \sqrt{2\Omega_{ij}\Omega_{ij}}, \tag{4b}$$

$$f_{v2} = 1 - \frac{\chi}{1 + \chi f_{v1}}, \tag{4c}$$

$$g = r + C_{w2}\left(r^6 - r\right), \tag{4d}$$

$$r = \min\left[\frac{\tilde{\nu}}{\tilde{S}k^2d^2}, 10\right], \tag{4e}$$

$$f_{v1} = \begin{cases} \frac{\chi^3}{\chi^3 + C_{v1}^3} & \text{if } \tilde{\nu} \geq 0 \\ 0 & \text{if } \tilde{\nu} < 0 \end{cases}, \tag{4f}$$

$$\tilde{S} = \begin{cases} S + \frac{\tilde{\nu}}{k^2d^2} f_{v2} & \text{if } \tilde{\nu} \geq 0 \\ S & \text{if } \tilde{\nu} < 0 \end{cases}, \tag{4g}$$

$$f_w = \begin{cases} g\left(\frac{1+C_{w3}^6}{g^6 + C_{w3}^6}\right)^{1/6} & \text{if } \tilde{\nu} \geq 0 \\ -1 & \text{if } \tilde{\nu} < 0 \end{cases}, \tag{4h}$$

$$f_n = \begin{cases} 1 & \text{if } \tilde{\nu} \geq 0 \\ \frac{C_{n1} + \chi^3}{C_{n1} - \chi^3} & \text{if } \tilde{\nu} < 0 \end{cases}. \tag{4i}$$

In equation (4b) Ω_{ij} is the vorticity, $\Omega_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right)$. The constants are given by: $C_{b1} = 0.1355$; $\sigma = 2/3$; $C_{b2} = 0.622$; $k = 0.41$; $C_{w2} = 0.3$; $C_{w3} = 2$; $C_{v1} = 7.1$; $C_{w1} = \frac{C_{b1}}{k^2} + \frac{1+C_{b2}}{\sigma}$.

The variable d denotes the distance from the field point to the nearest wall. When the distance to the wall increases the source terms decrease, mainly the destruction term goes to zero. The generation of turbulence depends on both the distance to the wall and the vorticity of the flow, since when the velocity gradients are higher there is more turbulence.

The last step is to define the boundary conditions used for the model. In walls, such as the airfoil surface, the turbulent viscosity is set to zero, therefore $\tilde{\nu} = 0$. The outlets are defined as convective outlets, meaning the normal component of the turbulent viscosity's gradient is zero, $\nabla\tilde{\nu} \cdot \boldsymbol{n} = 0$. In the inlet, ideally the condition is that $\tilde{\nu}$ is also set to zero, however many solvers have problems with this definition. For this reason, the condition in (5) is applied to the inlet.

$$\frac{\tilde{\nu}}{\nu}\frac{1}{\mathrm{Re}} = 10^{-6} \tag{5}$$

The Spalart–Allmaras model equation is coupled with the Navier–Stokes equations seen previously. Now, there is a new variable $\tilde{\nu}$ and a new trace variable, $\hat{\tilde{\nu}}$, to be solved in the global problem. A new $\tau_{\tilde{\nu}}$ parameter, now related to the jump of the SA working variable is created to enforce the continuity of flux of turbulence. However, since the convective and diffusive fluxes have the same order of magnitude for both the SA and the Navier–Stokes equations, and these fluxes define the order of the parameter τ, this parameter will be the same both the turbulence model and the Navier–Stokes equations.

The correct linearization of the source term is essential to obtain a stable method. The generation source term was assumed to be constant in each iteration and the destruction and cross diffusion source terms were linearized.

Finally, it is important that the meshes used for the SA model have a dimensionless wall distance y^+ approximately equal to one.

5.4. Flat Plate Turbulent Flow

The turbulent flow over the flat plate will be computed in this section. Due to the simplicity of the flow, it is very simple to understand if the turbulence model was correctly implemented.

It is possible to reach a Reynolds of 1×10^5 if the τ parameter is increased to 1900. For higher Reynolds numbers the required τ would make the solution invalid, consequently this is the maximum Reynolds achieved.

The drag coefficient is now 4.103×10^{-3}. Following the theoretical model in White (1998), the flow is laminar since the Reynolds is below 5×10^5, which is in accord with the HDG result. The theoretical drag coefficient has a value of 4.199×10^{-3}. In this case the HDG method has an underestimation of 2%. Figure 12 shows the turbulent variable of the SA equation.

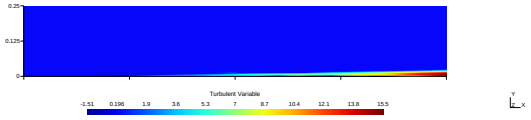

Figure 12: Flat plate turbulent variable $\tilde{\nu}/\nu$ field at Reynolds 1×10^5

The eddy viscosity reaches 15 times the kinematic viscosity, which implies the effective viscosity in some part of the domain is sixteen times the normal viscosity. However, this only occurs on the wake due to the low Reynolds number of the flow.

Although the turbulent variable $\tilde{\nu}$ is still not relevant over the plate, there is a clear increment in the variable along the plate. In this simulation, there is a negative value of the SA variable near the end of the plate, which is normally due to a lack of definition of the mesh used.

Figure 13 shows the shear stress distribution along the plate. The differences in this case are clearly bigger due to the turbulence model. Following the theoretical data, this flow is laminar, consequently the shear stress distribution was computed using the laminar equation. However, as seen from figure 12 the eddy viscosity in this case is no longer negligible, specially near the wake, where the difference is more substantial. The oscillation of the

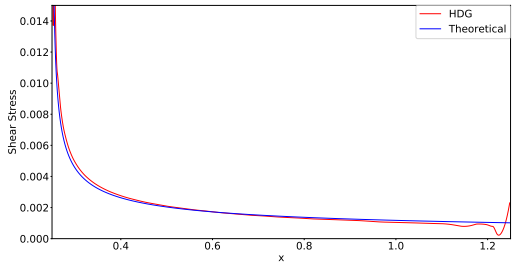

Figure 13: Flat plate shear stress at Reynolds 1×10^5

shear stress near the trailing edge are a consequence of the large Δx. The same refinement as in the leading edge should have been used in the end of the plate to reduce the oscillations.

Although the HDG method was not able to converge for flows with Reynolds number high enough for the flow to be turbulent, it correctly predicted the behaviour of the turbulence variable.

The results suggest that the coupling between the momentum and the turbulence model requires a more robust approach.

5.5. NACA 0012 Turbulent

The flow around the NACA0012 using the Spalart–Allmaras turbulence model is computed in this section.

The maximum Reynolds achieved was 1×10^4, above which the method was unable to converge even with an increment of the τ parameter. With this Reynolds a completely laminar flow is expected around the airfoil, with the turbulent variable increasing in the wake.

Due to the merger of the top and bottom boundary layers, the turbulence in the wake is much higher.

The drag coefficient is 3.699×10^{-2} and the lift obtained is -9.509×10^{-6}. Due to numerical errors it is not possible to obtain a lift value closer to zero.

The turbulent variable $\tilde{\nu}$ only increased in the wake of the airfoil, as expected. Figure 14 shows the turbulence field around the airfoil in greater detail.

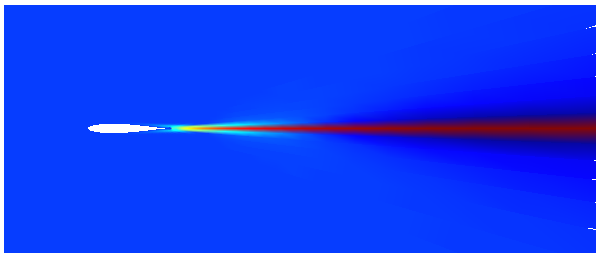

Figure 14: NACA0012 wake turbulent variable $\tilde{\nu}$ field at Reynolds 1×10^4

Near the trailing edge of the airfoil the increase of the turbulent variable $\tilde{\nu}$ begins, the maximum value of 5.64×10^{-4} for the turbulent variable almost triples the effective viscosity of the flow in that region. However, since the velocity stabilizes far from trailing edge, the generation of turbulence stops.

In addition to influencing the convergence and stability of the method, the linearization of the source terms also influences the solution obtained. As mentioned before, the results showed previously were computed linearizing the destruction and cross diffusion terms while the generation term remained constant in each iteration.

If the generation of turbulence is also linearized, the method becomes more unstable. The residual achieved is not as low as before (the incremental residual is only 1×10^{-2}). The most noticeable difference in this discretization is in the turbulent variable of the Spalart–Allmaras equation. In this case, the turbulent variable $\tilde{\nu}$ is negative in most of the wake. Due to the poorer convergence this implementation is not appropriate.

It is possible to only linearize the destruction term , with the remaining terms constant in each iteration. In this case, the method requires an increment of the τ parameter to be stable even at Reynolds 1×10^4. The parameter τ was increased to 1000 to converge the method. Due to the lack of linearization of the cross diffusion there is a decrease in the turbulent variable in the wake far from the airfoil. Here, negative values of the turbulent variable dominate most of the wake, consequently this solution is not appropriate either.

It is also possible to move all the source terms to the right side of the system, where they are constant in each iteration of the nonlinear solver. This handling makes the method even harder to converge. It was not possible to achieve a Reynolds of 1×10^4 even with an increment of the τ parameter.

The best implementation found was the one described in 5.3, where only the generation term is constant and the destruction and cross diffusion terms are linearized. There is an increase in the convergence difficulties if the source terms are linearized differently.

Manipulating the Spalart–Allmaras turbulence equation, it is possible to solve for the variable χ (equation (4a)) instead of the variable $\tilde{\nu}$. Due to the solution strategy being used, where the Reynolds number is incremented in steps and is imposed through the kinematic viscosity, the turbulent field in function of χ might provide a more adequate initial guess in each simulation. However, the convergence of the method with this implementation was not as good.

Relaxation of the velocity and the turbulence

variable was attempted to help stabilize the method. However, since the problem at high Reynolds number was not a stagnation of the residual but an increase of the latter, relaxation was unable to converge the method.

From both the analysis of the flat plate and the airfoil, it is possible to conclude the method is only able to converge when the coupling between the momentum equations and the turbulence model is weak, meaning the eddy viscosity does not affect the Navier–Stokes equations.

6. Conclusions

The HDG method has more degrees of freedom than the FVM for the same number of nodes. Despite this increase in variables, the method is still more computationally efficient to achieve the same goal. Even with a low-order approximation (only second-order), the method performs better since it requires less elements and the convergence is achieved in a fraction of the iterations, since Newton's method is used and the equations are coupled.

Although the method was able to converge for low Reynolds numbers using the one equation Spalart–Allmaras turbulence model without any stabilization, there were problems in the convergence of the implemented model. There was a clear destabilization in the flow over the flat plate when the turbulence model was implemented, which implies the method either requires stabilization in the turbulence model or Newton's nonlinear solver is unable to solve the equations coupled. The method only converges when the coupling between the Navier–Stokes equations and the turbulent model is weak.

In the future, stabilization terms should be added to this implementation of the method to test its robustness at high Reynolds numbers.

Acknowledgements

I would like to thank the professors at Técnico that helped me through this process. Professor Luís Eça and Professor Carlos Tiago were essential for the completion of this project. I would also like to thank Mahendra Paipuri for his brilliant program that was the basis of this thesis.

References

S. R. Allmaras, F. T. Johnson, and P. R. Spalart. Modifications and Clarifications for the Implementation of the Spalart–Allmaras Turbulence Model. *7th International Conference on Computational Fluid Dynamics*, July 2012.

B. Cockburn, B. Dong, J. Guzman, M. Restelli, and R. Sacco. A hybridizable discontinuous Galerkin method for steady-state convection-diffusion-reaction problems. *Journal on Scientific Computing*, 31(6):3827:3846, 2009a.

B. Cockburn, J. Gopalakrishnan, and R. Lazarov. Unified hybridization of discontinuous Galerkin, mixed and continuous Galerkin methods for second order elliptic problems. *SIAM Journal on Numerical Analysis*, 47(2):1319:1365, 2009b.

B. Cockburn, J. Gopalakrishnan, N. C. Nguyen, J. Peraire, and F. J. Sayas. Analysis of HDG methods for Stokes flow. *Mathematics of Computation*, 80(274):723:760, 2011.

Joel H. Ferziger and Milovan Peric. *Computational Methods for Fluid Dynamics*. Springer, third edition, 2002.

MARIN. ReFRESCO, 2015. URL `http://www.refresco.org/`. Accessed: 2018-09-18.

D. Moro, N. C. Nguyeny, and J. Peraire. Navier–Stokes Solution Using Hybridizable Discontinuous Galerkin methods. *AIAA 2011-3407*, 2011.

C. Nguyen, J. Peraire, and B. Cockburn. An implicit high-order hybridizable discontinuous Galerkin method the incompressible Navier–Stokes equations. *Journal of Computational Physics*, 230:1147:11705, 2011.

M. Paipuri. *Comparison and coupling of continuous and hybridizable discontinuous Galerkin methods: application to multi-physics problem*. PhD thesis, Instituto Superior Técnico, Universidade de Lisboa, 2018.

P. R. Spalart and S. R. Allmaras. A One-Equation Turbulence Model for Aerodynamic Flows. *Recherche Aerospatiale*, 1:5:21, 1994.

Frank M. White. *Fluid Mechanics*. WCB McGraw-Hill, fourth edition, 1998.