This is a postprint of the following article: Sabet Seraj, Anil Yildirim, Joshua L. Anibal, Joaquim R. R. A. Martins. "Dissipation and time step scaling strategies for low and high Mach number flows". *Journal of Computational Physics* (2023). The published article may differ from this postprint and is available at: https://doi.org/10.1016/j.jcp.2023.112358.

Dissipation and time step scaling strategies for low and high Mach number flows

Sabet Seraj, Anil Yildirim, Joshua L. Anibal, Joaquim R. R. A. Martins Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109, USA

Abstract

Aircraft aerodynamic design often involves evaluating flow conditions that span low subsonic to transonic or even supersonic Mach numbers. Compressible flow solvers are a natural choice for such design problems, but these solvers suffer from reduced accuracy and efficiency at low Mach numbers. In addition, simulations with supersonic conditions can be challenging to converge because of large gradients in the flow field. This paper presents three contributions to address these issues in the context of an approximate Newton–Krylov solver for the Reynolds-averaged Navier–Stokes equations. First, we propose a method for scaling the artificial dissipation terms in the Jameson–Schmidt–Turkel scheme to improve its accuracy at low Mach numbers while retaining the simplicity of the original scalar dissipation formulation. Second, we show that characteristic time-stepping combined with an approximate Newton method can accelerate convergence for low Mach number flows by reducing the stiffness in the linear system for each Newton iteration. Third, we introduce a dissipation-based continuation method for flows with shocks that improves robustness and accelerates convergence without sacrificing accuracy. These methods can make compressible flow solvers more accurate and efficient across low and high Mach number regimes.

1 Introduction

Aircraft aerodynamic design often requires evaluating performance over a wide range of flow conditions. For example, the flight envelope of a supersonic transport aircraft can span Mach numbers from 0.25 to 2.0 [1]. A conventional transonic aircraft will fly through Mach numbers from 0.2 to 0.85 in a typical mission [2]. In addition, certain applications produce flow fields with a range of Mach numbers at a single operating condition. Helicopter rotors experience nearly incompressible flow at blade roots and compressible flow at the tips [3]. Nacelles in crosswind can experience supersonic flow around the inlet lips despite low subsonic flow outside the nacelle [4]. The Mach number range in applications such as these motivates the need for CFD solvers that are accurate and efficient across a wide range of Mach numbers.

In theory, compressible flow solvers represent the actual physics across all Mach numbers, making them a natural choice for aerodynamic design. In practice, these solvers commonly suffer from reduced accuracy and speed at low Mach numbers. Modifying compressible flow solvers to resolve these issues is usually called *low-speed preconditioning*. Low-speed preconditioning has been studied in depth by several authors for both explicit and implicit time-stepping schemes. Turkel [5] introduced a class of preconditioning matrices for low-speed flows based on analysis of the Euler equations. This approach modifies the time-stepping terms to equalize the advective and acoustic wave speeds, accelerating convergence. The same preconditioners can improve accuracy at low Mach numbers if the artificial dissipation in the scheme is based on preconditioned fluxes [6]. This approach has also been applied to viscous flows [7, 8, 4]. van Leer et al. [9] proposed a preconditioner for all Mach numbers that is also based on the idea of minimizing the spread in wave speeds. They interpreted this approach as using different time steps along different characteristic directions, which led to the name "characteristic time-stepping". The authors demonstrated the preconditioner's effectiveness on the Euler equations at Mach numbers from 0.01 to 1.8 and on viscous flow at Mach 7.95. Lee [10] showed that preconditioning approaches based on the Euler equations work well for viscous flows in most practical cases.

Some work has also been done on low-speed preconditioning for Newton-based solvers. Newtonbased solvers are fully implicit and offer more flexibility in addressing the accuracy and speed issues independently without introducing unwelcome stability restrictions. To speed up convergence with Newton–Krylov solvers, Knoll et al. [11] considered multiplicative Schwarz preconditioning, and Weston et al. [12] used approximate block factorizations. Both of these papers take an algebraic approach to preconditioning. In contrast, Mary et al. [13] applied characteristic-based preconditioning to an approximate Newton solver. This work was limited to a maximum Reynolds number of 800 and primarily considered accuracy. These limitations suggest an unexplored potential for improved performance with Newton-based solvers using characteristic preconditioning methods.

In addition, high-speed flows have not been explored as extensively as low-speed cases. Kaushik et al. [14] and Olawsky et al. [15] found that Newton-based solvers can run into difficulties converging supersonic flows with second-order schemes compared to more dissipative first-order schemes. They resolved this by starting the solution with a first-order scheme before switching to a second-order scheme closer to convergence. This approach lacks reliability because it is not always clear when to switch to a second-order scheme for a given case. A related approach for improving convergence for flows with shocks is to increase the artificial dissipation in the discretization. For upwind schemes, this can be achieved by limiting the minimum eigenvalue of the flux Jacobian, commonly referred to as an *entropy fix*. For example, changing the entropy fix to be more dissipative is an option in the widely used FUN3D code [16]. The amount of dissipation in upwind schemes can also be increased by using a more dissipative flux limiter [17]. The equivalent for central schemes is increasing the artificial dissipation constants, as done by Öhrman [18] and Aprovitola et al. [19], for example. Increasing dissipation improves robustness in exchange for reduced accuracy. We desire an automated approach that increases robustness without sacrificing accuracy.

In this work, we introduce three methods to improve the efficiency and accuracy of Newton-based compressible flow solvers across low and high Mach number regimes. In Sec. 3, we describe a simple method for improving low-speed accuracy by scaling the artificial dissipation of a scalar dissipation scheme. This method preserves the simplicity of the original scheme and is easy to implement. We then formulate a characteristic time-stepping approach for an approximate Newton solver in Sec. 4. Characteristic time-stepping improves nonlinear convergence rates by reducing the stiffness of the linear system at each nonlinear iteration. Lastly, we present a dissipation-based continuation method in Sec. 5 that addresses nonlinear convergence difficulties that arise from high Mach number flows.

This method leverages the idea that more dissipative methods are easier to converge by starting with high artificial dissipation and smoothly reducing the dissipation to the desired level as the solution converges.

2 Baseline RANS solver

The CFD solver in this work is ADflow [20], a second-order finite volume code for multiblock and overset structured meshes. We solve the steady, compressible Reynolds-averaged Navier–Stokes (RANS) with the Spalart–Allmaras (SA) turbulence model [21]. The flow is assumed to be fully turbulent.

The Navier-Stokes equations can be written as

$$\frac{\mathrm{d}}{\mathrm{d}t} \int_{V} Q \mathrm{d}V + \oint_{\partial V} \vec{F}_{c} \cdot \vec{n} \mathrm{d}A - \oint_{\partial V} \vec{F}_{d} \cdot \vec{n} \mathrm{d}A = 0, \qquad (1)$$

where $Q = [\rho, \rho u, \rho v, \rho w, \rho E]^T$ represents the conservative variables, \vec{F}_c is the convective flux, and \vec{F}_d is the diffusive flux. This system is augmented by the SA equation, which adds the turbulent contribution to the diffusive flux. Discretizing in space with a cell-centered finite volume method, the equation for a cell *i* is

$$\frac{\mathrm{d}Q_i}{\mathrm{d}t}V_i + \sum_f \hat{F}_c \Delta A_f - \sum_f \hat{F}_d \Delta A_f = 0, \qquad (2)$$

where the discretized fluxes, \hat{F}_c and \hat{F}_d , are summed over the faces of the cell. We use the Jameson–Schmidt–Turkel (JST) scheme [22, 23] to discretize the convective flux and a central scheme with Green–Gauss gradients for the diffusive flux.

2.1 Artificial dissipation formulation

JST is a central scheme that combines second and fourth difference artificial dissipation terms for stability. We now describe the artificial dissipation formulation in the *I* mesh direction. The formulation for the *J* and *K* directions follow similarly. The second-difference dissipation coefficient at a face between cell *i* and i + 1 is given by

$$\epsilon_{i+\frac{1}{2}}^{(2)} = \kappa_2 \Upsilon_{i+\frac{1}{2}} \hat{\Lambda}_{i+\frac{1}{2}}^I, \tag{3}$$

and the fourth-difference coefficient is given by

$$\epsilon_{i+\frac{1}{2}}^{(4)} = \max\left(0, \kappa_4 \hat{\Lambda}_{i+\frac{1}{2}}^I - c_4 \epsilon_{i+\frac{1}{2}}^{(2)}\right).$$
(4)

By default, we use $\kappa_2 = 0.25$, $\kappa_4 = 0.0156$, $c_4 = 1$. $\Upsilon_{i+1/2}$ is the shock sensor at the face, which ensures that second-difference dissipation is only applied near shocks and fourth-difference dissipation is only applied away from shocks. The shock sensor at the face is computed as the maximum of the neighboring cell sensor values:

$$\Upsilon_{i+\frac{1}{2}} = \max\left(\Upsilon_i, \Upsilon_{i+1}\right). \tag{5}$$

For viscous flows, the shock sensor value in each cell is based on the local change in entropy, *s*:

$$\Upsilon_i = \frac{|s_{i+1} - 2s + s_{i-1}|}{|s_{i+1} + 2s + s_{i-1}|}.$$
(6)

 $\hat{\Lambda}_{i+1/2}^{I}$ is the spectral radius of the flux Jacobian in the *I* direction at the face and is computed as the average of the neighboring cell radii:

$$\hat{\Lambda}_{i+\frac{1}{2}}^{I} = \frac{1}{2} \left(\hat{\Lambda}_{i}^{I} + \hat{\Lambda}_{i+1}^{I} \right).$$
⁽⁷⁾

To compute the spectral radius in the *I* direction, we first compute the isotropic spectral radius, Λ , in each direction. The isotropic spectral radius in the *I* direction is

$$\Lambda^{I} = \vec{U} \cdot \vec{n}^{I} + c |\vec{n}^{I}|, \tag{8}$$

where $\vec{U} = [u, v, w]^T$ is the local velocity vector, *c* is the local speed of sound, and \vec{n}^I is the cell's average face area normal in the *I* direction. Similarly, for the *J* and *K* directions, we have

$$\Lambda^{J} = \vec{U} \cdot \vec{n}^{J} + c |\vec{n}^{J}|, \qquad (9)$$

$$\Lambda^{K} = \vec{U} \cdot \vec{n}^{K} + c |\vec{n}^{K}|. \tag{10}$$

We then compute a scaled spectral radius that is appropriate for high aspect ratio cells as suggested by Martinelli [24]:

$$\hat{\Lambda}^{I} = \Lambda^{I} \left(1 + \left(\frac{\Lambda^{J}}{\Lambda^{I}} \right)^{0.67} + \left(\frac{\Lambda^{K}}{\Lambda^{I}} \right)^{0.67} \right).$$
(11)

The artificial dissipation terms are proportional to the spectral radius, so appropriately scaling the spectral radius is critical for solver convergence and solution accuracy.

2.2 Solver algorithm

To converge the semi-discrete form (Eq. 2) to a steady solution, we define the residual in each cell, R_i , as the time-derivative term and drive it to zero:

$$R_i = \frac{\mathrm{d}Q_i}{\mathrm{d}t} V_i = \sum_f \hat{F}_d \Delta A_f - \sum_f \hat{F}_c \Delta A_f = 0.$$
(12)

The baseline nonlinear solver algorithm in this work is the approximate Newton–Krylov (ANK) method with pseudo-transient continuation implemented in ADflow [25]. The solver is primarily designed for robustness and speed at transonic flow conditions. The solution update at each iteration of the ANK solver, $\Delta Q^{(n)}$, is obtained by solving the linear system

$$\left[\left(T^{-1} \right)^{(n)} + \left(\frac{\partial R}{\partial Q} \right)^{(n)} \right] \Delta Q^{(n)} = -R(Q^{(n)}), \tag{13}$$

where *R* is the residual vector, $Q^{(n)}$ is the state vector, and *T* is the time-step matrix. *T* is a diagonal matrix, where the diagonal terms for each cell are the time step, Δt_i , that results in the desired global CFL number. The global CFL number is initially small and is increased as the solution converges. The SA turbulence equation is solved separately after each ANK update for the flow equations. A complete description of the ANK algorithm is given by Yildirim et al. [25].

2.3 Accuracy and cost comparisons

We compare the methods presented in this paper to the baseline solver in terms of accuracy and cost. Unless stated otherwise, all solutions are converged to a total residual of 10⁻¹⁰ relative to the freestream residual. To provide a hardware-independent measure of cost, we report total CPU times in TauBench work units [26]. We run all computations on NASA's Aitken supercomputer using Cascade Lake nodes, where one work unit equals 3.9078 processor seconds. We determine this by running TauBench ten times with the command mpiexec -np 1 ./TauBench -n 250000 -s 10 and computing the average run time.

3 Scaling artificial dissipation to improve accuracy at low Mach numbers

Artificial dissipation is essential for nonlinear convergence but can compromise accuracy. The JST scalar dissipation scheme is generally more dissipative than second-order upwind schemes [27]. Central schemes that are less dissipative than scalar dissipation include matrix dissipation [27] and the convective-upstream-split-pressure scheme [28]. However, these central and upwind schemes are primarily designed for transonic flows and run into accuracy issues at low Mach numbers because of improperly scaled artificial dissipation [6]. One common approach to scale the dissipation is to multiply the flux Jacobian by a local preconditioning matrix. For JST, this modifies the spectral radius and, consequently, the dissipation. Using a local preconditioning matrix can also accelerate convergence, which we discuss in Sec. 4. To improve low-speed accuracy, we opt to modify the spectral radius directly. The advantage of this approach over prior ones is that it requires minimal changes to the JST scalar dissipation formulation.

In the low Mach number limit, the advective contribution to the spectral radius (Eq. 8) is O(1), whereas the acoustic contribution is O(1/M) [29, 6]. Rieper [30] showed that artificial dissipation must be independent of the Mach number for accuracy in the incompressible limit. To scale the artificial dissipation for low Mach numbers, we introduce the acoustic scaling factor ζ in the isotropic spectral radius:

$$\Lambda^{I} = \vec{U} \cdot \vec{n}^{I} + \zeta c |\vec{n}^{I}|. \tag{14}$$

Using this formulation, we can select values for ζ that reduce or eliminate the dependence of artificial dissipation on the Mach number. Similar dissipation scaling approaches have been proposed in the context of upwind schemes [31] and kinetic energy preserving schemes [32].

To study the effect of ζ on low-speed accuracy, we consider a NACA 0012 airfoil at a Reynolds number of 10⁶, 3° angle of attack, and Mach numbers from 0.01 to 0.4. We use an O-mesh with 296 cells around the airfoil and 128 cells in the offwall direction, for a total of 37888 cells (Fig. 1). The initial offwall spacing is 10⁻⁶ chord lengths, and the farfield is 100 chord lengths away from the airfoil, which is consistent with the Drag Prediction Workshop guidelines [33].

To evaluate accuracy at different values of ζ , we compare the RANS results with predictions using XFOIL [34], a 2D panel method coupled with an integral boundary layer model. We use 300 panels for the inviscid discretization in the XFOIL analyses. Increasing the number of panels to 360 gives the same coefficient values to 4 significant digits. We force transition at the 0.3% chord location to match the fully turbulent RANS simulations. This is the smallest positive value we could use for the trip location without encountering numerical difficulties.



Figure 1: NACA 0012 O-mesh with 296 × 128 cells

Fig. 2a compares the drag for the baseline ($\zeta = 1$), no acoustic contribution ($\zeta = 0$), and scaling the acoustic contribution by the freestream Mach number ($\zeta = M_{\infty}$). With no scaling, drag diverges as the Mach number is reduced. The drag overprediction at low Mach numbers for the baseline is caused by a nonphysical suction peak on the trailing edge surface (Fig. 3). These nonphysical pressure fluctuations are a direct result of improperly scaled dissipation, as shown by Guillard and Viozat [35]. Chen et al. [36] also found similar trailing edge pressure spikes for low-speed inviscid flow. Scaling the acoustic contribution removes the trailing edge artifact and results in more accurate drag values. In addition, the change in drag flattens out as the Mach number is reduced. Scaling by the freestream Mach number yields the most accurate drag predictions compared to XFOIL.



Figure 2: Scaling artificial dissipation improves accuracy

The baseline errors in the lift at low Mach numbers are not as significant as for drag (Fig. 2b), but the lift still diverges as the Mach number approaches zero. Reducing artificial dissipation avoids



Figure 3: Reducing artificial dissipation removes the spurious pressure spikes at the trailing edge (Mach 0.01)

divergence. The value of ζ is not critical for lift prediction because the differences are primarily caused by discretization errors. We run a mesh refinement study at Mach 0.1 to demonstrate this effect. In addition to the 296 × 128 mesh, we use a 592 × 256 mesh and a 148 × 64 mesh for the study. Figure 4 shows the drag and lift for these three meshes, as well as the Richardson extrapolation [37, 38] using the two finest meshes. The Richardson extrapolation estimates the values we would obtain on an infinitely fine mesh. The extrapolated drag values depend on ζ . The extrapolated drag for $\zeta = 1$ is 7.5% higher than for $\zeta = 0$ and 5% higher than for $\zeta = M_{\infty}$. This means that adequately scaling the dissipation is essential for drag prediction, even for fine meshes. On the other hand, the extrapolated lift values are within 0.5% of each other. The differences in the lift with ζ are significant only for the 148 × 64 mesh. However, such a coarse mesh is outside the asymptotic range [38] in terms of mesh size and should not be used in general. The mesh refinement study also shows that drag for $\zeta < 1$ is less sensitive to mesh size than baseline. As a result, scaling the artificial dissipation can indirectly speed up analyses by allowing for coarser meshes at the same accuracy as the baseline.

Based on these results, we recommend using $\zeta = M_{\infty}$ for cases where a freestream Mach number exists. For other cases, ζ should equal the dominant or average Mach number. An alternative approach would be to use local values for ζ instead of using the same value in all cells. For example, we could use the local Mach number instead of the freestream Mach number in each cell. For the NACA 0012 case, this results in similar accuracy and convergence to a global value of $\zeta = 0$. This is because the local Mach number near the airfoil is close to zero, and the cells closest to the airfoil dictate the overall behavior of the solver. However, local values may be more applicable for cases where a freestream Mach number is not defined. The proposed scaling approach with global or local ζ is easy to implement and improves the accuracy of compressible solvers at low Mach numbers.



Figure 4: Mesh refinement at Mach 0.1

4 Accelerating convergence at low Mach numbers with characteristic time-stepping

In addition to lower accuracy, compressible flow solvers face slower convergence at low Mach numbers. For the NACA 0012 case from Sec. 3 with the 296 × 128 mesh and $\zeta = 1$, the solution for Mach 0.01 is three times slower than Mach 0.4 (Fig. 5). The slower convergence results from the difference in magnitude between the advective and acoustic wave speeds. This results in a more restrictive upper bound on the time step for explicit solvers. For Newton-based solvers, this is reflected in the stiffness of the linear system that is solved at each nonlinear iteration.



Figure 5: Cost increases as Mach number decreases (NACA 0012)

The baseline ANK solver defaults are tuned for compressible flow. However, we use settings that are beneficial for low Mach number flows for all cases with freestream Mach numbers of 0.4 or lower. The main changes are using more linear solver iterations and a stronger preconditioner for the linear system. Solving the linear system more tightly for higher Mach numbers and flows with shocks is

not as beneficial for improving nonlinear convergence.

We use a local preconditioning or characteristic time-stepping (CTS) approach to accelerate convergence at low Mach numbers further. This approach involves modifying the time-stepping terms to equalize the advective and acoustic wave speeds. We refer to this approach as characteristic time-stepping to emphasize that we are not modifying the artificial dissipation using the same preconditioner. To apply CTS to the ANK solver, we multiply the time-step matrix by the preconditioning matrix *P*. The linear system with the ANK step (Eq. 13) requires the inverse of the time-step matrix, so it is more convenient to work directly with the inverse of the preconditioner, P^{-1} :

$$\left[\left(P^{-1}T^{-1} \right)^{(n)} + \left(\frac{\partial R}{\partial Q} \right)^{(n)} \right] \Delta Q^{(n)} = -R(Q^{(n)}).$$
(15)

We make one change to adapt the preconditioner for use with pseudo-transient continuation. When the CFL number is low, the linear system is not stiff, and the baseline solver regularly outperforms CTS. A similar lack of stiffness was noted by Turkel and Vatsa [39] for small physical time steps in a dual time-stepping context. To take advantage of the baseline solver in this situation, we use a blended preconditioner based on the CFL number:

$$P_{\text{blend}}^{-1} = \left(\frac{\text{CFL}}{\text{CFL}_{\text{max}}}\right)P^{-1} + \left(1 - \frac{\text{CFL}}{\text{CFL}_{\text{max}}}\right)I.$$
(16)

With this blending, the solver initially acts like the baseline and reaches the fully preconditioned method when the CFL number reaches its maximum value.

The inverse preconditioner, P^{-1} , is a block diagonal matrix. Each diagonal block corresponds to one cell, and we denote the preconditioner for cell *i* as P_i^{-1} . Several different preconditioning matrices can be used in this formulation. We consider the Turkel preconditioner [5], denoted here as P_T , and the van Leer–Lee–Roe (VLR) preconditioner [9], denoted as P_{VLR} .

4.1 Turkel preconditioner

It is convenient to write the Turkel preconditioner in terms of the Euler symmetrizing variables,

$$\tilde{Q} = [p/(\rho c), u, v, w, p - c^2 \rho]^T.$$
(17)

With these variables, the preconditioner for cell i is

$$\tilde{P}_{\mathrm{T},i}^{-1} = \begin{bmatrix} \frac{1}{M_{\mathrm{T}}^2} & 0 & 0 & 0 & 0\\ \frac{\alpha u}{cM_{\mathrm{T}}^2} & 1 & 0 & 0 & 0\\ \frac{\alpha v}{cM_{\mathrm{T}}^2} & 0 & 1 & 0 & 0\\ \frac{\alpha w}{cM_{\mathrm{T}}^2} & 0 & 0 & 1 & 0\\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$
(18)

where we use $M_{\rm T}$ instead of the usual symbol β to avoid confusion with the β introduced for the VLR preconditioner in Sec. 4.2. As in prior work, we define $M_{\rm T}$ as a truncated form of the local Mach number:

$$M_{\rm T}^2 = \min\left(1, \max(M^2, \phi M_{\infty}^2)\right).$$
⁽¹⁹⁾

The lower bound on M_T^2 prevents the matrix from becoming ill-conditioned. Using $\phi > 0$ bounds M_T^2 away from zero in cells where the local Mach number is nearly zero. The ill-conditioning is mainly an issue at stagnation points. We find that $\phi = 10^{-4}$ provides the best performance while avoiding convergence issues. Larger values of ϕ do not show speedup compared to the baseline, whereas smaller values fail to resolve the stagnation point issue.

The upper bound on M_T^2 , combined with defining α as

$$\alpha = 1 - \left(M_{\rm T}^2\right)^{10},\tag{20}$$

eliminates preconditioning for locally supersonic flow. For $M \ge 1$, $M_T^2 = 1$ and $\alpha = 0$, which is equivalent to no preconditioning. $(M_T^2)^{10}$ quickly approaches 0 for subsonic Mach numbers. Consequently, $\alpha \approx 1$ for Mach numbers less than 0.8. This results in better conditioning than $\alpha = 0$. Prior time-stepping approaches typically used the preconditioner with $\alpha = 0$ for its improved robustness despite not providing the optimal conditioning in the low Mach number limit. The difference in robustness has been attributed to $\alpha = 1$ having lower artificial dissipation than $\alpha = 0$ [6]. However, we do not use the same scaling for the dissipation as for the time step, so we avoid this robustness issue.

Finally, we transform the matrix to conservatives variables through

$$P_{\mathrm{T},i}^{-1} = S\tilde{P}_{\mathrm{T},i}^{-1}S^{-1},\tag{21}$$

where *S* is the transformation matrix from symmetrizing to conservative variables:

$$S = \begin{bmatrix} \frac{\rho}{c} & 0 & 0 & 0 & -\frac{1}{c^2} \\ \frac{\rho u}{c} & \rho & 0 & 0 & -\frac{u}{c^2} \\ \frac{\rho v}{c} & 0 & \rho & 0 & -\frac{v}{c^2} \\ \frac{\rho w}{c} & 0 & 0 & \rho & -\frac{w}{c^2} \\ c\rho\left(\frac{M^2}{2} + \frac{1}{\gamma - 1}\right) & \rho u & \rho v & \rho w & -\frac{M^2}{2} \end{bmatrix}.$$
(22)

4.2 van Leer-Lee-Roe (VLR) preconditioner

The VLR preconditioner is usually expressed in terms of the Euler symmetrizing variables and in a flow-aligned coordinate frame:

$$\bar{\tilde{P}}_{\text{VLR},i}^{-1} = \begin{bmatrix} \frac{\beta^2 + \tau}{M_{\text{VLR}}^2 \tau} & \frac{1}{M} & 0 & 0 & 0\\ \frac{1}{M} & 1 & 0 & 0 & 0\\ 0 & 0 & \frac{1}{\tau} & 0 & 0\\ 0 & 0 & 0 & \frac{1}{\tau} & 0\\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$
(23)

where

$$\beta = \begin{cases} \sqrt{1 - M_{\rm VLR}^2}, M < 1\\ \sqrt{M_{\rm VLR}^2 - 1}, M \ge 1 \end{cases}, \qquad \tau = \begin{cases} \sqrt{1 - M_{\rm VLR}^2}, M < 1\\ \sqrt{1 - 1/M_{\rm VLR}^2} + \epsilon, M \ge 1 \end{cases}.$$
(24)

We define M_{VLR}^2 like M_T^2 for the Turkel preconditioner but without any restrictions for supersonic flow:

$$M_{\rm VLR}^2 = \max(M^2, \phi M_\infty^2). \tag{25}$$

We use $\phi = 10^{-4}$, the same value we use with the Turkel preconditioner. We also set $\epsilon = 10^{-4}$ when defining τ (Eq. 24) to avoid dividing by zero for sonic flow. We transform the preconditioner first to Cartesian coordinates with the rotation matrix *Z* and then to conservatives variables with the state transformation matrix *S* (Eq. 22):

$$P_{\text{VLR},i}^{-1} = SZ\bar{\tilde{P}}_{\text{VLR},i}^{-1}Z^{T}S^{-1}.$$
(26)

The rotation matrix [40] is

$$Z = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \cos \alpha \cos \theta & -\sin \theta & -\sin \alpha \cos \theta & 0 \\ 0 & \cos \alpha \sin \theta & \cos \theta & -\sin \alpha \sin \theta & 0 \\ 0 & \sin \alpha & 0 & \cos \alpha & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$
 (27)

where

4.3 The importance of dissipation scaling

When introducing the VLR preconditioner, van Leer et al. [9] used an upwind scheme and explicit time-stepping. They found that applying preconditioning to only the time-stepping terms resulted in severe restrictions on the time step. To avoid these restrictions, it was necessary to use a preconditioned form of the dissipation matrix in the Roe flux. However, they were unclear on whether this necessity extends to central schemes such as scalar dissipation. We show that dissipation scaling is also necessary for scalar dissipation schemes and Newton-based solvers.

We first apply CTS to the NACA 0012 case at Mach 0.01. For this section, we only show results with the Turkel preconditioner. However, we see similar trends with VLR. We find that the performance strongly depends on the choice of ζ (Fig. 6). Using $\zeta = M_{\infty}$ results in faster convergence than the baseline, whereas CTS is slower than the baseline for $\zeta = 1$ and $\zeta = 0$. For $\zeta = 1$, CTS takes more than twice as long as the baseline to converge and exhibits noisy convergence below a relative residual of 10^{-7} . Convergence with CTS is noisier than the baseline, even for cases where CTS is faster. This is related to the preconditioner becoming ill-conditioned at stagnation points. The noise is worse for lower freestream Mach numbers (Fig. 7). In addition, using larger values of ϕ reduces the noise but results in slower convergence. The noise is most significant for $\zeta = 1$ because of the nonphysical stagnation regions shown in Fig. 3a. We conclude that applying local preconditioning to the time step is problematic when the flow solution is locally nonphysical. The slowdown with CTS is less severe for $\zeta = 0$ than $\zeta = 1$ but still indicates that the solver is not working as intended. As mentioned before, removing too much dissipation can compromise solver robustness. In this case, the solver still converges but performs poorly.



Figure 6: CTS reduces cost for $\zeta = 10^{-2}$ (NACA 0012 at Mach 0.01)

We can look at linear convergence to see how CTS achieves speedup for $\zeta = M_{\infty}$. The linear system for each ANK step is solved using the generalized minimum residual (GMRES) method [41]. The target linear residual is 0.05. We also set the maximum number of GMRES iterations to 100. If the iteration limit is reached before the linear residual target is met, the solver will continue with the partially converged step. We use the linear residual as an indicator of the linear system stiffness. For $\zeta = M_{\infty}$, CTS reduces the linear residuals such that they are all less than 0.1 (Fig. 8a). This results in faster nonlinear convergence. However, the same is not true for $\zeta = 0$ (Fig. 8b). Despite reducing the



Figure 7: CTS convergence is noisier at lower Mach numbers (NACA 0012, $\zeta = M_{\infty}$, Turkel)

linear residuals, nonlinear convergence is not improved with CTS. As a result, we use $\zeta = M_{\infty}$ for all baseline and CTS cases with freestream Mach numbers of 0.4 or lower in Sec. 4.4 and Sec. 4.5.



Figure 8: CTS improves linear convergence, but this does not guarantee improved nonlinear convergence (NACA 0012 at Mach 0.01)

4.4 NACA 0012 Mach number sweep

We now compare Turkel and VLR preconditioning with the baseline for a range of Mach numbers. Both preconditioners converge faster than the baseline for all Mach numbers we tried between 0.01 to 0.4 (Fig. 9a). The most significant speedup results are for Mach 0.1 and lower, where Turkel is 41%-64% faster than baseline. VLR is consistently faster than Turkel and achieves 49%-71% speedup compared to the baseline in this regime. There is also a considerable speedup for Mach 0.4, but we do not expect to see the same speedup if the baseline solver is tuned for compressible flow.

We also compare CTS and baseline for Mach numbers between 0.6 to 1.4. We use $\zeta = 1$ for

these cases because we are no longer in the low Mach number regime. We also tune the solver for compressible flow. These higher Mach number cases converge faster if we switch to a Newton–Krylov (NK) solver after converging to a relative residual of 10^{-5} with ANK. The NK solver has no time-step matrix and solves the flow and turbulence equations simultaneously in the same Newton system. The lack of the time-step matrix means that CTS has no effect. However, this switch ensures a fair cost comparison in which the baseline solver is not restricted to less effective methods. The performance of Turkel and VLR is similar to the baseline at higher Mach numbers (Fig. 9b). This is the expected behavior for the Turkel preconditioner because we turn off preconditioning for locally supersonic flow. VLR is designed to be an all-speed preconditioner but performs similarly to baseline for higher Mach numbers. This is because the difference in wave speeds and the resulting linear system stiffness is not an issue for Mach numbers around 1. We could try Mach numbers much higher than 1, but this runs into other challenges. We discuss these challenges and propose a solution in Sec. 5.



Figure 9: CTS reduces cost for low Mach numbers (NACA 0012)

4.5 Delta wing

There are also low Mach number cases for which CTS does not result in any speedup. We demonstrate this on a delta wing with an aspect ratio of two at a Mach number of 0.083, Reynolds number of 5.9×10^5 , and angle of attack of 5° (Fig. 10). The geometry and flow conditions are taken from experimental work by Jarrah and Ashley [42]. We use an overset mesh with 1.9 million cells (Fig. 10a). More details on the mesh are given by Seraj and Martins [43]. For this case, stagnation point flow is restricted to a small region near the wing apex (Fig. 10b). Despite the low freestream Mach number, the flow field does not have enough cells with nearly zero local Mach number to make the global linear system stiff. As a result, CTS produces nearly identical convergence behavior as the baseline (Fig. 11). This also explains why the linear residuals are lower than the NACA 0012 baseline case. We show only Turkel results for clarity, but VLR also gives nearly identical results.



(a) Overset mesh with 1.9 million cells

(b) Pressure contours at Mach 0.083, $\alpha = 5^{\circ}$

Figure 10: Delta wing geometry from Jarrah and Ashley [42]



Figure 11: CTS is nearly identical to baseline for the delta wing

5 Dissipation-based continuation for flows with shocks

Flows with shocks pose different challenges for solvers than the low Mach number flows we have discussed. Shocks can move within the domain as the solution converges, causing significant changes in cells' states. To improve robustness, the ANK solver employs a physicality check that limits the update step size such that the density and energy do not change by more than 20% in any cell. A critical outcome of this approach is that the density and energy cannot become negative. Despite this added robustness, ANK can still perform poorly or fail to converge when the step sizes become small.

We demonstrate the effect of step size on convergence using transonic and supersonic flow over a NACA 0012 airfoil. At Mach 3.0, the step sizes are severely limited compared to Mach 0.75, which slows down convergence (Fig. 12). For the NACA 0012 cases in this section, we use the 296 × 128 mesh from Sec. 3. The Reynolds number is 10^7 , and the angle of attack is 3° unless stated otherwise. As with the high Mach number cases in Sec. 4.4, we use the NK solver during the final stages of convergence for all cases in this section. We switch to the NK solver after converging to a relative residual of 10^{-8} with ANK. We could have converged these cases using only ANK. However, switching to NK results in faster convergence and represents typical settings for the baseline solver.



Figure 12: Limited step sizes slow down convergence for supersonic flows (NACA 0012)

One approach to accelerating convergence is to loosen the physicality check to accept larger changes in the states. However, this reduces the robustness of the solver and can often lead to stalled convergence. Another approach is to increase the second-difference artificial dissipation constant, κ_2 . This improves convergence by reducing the spatial gradient in the states but has the disadvantage of reducing solution accuracy. To accelerate convergence while preserving accuracy, we propose a dissipation-based continuation (DBC) method where κ_2 starts high and is continuously reduced as the solution converges.

DBC was previously explored by Hicken et al. [44, 45]. The DBC approach presented here differs from this prior work in that we do not use it as a general globalization strategy. We retain the shock sensor (Eq. 6) for the second-difference dissipation, so DBC is only applied near shocks. In addition, we use DBC in combination with pseudo-transient continuation for globalization. A related but less automated approach is to start the solution with a first-order scheme and switch to a second-order scheme at some point during the solution [14, 15].

5.1 DBC formulation

We write the DBC formulation as

$$\kappa_2^{(n)} = \kappa_2 + f_c^{(n)} \kappa_2^c, \tag{28}$$

where $\kappa_2^{(n)}$ is the dissipation constant at nonlinear iteration n, κ_2 is the desired final dissipation constant, κ_2^c is the initial value of the additional dissipation, and $f_c^{(n)}$ is the continuation parameter at iteration n. We require that the continuation parameter starts at 1 and approaches 0 as the solution converges. We choose to compute the continuation parameter using a generalized sigmoid function:

$$f_c^{(n)} = \frac{1}{1 + e^{-\sigma \left(\log_{10}\left(\eta_{\rm rel}^{(n)}\right) + \lambda\right)}},\tag{29}$$

where σ is the sharpness parameter, λ is the midpoint parameter, and $\eta_{rel}^{(n)}$ is the relative convergence at iteration *n*. The parameters σ and λ determine the shape of the sigmoid. Larger values of σ will result in a steeper drop in dissipation. We call λ the midpoint parameter because the additional dissipation is equal to $0.5\kappa_2^c$ at a relative residual of $10^{-\lambda}$. We use $\sigma = 3$ and $\lambda = 3$, which results in the continuation function shown in Fig. 13. When the solution is converged, the additional dissipation is small enough to recover the solution of the original problem. The shape parameters generally do not need to be tuned for each case. However, the appropriate value for κ_2^c depends on the Mach number. We find that $\kappa_2^c = 0.2M_{\infty}$ works well, and we use this value for all cases presented here.



Figure 13: The additional dissipation decreases smoothly as the flow converges

5.2 NACA 0012

We first compare the baseline solver and DBC for the NACA 0012 case at Mach numbers from 0.75 to 3.0 at a 3° angle of attack (Fig. 14a). DBC is between 3% and 13% faster than baseline for Mach numbers of 2.0 and above. Similarly, DBC is between 4% and 8% faster than baseline at Mach 2.0 and angles of attack from 0° to 6° (Fig. 14b). DBC increases the number of near-unit steps, which improves convergence (Fig. 15). The converged lift and drag values with DBC match the baseline values to 8 significant digits.



Figure 14: DBC offers slight speedup for the NACA 0012 case



Figure 15: DBC improves convergence by increasing step sizes (NACA 0012 at Mach 2.5)

5.3 Supersonic transport configuration

The advantages of DBC are more pronounced when applied to a more complex case. We now consider a supersonic transport (SST) configuration with a cranked arrow wing, T-tail, and canard (Fig. 16). We consider flow at Mach numbers of 1.8 and 0.95, representing realistic supersonic and transonic cruise conditions for the next generation of SSTs [1]. Both cases have a 4° angle of attack. DBC converges both cases, whereas the baseline solver does not converge either (Fig. 17). We also compare DBC to the baseline solver with κ_2 increased from the default value of 0.25 to the initial DBC value of $0.25 + 0.2M_{\infty}$. For Mach 1.8, increasing κ_2 converges faster than DBC. However, this solution is also less accurate because of the increased dissipation near shocks. For Mach 0.95, the solver still stalls with the higher κ_2 , albeit after about three more orders of convergence.

In some cases, robustness is more valuable than faster convergence. The SST mesh used here is the same as that used by Seraj and Martins [46] for aerodynamic shape optimization studies. Reliably converging the flow at different design points is essential for a well-behaved optimization [47]. DBC offers more robustness than the baseline solver for complex geometries at high Mach numbers, making it more suitable for many-query scenarios such as parameter sweeps or design optimization.



Figure 16: SST overset mesh with pressure contours at Mach 1.8

6 Conclusions

Using compressible flow solvers for low Mach number flows typically results in poor accuracy and speed. In addition, high Mach number flows can be challenging to converge reliably without increasing the amount of dissipation in the spatial discretization. We present three contributions to address these challenges.

First, we propose a simple modification to the Jameson–Schmidt–Turkel scheme to make the artificial dissipation appropriate for low Mach number flows. We show that scaling down the acoustic contribution of the spectral radius improves accuracy and reduces sensitivity to mesh size. We recommend using a scaling factor equal to the freestream Mach number. Second, we demonstrate the effectiveness of a characteristic time-stepping method for approximate Newton–Krylov solvers.



Figure 17: DBC converges the SST cases without sacrificing accuracy

This approach reduces the stiffness of the linear system at low Mach numbers by modifying the timestep matrix. We show that the artificial dissipation must be adequately scaled to achieve speedup with characteristic time-stepping. In addition, cases with small regions of stagnation point flow do not benefit from characteristic time-stepping. Third, we present a dissipation-based continuation method for flows with shocks. The continuation approach introduces additional dissipation near shocks during the initial stages of the solution and smoothly reduces the dissipation as the solution converges. This approach is faster and more robust than the baseline solver, particularly for more complex geometries. These contributions enable more accurate and efficient flow solvers across low and high Mach numbers, making CFD more dependable for applications involving a wide range of flow conditions, such as aircraft design.

Acknowledgments

This work was supported by NASA's Commercial Supersonic Technology project (grant number 80NSSC19K1661). S. Seraj was also supported by the Natural Sciences and Engineering Research Council of Canada (funding reference number PGSD3-545678-2020). Computational resources were provided by the NASA High-End Computing Program through the NASA Advanced Supercomputing Division at Ames Research Center.

References

[1] Morgenstern, J., Buonanno, M., Yao, J., Murugappan, M., Paliath, U., Cheung, L., Malcevic, I., Ramakrishnan, K., Pastouchenko, N., Wood, T., Martens, S., Viars, P., Tersmette, T., Lee, J., Simmons, R., Plybon, D., Alonso, J., Palacios, F., Lukaczyk, T., and Carrier, G., "Advanced Concept Studies for Supersonic Commercial Transports Entering Service in the 2018-2020 Period Phase 2," Tech. Rep. NASA/CR—2015-218719, NASA, Glenn Research Center, Cleveland, OH, 2015.

- [2] Bons, N. P., and Martins, J. R. R. A., "Aerostructural Design Exploration of a Wing in Transonic Flow," Aerospace, Vol. 7, No. 8, 2020, p. 118. doi:10.3390/aerospace7080118.
- [3] Gleize, V., and Costes, M., "Low-Mach-Number Preconditioning Applied to Turbulent Helicopter Fuselage Flowfield Computation," AIAA Journal, Vol. 41, No. 4, 2003, pp. 653–662. doi:10.2514/2.1995.
- [4] Colin, Y., Deniau, H., and Boussuge, J.-F., "A robust low speed preconditioning formulation for viscous flow computations," *Computers & Fluids*, Vol. 47, No. 1, 2011, pp. 1–15. doi:10.1016/j.compfluid.2011.01.015.
- [5] Turkel, E., "Preconditioned methods for solving the incompressible and low speed compressible equations," *Journal of Computational Physics*, Vol. 72, No. 2, 1987, pp. 277–298. doi:10.1016/0021-9991(87)90084-2.
- [6] Turkel, E., "Preconditioning techniques in computational fluid dynamics," Annual Review of Fluid Mechanics, Vol. 31, No. 1, 1999, pp. 385–416. doi:10.1146/annurev.fluid.31.1.385.
- [7] Weiss, J. M., and Smith, W. A., "Preconditioning applied to variable and constant density flows," *AIAA Journal*, Vol. 33, No. 11, 1995, pp. 2050–2057. doi:10.2514/3.12946.
- [8] Nemec, M., and Zingg, D. W., "Aerodynamic Computations Using the Convective-Upstream Split-Pressure Scheme with Local Preconditioning," *AIAA Journal*, Vol. 38, No. 3, 2000, pp. 402–410. doi:10.2514/2.998.
- [9] van Leer, B., Lee, W.-T., and Roe, P., "Characteristic time-stepping or local preconditioning of the Euler equations," *10th Computational Fluid Dynamics Conference*, 1991. doi:10.2514/6.1991-1552.
- [10] Lee, D., "The Design of Local Navier-Stokes Preconditioning for Compressible Flow," Journal of Computational Physics, Vol. 144, No. 2, 1998, pp. 460–483. doi:10.1006/jcph.1998.5994.
- [11] Knoll, D. A., McHugh, P. R., and Keyes, D. E., "Newton–Krylov methods for low-Mach-number compressible combustion," AIAA Journal, Vol. 34, No. 5, 1996, pp. 961–967. doi:10.2514/3.13174.
- [12] Weston, B., Nourgaliev, R., Delplanque, J.-P., and Barker, A. T., "Preconditioning a Newton– Krylov solver for all-speed melt pool flow physics," *Journal of Computational Physics*, Vol. 397, 2019, p. 108847. doi:10.1016/j.jcp.2019.07.045.
- [13] Mary, I., Sagaut, P., and Deville, M., "An algorithm for low Mach number unsteady flows," *Computers & Fluids*, Vol. 29, No. 2, 2000, pp. 119–147. doi:10.1016/S0045-7930(99)00007-9.
- [14] Kaushik, D. K., Keyes, D. E., and Smith, B. F., "Newton-Krylov-Schwarz Methods for Aerodynamics Problems: Compressible and Incompressible Flows on Unstructured Grids," 11th International Conference on Domain Decomposition Methods, 1998.
- [15] Olawsky, F., Infed, F., and Auweter-Kurtz, M., "Preconditioned Newton Method for Computing Supersonic and Hypersonic Nonequilibrium Flows," *Journal of Spacecraft and Rockets*, Vol. 41, No. 6, 2004, pp. 907–914. doi:10.2514/1.4010.

- [16] Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Gnoffo, P. A., Hammond, D. P., Jacobson, K. E., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Nielsen, E. J., Park, M. A., Rumsey, C. L., Thomas, J. L., Thompson, K. B., Walden, A. C., Wang, L., and Wood, W. A., "FUN3D Manual: 13.7," Tech. Rep. NASA/TM–20205010139, NASA Langley Research Center, Hampton, Virginia, November 2020.
- [17] Berger, M., Aftosmis, M. J., and Murman, S. M., "Analysis of Slope Limiters on Irregular Grids," 43rd AIAA Aerospace Sciences Meeting and Exhibit, 2005. doi:10.2514/6.2005-490.
- [18] Öhrman, J., "Evaluation of a CFD method for estimating aerodynamic loads on external stores on JAS 39 Gripen," Ph.D. thesis, Umeå University, 2011.
- [19] Aprovitola, A., Dyblenko, O., Pezzella, G., and Viviani, A., "Aerodynamic Analysis of a Supersonic Transport Aircraft at Low and High Speed Flow Conditions," *Aerospace*, Vol. 9, No. 8, 2022, p. 411. doi:10.3390/aerospace9080411.
- [20] Mader, C. A., Kenway, G. K. W., Yildirim, A., and Martins, J. R. R. A., "ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization," *Journal of Aerospace Information Systems*, Vol. 17, No. 9, 2020, pp. 508–527. doi:10.2514/1.1010796.
- [21] Spalart, P., and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," La Recherche Aerospatiale, Vol. 1, 1994, pp. 5–21.
- [22] Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solution of the Euler Equations by Finite Volume Methods Using Runge–Kutta Time Stepping Schemes," 14th Fluid and Plasma Dynamics Conference, 1981. doi:10.2514/6.1981-1259.
- [23] Jameson, A., "Origins and Further Development of the Jameson–Schmidt–Turkel Scheme," AIAA Journal, Vol. 55, No. 5, 2017, pp. 1487–1510. doi:10.2514/1.J055493.
- [24] Martinelli, L., "Calculations of Viscous Flows with a Multigrid Method," Ph.D. thesis, Princeton University, 1987.
- [25] Yildirim, A., Kenway, G. K. W., Mader, C. A., and Martins, J. R. R. A., "A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations," *Journal of Computational Physics*, Vol. 397, 2019, p. 108741. doi:10.1016/j.jcp.2019.06.018.
- [26] Wang, Z., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, R., Hillewaert, K., Huynh, H., Kroll, N., May, G., Persson, P.-O., van Leer, B., and Visbal, M., "High-order CFD methods: current status and perspective," *International Journal for Numerical Methods in Fluids*, Vol. 72, No. 8, 2013, pp. 811–845. doi:10.1002/fld.3767.
- [27] Swanson, R. C., and Turkel, E., "On central-difference and upwind schemes," *Journal of Compu*tational Physics, Vol. 101, No. 2, 1992, pp. 292–306. doi:10.1016/0021-9991(92)90007-L.
- [28] Jameson, A., "Analysis and Design of Numerical Schemes for Gas Dynamics 1—Artificial Diffusion, Upwind Biasing, Limiters and Their Effect on Accuracy and Multigrid Convergence," *International Journal of Computational Fluid Dynamics*, Vol. 4, No. 3–4, 1995, pp. 171–218. doi:10.1080/10618569508904524.

- [29] Gustafsson, B., "Unsymmetric hyperbolic systems and the Euler equations at low Mach numbers," *Journal of Scientific Computing*, Vol. 2, No. 2, 1987, pp. 123–136. doi:10.1007/BF01061482.
- [30] Rieper, F., "A low-Mach number fix for Roe's approximate Riemann solver," Journal of Computational Physics, Vol. 230, No. 13, 2011, pp. 5263–5287. doi:10.1016/j.jcp.2011.03.025.
- [31] Liou, M.-S., "A sequel to AUSM, Part II: AUSM+-up for all speeds," *Journal of Computational Physics*, Vol. 214, No. 1, 2006, pp. 137–170. doi:10.1016/j.jcp.2005.09.020.
- [32] Berberich, J. P., and Klingenberg, C., "Entropy Stable Numerical Fluxes for Compressible Euler Equations Which Are Suitable for All Mach Numbers," *Recent Advances in Numerical Methods for Hyperbolic PDE Systems*, 2021. doi:10.1007/978-3-030-72850-2_8.
- [33] Vassberg, J. C., Tinoco, E. N., Mani, M., Brodersen, O. P., Eisfeld, B., Wahls, R. A., Morrison, J. H., Zickuhr, T., Laflin, K. R., and Mavriplis, D. J., "Abridged Summary of the Third AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 45, No. 3, 2008, pp. 781–798. doi:10.2514/1.30572.
- [34] Drela, M., "XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils," Low Reynolds Number Aerodynamics, edited by T. J. Mueller, Springer Berlin Heidelberg, Berlin, Heidelberg, 1989, pp. 1–12. doi:10.1007/978-3-642-84010-4_1.
- [35] Guillard, H., and Viozat, C., "On the behaviour of upwind schemes in the low Mach number limit," *Computers & Fluids*, Vol. 28, No. 1, 1999, pp. 63–86. doi:10.1016/j.camwa.2018.02.028.
- [36] Chen, S., Yan, C., and Xiang, X., "Effective low-Mach number improvement for upwind schemes," *Computers & Mathematics with Applications*, Vol. 75, No. 10, 2018, pp. 3737–3755. doi:10.1016/j.camwa.2018.02.028.
- [37] Zingg, D. W., "Viscous airfoil computations using Richardson extrapolation," 10th Computational Fluid Dynamics Conference, 1991. doi:10.2514/6.1991-1559.
- [38] Roache, P. J., "Perspective: A Method for Uniform Reporting of Grid Refinement Studies," Journal of Fluids Engineering, Vol. 116, No. 3, 1994, pp. 405–413. doi:10.1115/1.2910291.
- [39] Turkel, E., and Vatsa, V. N., "Local preconditioners for steady and unsteady flow applications," ESAIM: Mathematical Modelling and Numerical Analysis, Vol. 39, No. 3, 2005, pp. 515–535. doi:10.1051/m2an:2005021.
- [40] Moragues Ginard, M., Vázquez, M., and Houzeaux, G., "Local preconditioning and variational multiscale stabilization for Euler compressible steady flow," *Computer Methods in Applied Mechanics and Engineering*, Vol. 305, 2016, pp. 468–500. doi:10.1016/j.cma.2016.02.027.
- [41] Saad, Y., and Schultz, M. H., "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. doi:10.1137/0907058.
- [42] Jarrah, M. A., and Ashley, H., "Impact of flow unsteadiness on maneuvers and loads of agile aircraft," 30th Structures, Structural Dynamics and Materials Conference, 1989. doi:10.2514/6.1989-1282.

- [43] Seraj, S., and Martins, J. R. R. A., "Predicting the High-Angle-of-Attack Characteristics of a Delta Wing at Low Speed," *Journal of Aircraft*, Vol. 59, No. 4, 2022, pp. 1071–1081. doi:10.2514/1.C036618.
- [44] Hicken, J., and Zingg, D., "Globalization Strategies for Inexact-Newton Solvers," 19th AIAA Computational Fluid Dynamics, American Institute of Aeronautics and Astronautics, 2009. doi:10.2514/6.2009-4139.
- [45] Hicken, J., Buckley, H., Osusky, M., and Zingg, D., "Dissipation-based continuation: a globalization for inexact-Newton solvers," 20th AIAA Computational Fluid Dynamics Conference, American Institute of Aeronautics and Astronautics, 2011. doi:10.2514/6.2011-3237.
- [46] Seraj, S., and Martins, J. R. R. A., "Aerodynamic Shape Optimization of a Supersonic Transport Considering Low-Speed Stability," AIAA SciTech Forum, 2022. doi:10.2514/6.2022-2177.
- [47] Martins, J. R. R. A., "Aerodynamic Design Optimization: Challenges and Perspectives," Computers & Fluids, Vol. 239, 2022, p. 105391. doi:10.1016/j.compfluid.2022.105391.