

Matching Multistage Schemes to Viscous Flow

by Bil Kleb, Bram van Leer, and Bill Wood

Multistage, explicit time stepping can be tailored to accelerate convergence for scalar advection-diffusion problems by using optimized multistage coefficients that vary with the local cell Reynolds number. And, when combined with local preconditioning, variable-coefficient multistage schemes for computational fluid dynamics codes can also provide an order of magnitude faster convergence, relative to standard, fixed-coefficient schemes, for the Navier-Stokes system of equations.

1 Introduction

Slow convergence to steady state of explicit time-stepping CFD codes for viscous flow problems motivates this work. The development of local preconditioning methods¹ and the flexibility of multistage time-marching schemes provide the tools to significantly improve explicit convergence rates.²

An explicit time marching scheme is typically limited by the minimum of an inviscid time step and a viscous time step. In terms of Von Neumann stability analysis these two time steps correspond to the extents of the Fourier footprint³ along the imaginary axis and negative real axis, respectively, as depicted in figure 1.

In the beginning of the “Euler era” (early 1980s), multistage time-marching schemes were optimized for the maximum *imaginary* extent of the stability domain, given the number of stages of the scheme.⁴ This allows inviscid problems to be solved efficiently on a single grid. Very soon, multigrid relaxation made its entrance, and the emphasis in multistage design shifted to choosing the coefficients such that the scheme provides the maximum high-frequency damping for a given number of stages—typically at half the maximum allowable time step.

Originally, this optimization technique was based on a scalar analysis. It thus lost potency when applied to a system of equations in which the high frequencies of the Fourier footprint are no longer concentrated in a single region. The introduction of local preconditioning for the Euler equations, which tends to equalize the convective time scales and concentrate the Fourier footprint, finally made it possible to optimize high-frequency damping for all types of waves admitted by the inviscid system.⁵

For a typical viscous flow problem, however, a majority of the computational cells’ time steps are limited by the

¹ Eli Turkel, “Preconditioning Techniques in Computational Fluid Dynamics”, *Annual Review in Fluid Mechanics*, 31, 1999, pp. 385–416.

² And, as an added bonus, certain local preconditioners also restore solution accuracy for low-Mach-number flows.

³ The locus of the Fourier transform in the complex plane.

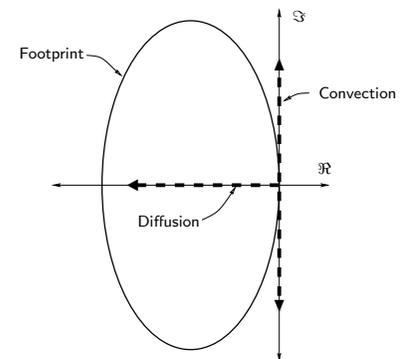


Figure 1: Spatial operator Fourier components.

⁴ P. Sonneveld and Bram van Leer, “A Minimax Problem along the Imaginary Axis”, *Nieuw Archief voor Wiskunde*, 4, 1985, (see also Delft Report of the Department of Mathematics and Informatics 84–25).

⁵ See, for example, Turkel’s review cited above.

viscous criterion due to the limited extent of the temporal stability boundary along the negative real axis. By making the multistage coefficients a function of the local cell Reynolds number,⁶ the stability boundary can be reshaped to alleviate this restriction. Local preconditioning remains necessary for equalizing both convective and dissipative time scales admitted by the system. However, the analytical preconditioning analysis for the Navier-Stokes equations is much more complicated than for the Euler equations and has only been solved for the one-dimensional case.⁷ As a result, the current work employs a simple viscous Jacobi block to extend the Euler preconditioner to viscous problems.

In the work of Lynn,⁸ multistage methods are optimized based on the precise Fourier footprint of the spatial operator. The optimization must be performed separately for each choice of cell Reynolds number. The multistage time steps can be tabulated as functions of Re_h , or, more practically, given as functions that closely fit the data. However, neither Lynn nor, later, Lee⁹ actually carry out Navier-Stokes calculations with Re_h -dependent multistage coefficients.

In the present work the optimization procedure is reversed, and thereby greatly simplified. After covering the multistage scheme design, results will be shown for a *scalar* model problem and then the Navier-Stokes *system* of equations.

2 Multistage Scheme Design

Designing multistage time-marching schemes is most easily accomplished by working in the complex-variable space of the Fourier transform via Von Neumann stability analysis.

To achieve stability, the Fourier footprint of the spatial discretization must be contained by the stability boundary of the time-marching scheme. A secondary goal is to provide some level of high-frequency damping. When using single-grid relaxation, the tradition is to go for the largest time step possible. For multigrid relaxation, however, it is necessary to balance the desire for large time steps with good overall high-frequency damping. A modified version of the latter strategy actually turns out to also be more efficient for single-grid relaxation, as was shown by Tai.¹⁰ Although, maximizing high-frequency damping without setting a target level for the damping is not a good idea for a viscous equation solver because this reduces the time step to unacceptably low values.

Once a discretization of the spatial terms is chosen, the shape of its Fourier footprint is fixed; and so, the next subsection rigorously defines the spatial discretization chosen for this study. Afterward, the freedoms allowed in the temporal operator will be explored. Finally, the procedure for marrying

⁶ Where local cell Reynolds number, Re_h , is the Reynolds number based on a cell length scale (defined later), and the cell's velocity and viscosity.

⁷ Christopher Depcik and Bram van Leer, *In Search of an Optimal Local Navier-Stokes Preconditioner*, AIAA Paper 2003-3703, 2003.

⁸ John Francis Lynn, *Multigrid Solution of the Euler Equations with Local Preconditioning*, Ph.D. thesis, University of Michigan, 1995.

⁹ Dohyung Lee, et al., *A local Navier-Stokes Preconditioner for all Mach and Cell Reynolds Numbers*, AIAA Paper 97-2024, 1997.

¹⁰ C. H. Tai, *Acceleration Techniques for Explicit Euler Codes*, Ph.D. thesis, University of Michigan, 1990.

the temporal operator with the spatial operator is given.

Spatial Scheme

Consider the one-dimensional advection-diffusion equation,

$$S_t + a S_x = \frac{1}{Re} S_{xx}, \quad (1)$$

where S is some conserved, scalar quantity; a is the advection velocity; and Re is the Reynolds number. Discretizing this via central differencing for the diffusive term and first-order upwind for the advective term yields the footprint,

$$\mathcal{F} = - \left(\frac{a\Delta t}{\Delta x} + 2 \frac{\Delta t}{Re\Delta x^2} \right) (1 - \cos \beta) - i \frac{a\Delta t}{\Delta x} \sin \beta, \quad (2)$$

where $i = \sqrt{-1}$, β is the Fourier frequency, and Δ represents a discrete quantity, e.g., $\Delta x = x_j - x_{j-1}$.

Typically, parameters are defined according to the advective and diffusive terms, e.g.,

$$\nu = \frac{a\Delta t}{\Delta x} \quad \text{and} \quad \sigma = \frac{\Delta t}{Re\Delta x^2},$$

where ν is the Courant number and σ is the Von Neumann number. This parameter definition gives an elliptical footprint¹¹ in Fourier space as shown in figure 2.

Fixing the location of the negative real extent of this footprint is accomplished by examining the high frequency limit ($\beta = \pi$) of equation 2,

$$-R_S = -2 \left(\frac{a\Delta t}{\Delta x} + 2 \frac{\Delta t}{Re\Delta x^2} \right). \quad (3)$$

Solving for the time step and writing in terms of cell Reynolds number, yields,

$$\Delta t = \frac{R_S \Delta x}{2} \frac{Re_{\Delta x}}{a Re_{\Delta x} + 2}, \quad (4)$$

where $Re_{\Delta x} = Re a \Delta x$.

After inserting the definition of Δt from equation 4 into equation 2, the overall shape of the footprint is now seen to be dependent on only two parameters: R_S and the cell Reynolds number,

$$\mathcal{F} = - \frac{R_S}{2} \left[(1 - \cos \beta) + i \frac{Re_{\Delta x}}{Re_{\Delta x} + 2} \sin \beta \right]. \quad (5)$$

R_S sets the negative real extent and the cell Reynolds number dictates the semi-minor axis of the ellipse.

¹¹ A higher-order κ -scheme yields an ellipse only for $\kappa = 1$. In general, the higher-order scheme is egg-shaped with the widest end to the left.

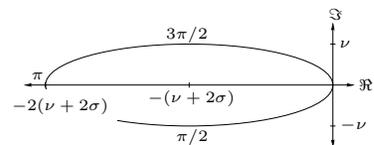


Figure 2: Fourier footprint for the one-dimensional Advection-Diffusion equation using first-order upwind spatial discretization for the advective term and central differencing for the diffusion term ($\beta \in [0, 2\pi]$).

WHILE the one-dimensional case is relatively straightforward in terms of parameter choices, additional spatial dimensions create ambiguities that are not resolved in the CFD community.

In two-dimensions, the governing equation is now,

$$S_t + a S_x + b S_y = \frac{1}{Re} (S_{xx} + S_{yy}), \quad (6)$$

which yields a footprint of the form,

$$\begin{aligned} \mathcal{F} = & - \left(\frac{a\Delta t}{\Delta x} + 2\frac{\Delta t}{Re\Delta x^2} \right) (1 - \cos \beta_x) - i\frac{a\Delta t}{\Delta x} \sin \beta_x \\ & - \left(\frac{b\Delta t}{\Delta y} + 2\frac{\Delta t}{Re\Delta y^2} \right) (1 - \cos \beta_y) - i\frac{b\Delta t}{\Delta y} \sin \beta_y, \end{aligned} \quad (7)$$

for first-order upwind and central differenced viscous terms. Note there are now two Fourier frequency components, β_x and β_y .

From this point, there are many avenues to follow. One of the most prevalent is to define component-wise Courant numbers and Von Neumann numbers, e.g.,

$$\begin{aligned} \nu_x &= \frac{a\Delta t}{\Delta x}, & \sigma_x &= \frac{\Delta t}{Re\Delta x^2}, \\ \nu_y &= \frac{b\Delta t}{\Delta y}, & \text{and } \sigma_y &= \frac{\Delta t}{Re\Delta y^2}. \end{aligned}$$

Instead, the parameters chosen are more closely linked to the flow physics. The flow quantities are written in terms of polar coordinates (q, ϕ) and the geometric quantities are written in terms of Δx , the cell aspect ratio, $q = \sqrt{a^2 + b^2}$, $\phi = \arctan(b/a)$, and $\mathcal{R} = \Delta x/\Delta y$. Figure 3 depicts these orientations.

Applying these definitions to equation 7 gives,

$$\begin{aligned} \mathcal{F} = & - \left(\frac{q \cos \phi \Delta t}{\Delta x} + 2\frac{\Delta t}{Re\Delta x^2} \right) (1 - \cos \beta_x) \\ & - \left(\mathcal{R} \frac{q \sin \phi \Delta t}{\Delta x} + 2\mathcal{R}^2 \frac{\Delta t}{Re\Delta x^2} \right) (1 - \cos \beta_y) \quad (8) \\ & - i\frac{q \cos \phi \Delta t}{\Delta x} \sin \beta_x - i\mathcal{R} \frac{q \sin \phi \Delta t}{\Delta x} \sin \beta_y. \end{aligned}$$

And, as for the one-dimensional case, examining the high frequency limit ($\beta_x = \beta_y = \pi$) of the Fourier footprint reveals,

$$-R_S = -2\Delta t \left(q \frac{\cos \phi + \mathcal{R} \sin \phi}{\Delta x} + \frac{2}{Re} \frac{1 + \mathcal{R}^2}{\Delta x^2} \right). \quad (9)$$

In this form, two length scales are discernible. These will be labeled l and h as follows:

$$l = \frac{\Delta x}{\cos \phi + \mathcal{R} \sin \phi}, \quad (10)$$

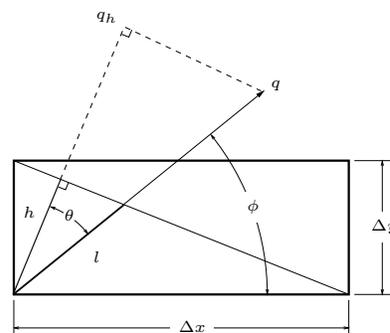


Figure 3: Time-step geometry.

which is the length from a cell corner to the diagonal, along the advection velocity direction as previously reported by Tai¹² and,

$$h = \frac{\Delta x}{\sqrt{1 + \mathcal{R}^2}} = \frac{\Delta x \Delta y}{\sqrt{\Delta x^2 + \Delta y^2}}, \quad (11)$$

which is the ratio of cell area to the length of the diagonal.

Furthermore, h and l can be related by $l = h / \cos \theta$ where θ is the angle between h and the advection velocity. Using the new length scale, h , and the grid-to-flow angle, θ , a time step can be found which fixes the high frequency portion of the footprint at $-R_S$,

$$\Delta t = \frac{R_S}{2} \frac{h}{q_h} \frac{Re_h}{Re_h + 2}. \quad (12)$$

where $Re_h = q_h h Re$ and $q_h = q \cos \theta$. Compare this to one-dimensional results given by equation 4 on page 3. In the two-dimensional case Δx is replaced by the new length scale, h , which gives a form similar to the harmonic sum of the two length scales,

$$h = 1 / \sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}}. \quad (13)$$

This weighted sum favors the smaller of the two components which is similar to the common practice of using the minimum of Δx and Δy or using their harmonic sum, $1 / (\Delta x^{-1} + \Delta y^{-1})$. A departure from current practice and the one-dimensional results is that the flow speed component in the h direction, q_h , governs the advective portion and not the full flow velocity, q .

Substituting the definition of Δt from equation 12 into the Fourier footprint given by equation 8 and using the cell Reynolds number definition, yields,

$$\begin{aligned} \mathcal{F} = & -\frac{R_S}{2} \frac{1}{\sqrt{1 + \mathcal{R}^2}} \frac{Re_h}{Re_h + 2} \\ & \left[\left(\cos \phi + \frac{2}{Re_h \sqrt{1 + \mathcal{R}^2}} \right) (1 - \cos \beta_x) \right. \\ & + \mathcal{R} \left(\sin \phi + \frac{2\mathcal{R}}{Re_h \sqrt{1 + \mathcal{R}^2}} \right) (1 - \cos \beta_y) \\ & \left. + i(\cos \phi \sin \beta_x + \sin \phi \sin \beta_y) \right], \quad (14) \end{aligned}$$

which, in addition to depending on the frequencies, the negative real extent, and the cell Reynolds number, as in the one-dimensional case, is also dependent on the flow angle, ϕ , and cell aspect ratio, \mathcal{R} . This form of the Fourier footprint is not as straightforward as the one-dimensional result, but

¹² C. H. Tai, *Acceleration Techniques for Explicit Euler Codes*, Ph.D. thesis, University of Michigan, 1990.

by considering the case of unity aspect ratio, a bisecting flow angle ($\phi = 45^\circ$), and concurrent frequencies ($\beta_x = \beta_y = \beta$), it collapses to the one-dimensional result given by equation 5 on page 3.

Now that the spatial operator footprint is conveniently scaled, a suitable temporal stability boundary can be found.

Temporal Scheme

Multistage schemes consist of a series of predictor-corrector steps. The amplification factor of a four-stage scheme, for example, has the form,

$$P_4(z) = 1 + \alpha_4 z + \alpha_4 \alpha_3 z^2 + \alpha_4 \alpha_3 \alpha_2 z^3 + \alpha_4 \alpha_3 \alpha_2 \alpha_1 z^4, \quad (15)$$

with $\alpha_4 = 1$. The coefficients, α_i , can be selected with considerable freedom, especially if high-order time accuracy properties are not desired.

In fact, this entire study was inspired by the large, negative *real* extent of the stability boundary made possible by using multistage coefficients derived from Chebyshev polynomials.¹³ Chebyshev polynomials are given by a recursion relation,

$$T_{n+1}(x) = 2x T_n(x) - T_{n-1}(x), \quad n > 1, \quad (16)$$

with $T_0(x) = 1$ and $T_1(x) = x$ where $x \in [0, 1]$. So, for example, the 4th-order Chebyshev polynomial is,

$$T_4(x) = 8x^4 - 8x^2 + 1. \quad (17)$$

By equating the Chebyshev polynomial with the form of the multistage predictor given by equation 15, one can find the resulting multistage coefficients. Figure 4 shows the stability boundaries for Chebyshev polynomials corresponding to 2-, 3-, and 4-stage multistage schemes. Note that the negative real extent is given by $-2n^2$ where n is the number of stages.

Sonneveld, Van Leer, Turkel, and others had experimented with these polynomials in the 1970's and 1980's prompted by the minimax problem posed by Van der Houwen.¹⁴ Their focus, however, was to achieve the largest possible *imaginary* extent. It was found that the imaginary extent varies linearly as the order of the polynomial, which is equal to the number of stages in a multistage scheme. When applying Chebyshev polynomials along the negative real axis, however, their extent varies as the order of the polynomial squared. Thus, with a moderate number of stages, a significant extent of the real axis can be encompassed. Note, however, that the pure Chebyshev polynomial boundaries depicted in figure 4 are of little practical use since they have neutral stability points distributed along the real axis.

¹³ W. Riha, "Optimal Stability Polynomials", *Computing*, 9, 1972, pp. 37-43.

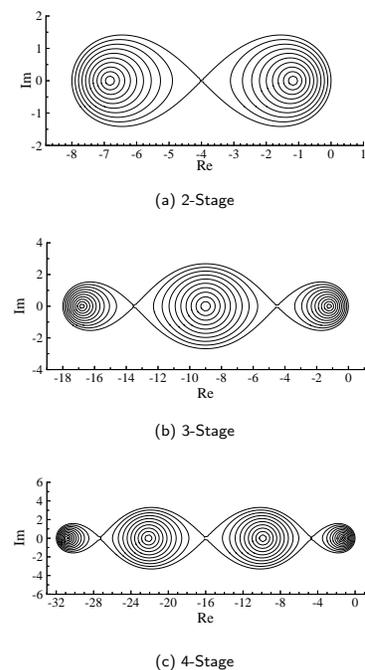


Figure 4: Chebyshev polynomial stability boundaries given by equation 16.

¹⁴ P. J. van der Houwen, "The Minimax Problem", *Nieuw Archief voor Wiskunde*, 26(3), 1978, p. 339.

Based on the work of Lorber et al.,¹⁵ a three-parameter family of Chebyshev polynomial transformations was used that provides a more practical stability region. This transformation, attributed to Manteuffel,¹⁶ consists of a shift and scaling of the Chebyshev polynomials,

$$P_n(z) = \frac{T_n\left(\frac{d-z}{\varepsilon}\right)}{T_n\left(\frac{d}{\varepsilon}\right)}, \quad (18)$$

such that the resulting polynomial guarantees an elliptical stability region with a semi-major axis of length d and a semi-minor axis of length ε . However, the actual stability region is slightly larger and typically much more irregular.

The coefficients of $P_n(z)$ are chosen such that the stability boundary given by $|P_n(z)| = 1$ remains simply connected and the polynomial is scaled properly, i.e., $P_n(0) = 1$, and $P'_n(0) = 1$, so that $P_n(2) \approx 1 + 2$ near the origin, which gives first-order temporal accuracy.

With only two-stages, the extent along the negative real axis is not very impressive, while odd-number-stage schemes do not provide sufficient coverage of imaginary axis necessary for high-order spatial operators. Thus, the smallest number of stages for practical application is a four-stage scheme. This scheme is given by

$$P_4(z) = 1 - \frac{32d^3 - 16\varepsilon^2d}{\mathcal{D}_4}z + \frac{48d^2 - 8\varepsilon^2}{\mathcal{D}_4}z^2 - \frac{32d}{\mathcal{D}_4}z^3 + \frac{8}{\mathcal{D}_4}z^4, \quad (19)$$

where $\mathcal{D}_4 = 8d^4 - 8d^2\varepsilon^2 + \varepsilon^4$ and the $P'(0) = 1$ constraint gives,

$$\varepsilon = \sqrt{4d(2+d) - \sqrt{16d^2(2+d)^2 - 8d^3(d+4)}}.$$

Figure 5 shows the transition of the stability boundary and damping properties from the inviscid limit to the viscous limit for this four-stage scheme. Subfigures 5a and 5c show an over an order of magnitude difference in negative real extent through the use of Manteuffel's Chebyshev family.

And finally, the gap between Manteuffel transformation variables and those employed for multistage schemes is bridged by equating the multistage coefficients, α_k , of equation 15 with those in equation 19.

The next step is mating these spatial and temporal results to form a robust and efficient method for viscous flow problems.

¹⁵ Alfred A. Lorber, et al., "ODE Recursions and Iterative Solvers for Linear Equations", *SIAM Journal of Scientific Computing*, 17(1), 1996, pp. 65-77, which was found by Eli Turkel during his stay at NASA Langley Research Center the summer of 1998.

¹⁶ Thomas A. Manteuffel, "The Tchebychev Iteration for Nonsymmetric Linear Systems", *Numerische Mathematik*, 28, 1977, pp. 307-327.

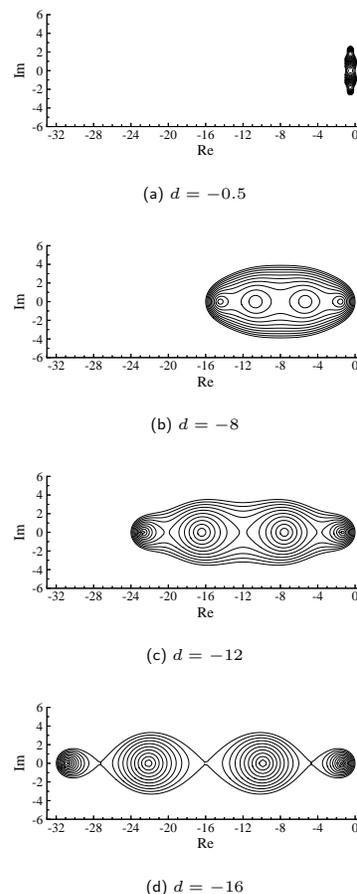


Figure 5: Four-stage Manteuffel Chebyshev blending as a function of d , contours of damping from 0.1 to 1.0 by 0.1 increments.

Fitting

In summary, the procedure begins with selecting a stability boundary from the family given by the Manteuffel transformation. Next, a high-frequency damping level is located on the negative real axis; and then the cell Reynolds number is increased until some part of the spatial operator's Fourier footprint reaches a specified damping contour of the temporal operator. Figure 6 depicts the parameters, the stability boundary, and the spatial operator Fourier footprint.

Specifically, the fitting algorithm is as follows:

1. Pick successive values for d , the blending parameter of the Manteuffel mapping given by equation 18. This sets the shape of the temporal stability boundary.
2. Compute $-R_T$, the negative real extent of the temporal operator. ($R_T = 2d$.)
3. Given the negative real extent of the stability domain, $-R_T$, as a starting point, find the negative real extent of the spatial operator, $-R_S$, by moving toward the origin along the real axis until $|P_n(z)|$ matches a specified high-frequency damping rate, σ_π .¹⁷
4. Using the viscous limit ($Re_h \rightarrow 0$) as an initial guess, increase the cell Reynolds number of the spatial operator until reaching the prescribed damping criteria.

For single-grid relaxation, the damping criterion is chosen to achieve a specified rate for the highest frequency combination, (π, π) and less stringent damping properties over other frequency ranges. For example, figure 7 shows the resulting family of Fourier design plots for the case of 0.5 high-frequency damping, damping of 0.98 over the range $[\pi/3, \pi]$, and stable elsewhere. (Other designs are discussed in the dissertation.¹⁸) This design places a large emphasis on maximizing the time step, but still affords nice high-frequency damping properties. Notice the large negative extent of the stability domain along the real axis, made possible by the polynomial. The domain maximally extends to $-2n^2$ for an n -stage scheme; however, at this bound, the scheme is only stable for the limit of $Re_h = 0$.

Figure 8 and table 1 on page 10 show the variation of the multistage coefficients and the negative-real extent of the spatial operator as a function of cell Reynolds number for the single-grid relaxation design shown in figure 7. Attempts were made to curve fit these results to reduce the computational overhead of the lookup needed for each computational cell. However, the stability boundaries were found to be sensitive to errors in the coefficients and the slight kink in the values around $Re_h = 4$ (0.6 on \log_{10} scale) prevented the use of simple functions like the hyperbolic tangent. As a result, the

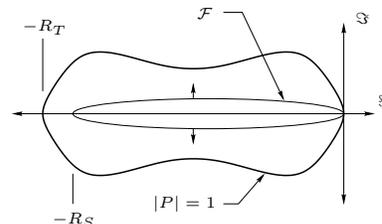
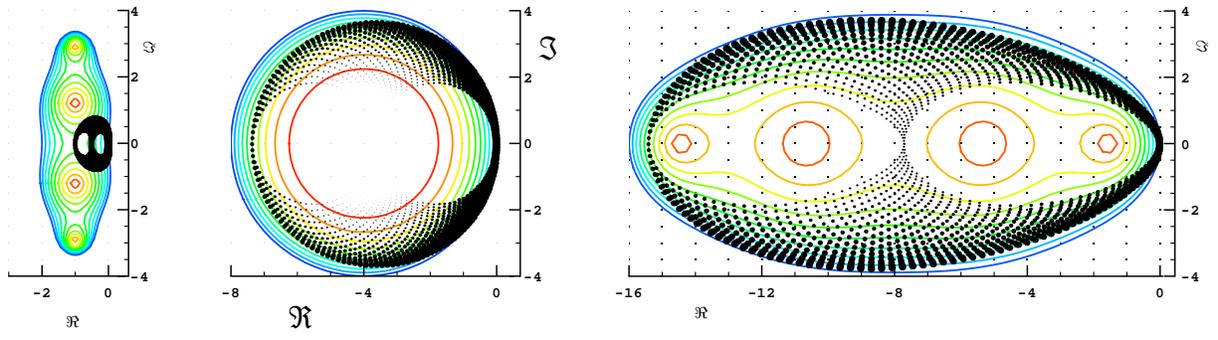


Figure 6: Sketch of optimization procedure where $|P| = 1$ is the temporal stability boundary and \mathcal{F} is the spatial footprint.

¹⁷ Or, if unable to obtain the requested damping because it is too low or too close to the inviscid limit, use the maximum damping available.

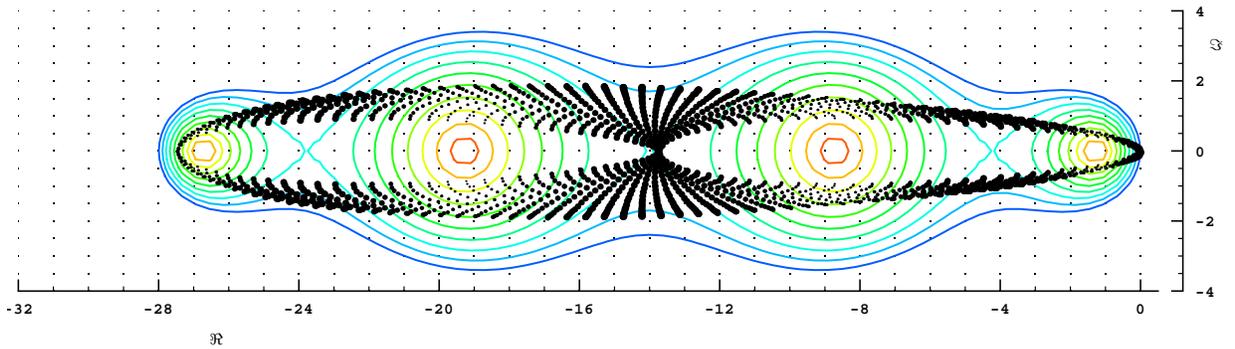
¹⁸ William L. Kleb, *Matching Multi-stage Schemes to Viscous Flow*, Ph.D. thesis, University of Michigan, 2004.



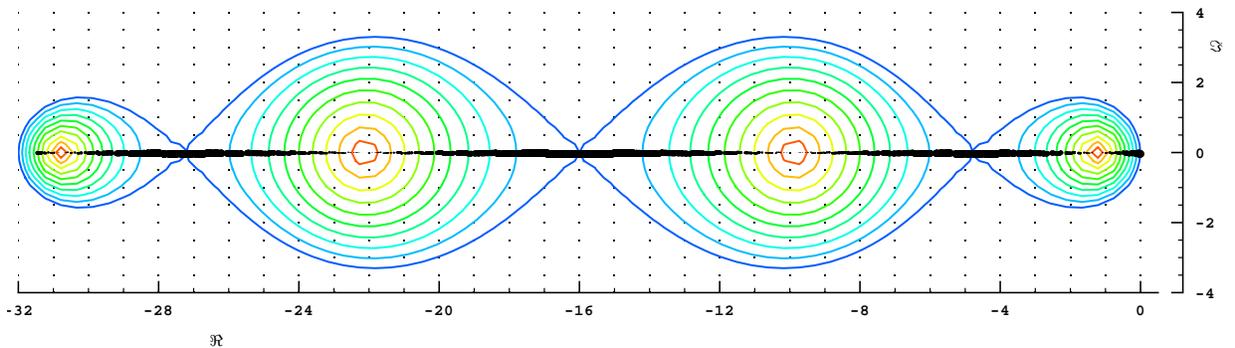
(a) $d = -1$

(b) $d = -4$

(c) $d = -8$



(d) $d = -14$



(e) $d = -16$

Figure 7: Fourier design plots as a function of the temporal operator parameter, d , for Fromm's Scheme ($\kappa = 0$) with a prescribed high-frequency damping of 0.5, at least 0.98 damping for $\beta_x, \beta_y \in [\pi/3, \pi]$, and stable otherwise. Damping contours are shown every 0.1. The frequency symbols are sized by their respective damping.

coefficients are stored in tabular form and queried based the cell Reynolds number. On the order of 500 Reynolds numbers are used to span the domain of significant variation (0.01 to 100), and zeroth order extrapolation is used for values outside the lookup table's domain.

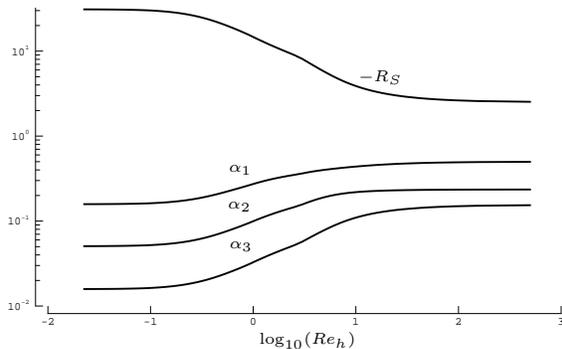


Figure 8: The negative real Fourier extent of the spatial operator and the multistage coefficients as a function of the cell Reynolds number for damping criteria suitable for single-grid relaxation: 0.5 at (π, π) with 0.98 for β_x or $\beta_y \in [\pi/3, \pi]$.

Table 1: A sampling of the data in figure 8.

Re_h	$-R_S$	α_1	α_2	α_3	α_4
100	2.53	0.1538	0.2352	0.4981	1.0000
10	3.88	0.1099	0.2191	0.4381	1.0000
5	5.70	0.0787	0.1913	0.4011	1.0000
2	9.98	0.0472	0.1355	0.3349	1.0000
1	14.5	0.0332	0.1008	0.2754	1.0000
0.5	20.5	0.0237	0.0742	0.2178	1.0000
0.1	29.5	0.0166	0.0531	0.1647	1.0000
0.01	30.4	0.0162	0.0516	0.1607	1.0000

To accommodate multigrid relaxation, a range of damping criteria were explored to find one that minimized the maximum damping over the hi-hi frequency range but still resulted in parameters that spanned the cell Reynolds number range from the inviscid limit to the viscous limit. Figure 9 represents the best compromise found to date. By comparing tables 1 and 2, the trade off between damping properties and the time step (proportional to the negative-real extent, R_S) is readily apparent.

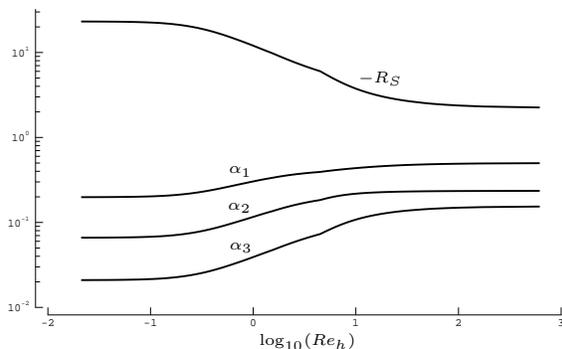


Figure 9: The negative real Fourier extent of the spatial operator and the multistage coefficients as a function of the cell Reynolds number for damping criteria supporting multigrid relaxation: 0.4 at (π, π) with 0.6 for β_x and $\beta_y \in [\pi/2, \pi]$.

Table 2: A sampling of the data in figure 9.

Re_h	$-R_S$	α_1	α_2	α_3	α_4
100	2.37	0.1493	0.2342	0.4912	1.0000
10	3.72	0.1087	0.2183	0.4367	1.0000
5	5.61	0.0775	0.1897	0.3994	1.0000
2	8.55	0.0535	0.1493	0.3540	1.0000
1	12.0	0.0392	0.1164	0.3043	1.0000
0.5	16.3	0.0293	0.0902	0.2537	1.0000
0.1	22.6	0.0215	0.0677	0.2023	1.0000
0.01	23.2	0.0210	0.0661	0.1983	1.0000

The following section provides a numerical demonstration of these Fourier designs for a scalar model problem while section 5 explores their use for the Navier-Stokes equations.

3 Scalar Results

This section presents numerical results obtained from applying the variable-coefficient multistage schemes described in the previous section.¹⁹ The scalar advection-diffusion model problem is solved by using upwind, finite-volume techniques on a uniform Cartesian mesh.²⁰ This serves to focus attention on the effects of the variable-coefficient multistage time stepping without the additional complexity of local preconditioning necessary for systems of equations, which will be covered in the following section.

Specifically, a circular-advection velocity field²¹ was chosen because it contains a wide range of cell Reynolds number and flow angles.²² The domain consists of 64 equally-spaced cells along the x direction from -1 to 1 and 32 equally-spaced cells along the y direction from 0 to 1, with a cosine profile from $x = -0.8$ to -0.2 at the lower left rotating clockwise to exit at the lower right.

Figure 10 shows a representative solution for a Reynolds number of 500 based on unit length and velocity. For this problem, the cell Reynolds numbers are a factor of 32 (the inverse of the grid spacing, h) less than the global Reynolds number because $Re_h \propto Re \cdot h \cdot q$. In addition to the factor due to cell size, the local cell Reynolds numbers also vary by factors of near zero at the center to $\sqrt{2}$ at the corners due to the advection velocity field. So, if the global Reynolds number is 10^4 , the local cell Reynolds numbers will vary from near 0 to $\sqrt{2} \times 10^4/h$ depending on their radial location.

The work to reduce the L_2 error norm of the solution to 10^{-10} are recorded in tables 3 and 4 for a range of Reynolds numbers for both the variable-coefficient multistage scheme and the standard, fixed-coefficient scheme. The fixed-coefficient scheme uses advection-optimized multistage coefficients that maximize damping over frequencies from $\pi/4$ to π due to Tai,²³ while the variable-coefficient scheme uses the cell-Reynolds-number-dependent coefficients that were developed in the proceeding analysis.

For single-grid relaxation, table 3 clearly demonstrates the devastating effect that low Reynolds numbers (and hence low cell Reynolds numbers) have on the standard advection-optimized time marching scheme by increasing time to convergence by a factor of 100. Meanwhile, the variable-coefficient scheme retains a healthy convergence rate even with local cell Reynolds numbers well below 0.01 and shows an order-of-magnitude improvement over the fixed-coefficient scheme. For the larger Reynolds numbers (≥ 100), however, the fixed-coefficient scheme is slightly more efficient than the variable-coefficient scheme because the fixed-coefficient

¹⁹No CPU timing data is presented for this proof-of-concept study because no optimization of the new implementation has been attempted.

²⁰The code was written using test-driven development as part of an Extreme Programming pathfinder project and is available from William A. Wood and William L. Kleb, "Multi-Stage Runge-Kutta Circular Advection Solver, Release 3.0", *NASA Tech Briefs*, 2002, p. 43.

²¹By setting $a = -y$ and $b = x$ in equation 6 on page 4.

²²The related Smith-Hutton problem was solved as a code verification case.

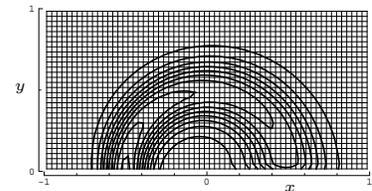


Figure 10: Circular advection of a half cosine wave from $x = -0.8$ to -0.2 ($y = 0$) for a Reynolds number of 500 based on unity length and velocity. Solution contours are shown every 0.1 from 0.05 to 0.95.

²³C. H. Tai, *Acceleration Techniques for Explicit Euler Codes*, Ph.D. thesis, University of Michigan, 1990.

Table 3: Comparison of *single-grid relaxation* for fixed- and variable-coefficient multistage schemes for circular advection on a 64×32 grid.

Re	Iterations	
	Varied	Fixed*
0.0001	1,400	9,600
0.001	560	8,900
0.01	510	8,100
0.1	470	7,300
1	420	6,600
10	300	4,400
100	77	610
1,000	110	100
10,000	120	87
100,000	120	87
1,000,000	120	86

*With $\nu = 1.29$ and coefficients: 0.1667, 0.3027, 0.5276, and 1.000.

scheme has been optimized for the inviscid limit whereas the variable-coefficient scheme merely uses the inviscid end of the Manteuffel family.²⁴

Table 4 shows similar results for a multigrid scheme using variable and fixed coefficients optimized over the frequency range of $\pi/2$ to π . Again, the variable-coefficient multistage scheme shows less variation due to Reynolds number. However, in this case, the high-frequency damping properties hold more sway for multigrid relaxation than large time steps, and the variable-coefficient scheme, while still beneficial at lower Reynolds numbers, does not yield the order-of-magnitude improvement seen for the single-grid scheme. The fixed-coefficient scheme is again slightly better for $Re \geq 1000$ because the Manteuffel transformation used to create the variable-coefficient scheme does not give optimal properties for the inviscid limit, but these local Reynolds numbers are well above those that typically reside in the computationally-stiff regions of aerodynamic flow, i.e., viscous boundary layers.

4 Local Preconditioning

Preconditioning appears in the CFD community in two very distinct contexts. In one context, preconditioning is applied to a *global* system of *discrete* equations to improve solution methods such as the conjugate-gradient method. Whereas in the other context, *local* preconditioning is directly applied to the *analytic* governing equations. This study employs local preconditioning, which is a technique to remove stiffness from a system of equations.

Convergence to steady-state for inviscid calculations is impaired in some flow regimes due to the spread in the characteristic wave speeds. For example, for waves moving in the streamwise direction, the wave speeds are q , $q+a$, and $q-a$, where q is the total flow speed and a is the speed of sound. The ratio of the largest to smallest wave speed is termed the condition number, C , and serves as a measure of “stiffness.” Figure 11 shows the condition number as a function of Mach number. The solid line represents the original, unconditioned Euler equations. Infinite stiffness occurs in both the subsonic and transonic regimes, and as the Mach number increases supersonically, the condition number asymptotes to the ideal value of one.

The dashed line in figure 11 represents the Euler equations preconditioned with the Van Leer-Lee-Roe preconditioner.²⁵ This preconditioner comes closest to achieving equalization of wave speeds without affecting their direction of propagation or their hyperbolicity. The dashed line shows it is possible to completely eliminate the stiffness as the Mach number

²⁴ For more information, see William L. Kleb, *Matching Multistage Schemes to Viscous Flow*, Ph.D. thesis, University of Michigan, 2004.

Table 4: Comparison of *multigrid relaxation* for fixed- and variable-coefficient multistage schemes for circular advection on a 64×32 grid.

Re	Work Units	
	Varied	Fixed*
0.0001	210	670
0.001	180	610
0.01	170	540
0.1	150	500
1	140	440
10	110	370
100	86	170
1,000	130	110
10,000	160	100
100,000	160	100
1,000,000	160	100

* With $\nu = 1.39$ and coefficients: 0.1400, 0.2939, 0.5252, and 1.000.

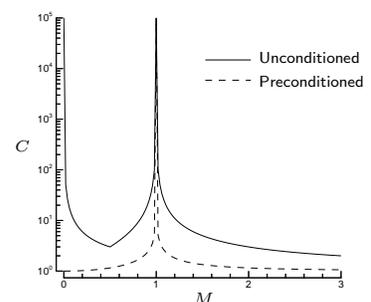


Figure 11: Condition number as a function of Mach number for the Euler equations.

²⁵ Bram van Leer, et al., *Characteristic Time-Stepping or Local Preconditioning of the Euler Equations*, AIAA Paper 91-1552, 1991.

approaches zero, reduce the order of the transonic singularity, and in general, substantially lower the condition number for the system of equations. This preconditioning makes the system behave much like a scalar equation that has a single wave speed and facilitates the application of time-marching schemes designed for scalar equations.

So, instead of marching to the steady state with,

$$\vec{U}_t = Res(\vec{U}), \quad (20)$$

march with a preconditioned residual,

$$\vec{U}_t = \mathbf{P}(\vec{U})Res'(\vec{U}), \quad (21)$$

where the residual is also changed slightly to obtain correct behavior in the incompressible limit. This can also be viewed as marching with a local matrix time step as opposed to a local scalar time step.

5 System Results

Continuing the numerical experiments begun in section 3 for a *scalar* equation, this section provides results for a *system* of equations, namely the Navier-Stokes equations. In order to apply the variable-coefficient multistage schemes developed in section 2, the equations are preconditioned according to the method discussed in the previous section. This preconditioning

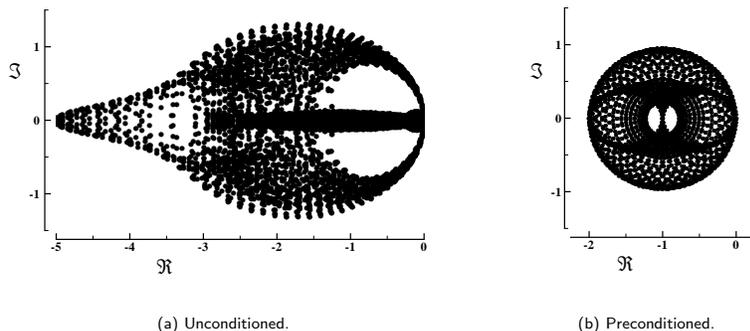


Figure 12: Fourier footprint of the two-dimensional, quasi-linear Navier-Stokes equations, discretized with first-order upwind for $M = 0.1$, $Re_h = 1$, $\mathcal{R} = 1$, $\phi = 45^\circ$, and $\beta_{x,y} = [0, \pi : 21]$.

enables the direct application of the multistage schemes derived in section 3 because the footprint again becomes compact and a function of only a few parameters, i.e., Re_h , \mathcal{R} , M , and ϕ , as shown in figure 12.

The flow solver developed for this portion of the study was purposely kept as simple as possible to isolate the effect of the variable multistage coefficients. The solver consists of a cell-centered, finite-volume-based scheme that uses Roe's

Flux-Difference Splitting²⁶ for the inviscid fluxes and central differencing for the viscous terms on a Cartesian grid of quadrilaterals.

The numerical test problem chosen is that of a laminar boundary-layer flow; specifically, two-dimensional subsonic flow over a flat plate. The unit-length Reynolds number is set at 10,000 and three different freestream Mach numbers are used (0.3, 0.1, and 0.05) to verify that the local preconditioning was implemented correctly. The computational domain and mesh are shown in figure 13. The plate is 4 units long, with

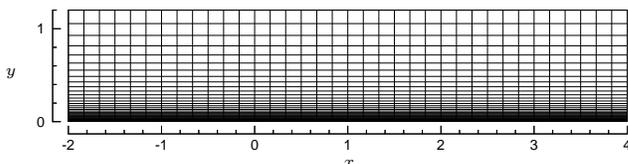


Figure 13: Boundary layer computational domain and grid.

the upstream boundary 2 units away from the leading edge and the upper boundary is placed at 1.2 units. There are 36 cells evenly distributed along the x -direction (24 cells on the plate and 12 upstream) and 40 cells exponentially stretched in the y -direction.²⁷ By employing characteristic boundary conditions at inflow and outflow, neither second-order boundary conditions nor solution-assumed grids are necessary to capture the boundary-layer gradients as shown in figure 14.

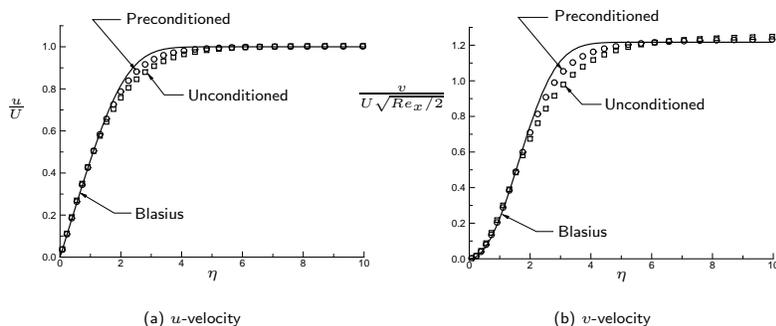


Figure 14: Boundary layer solution profiles at $x/L = 0.521$ for Mach 0.1 flow.

Table 5 shows the number of iterations to converge the L_2 error norm to 10^{-6} for a range of Mach numbers for both the fixed-coefficient multistage scheme and the variable-coefficient multistage scheme. The fixed-coefficient scheme was run with and without preconditioning.

For the fixed-coefficient scheme, the results affirm the claims of preconditioning by providing nearly Mach-number-independent convergence. Some degradation is present,

²⁶ Philip L. Roe, “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes”, *Journal of Computational Physics*, 43(2), 1981, pp. 357–372.

²⁷ For details, see William L. Kleb, et al., *Efficient Multi-Stage Time Marching for Viscous Flows via Local preconditioning*, AIAA Paper 99-3267, 1999.

Table 5: Comparison of iterations to convergence for fixed- and variable-coefficient multistage schemes for viscous flow over a flat plate.

M	Fixed		Varied*
	Uncond.	Precond.	
0.3	21,000	7,700	3,100
0.1	45,000	7,500	1,900
0.05	97,000	8,200	1,500

* Requires preconditioning.

however, as the cell-Reynolds numbers decrease with Mach number. This effect was predicted by Lee²⁸ and shows the limitation of simply using the viscous Jacobians to augment the Euler preconditioning. Depcik and Van Leer have recently presented a better-conditioned viscous preconditioner.²⁹

The results also show that the variable-coefficient scheme provides further improvement in convergence efficiency compared to the fixed-coefficient scheme. The variable-coefficient scheme reaches convergence in nearly a factor of five fewer iterations than the fixed-coefficient scheme for the most difficult case (Mach 0.05).

6 Concluding Remarks

Beginning with the notion that multistage time-marching schemes have considerable malleability with respect to time step coefficients when no longer held to rigid time accuracy constraints, it has been shown that they can indeed be made effective in what was formerly the exclusive domain of implicit schemes: viscous flows.

A set of multistage coefficients were created for a given spatial discretization as a function of cell Reynolds number. As part of this process, a precise definition for the numerical time step was derived.

The next step was to numerically test these optimized sets of multistage coefficients that vary timestep to timestep and cell to cell. To isolate the effects of the variable multistage coefficients the test case chosen was very simple: circular advection-diffusion. The numerical results supported the analytical analysis by demonstrating an order-of-magnitude improvement in convergence rate for single-grid relaxation and a factor of 3 for multigrid relaxation.

With the success of the scalar case, preconditioning was applied to make the Navier-Stokes system of equations behave more nearly as a single scalar equation. Then, by applying the variable multistage coefficient scheme to a typical boundary-layer flow problem, the results affirmed the benefits of local preconditioning and showed up to a factor of 5 savings over a fixed-coefficient scheme.

However, because this is a proof-of-concept study, many variables were purposely held constant or ignored entirely,³⁰ so the effects of the cell Reynolds number-dependent multistage coefficients could be ascertained. Given the improvement over the standard fixed coefficient method, however, this approach merits further exploration.

²⁸ Dohyung Lee, *Local Preconditioning of the Euler and Navier-Stokes Equations*, Ph.D. thesis, University of Michigan, 1996.

²⁹ Christopher Depcik and Bram van Leer, *In Search of an Optimal Local Navier-Stokes Preconditioner*, AIAA Paper 2003-3703, 2003.

³⁰ Notably, the computational overhead associated with obtaining and using a different set of multistage coefficients for each cell in the mesh.

Acknowledgments

The authors would like to acknowledge the help of Ken Powell, and Philip Roe of the University of Michigan; Eli Turkel of Tel Aviv University, and Peter Gnoffo of NASA.

About the Authors



BIL KLEB, a lifetime senior AIAA member, has worked in the area of computational aerothermodynamics at NASA's Langley Research Center for the past 15 years. During this time, he pioneered the first full-vehicle reentry simulation of the Shuttle Orbiter and earned an MBA from The College of William and Mary and a PhD of Aerospace Engineering from the University of Michigan. For the past two years, Bil has been a steward of the FUN3D software development team.³¹

Email: Bil.Kleb@NASA.gov

³¹ fun3d.larc.nasa.gov, last accessed June 1st, 2005.



BRAM VAN LEER has worked in the field of CFD since transitioning from computational astrophysics shortly after receiving his Ph.D. from Leiden State University. Bram has been a professor of Aerospace Engineering since 1986 and was elected an AIAA Fellow in 1995.

Email: Bram@UMich.edu



BILL WOOD, a lifetime senior AIAA member with a Ph.D. in Aerospace Engineering from Virginia Tech, has worked in the field of CFD at NASA Langley Research Center's Aerothermodynamics Branch since 1991.

Email: Bill.Wood@NASA.gov

Colophon

This document was typeset in Computer Modern font with the L^AT_EX typesetting system using the `handout` option of the AIAA package,³² version 3.8. The `handout` class option used here strives toward the layout style espoused by the visual information design expert Edward Tufte.³³ Also employed were the `color`, `subfigure`, `subfigmat`, `psfrag`, `amsmath`, `bm`, `array`, `xspace`, `textcomp`, `bibentry`, `booktabs`, `multirow`, `threeparttable`, `varioref`, `wrapfig`, `hyperref`, and `nohyperref` packages.

³² www.ctan.org, last accessed June 1st, 2005.

³³ Edward R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, 1983; Edward R. Tufte, *Envisioning Information*, Graphics Press, 1990; and Edward R. Tufte, *Visual Explanations: Images and Quantities, Evidence and Narrative*, Graphics Press, 1997.