JOURNAL OF COMPUTATIONAL PHYSICS **127,** 346–362 (1996)
ARTICLE NO. 0180

The Importance of Eigenvectors for Local Preconditioners of the Euler Equations

D. L. DARMOFAL* AND P. J. SCHMID†

*Department of Aerospace Engineering, Texas A & M University, College Station, Texas 77843; †Department of Applied Mathematics, University of Washington, Seattle, Washington 98195

Received September 27, 1995; revised March 27, 1996

The design of local preconditioners to accelerate the convergence to a steady state for the compressible Euler equations has so far been solely based on eigenvalue analysis. However, numerical evidence exists that the eigenvector structure also has an influence on the performance of preconditioners and should therefore be included in the design process. In this paper, we present the mathematical framework for the eigenvector analysis of local preconditioners for the multi-dimensional Euler equations. The non-normality of the preconditioned system is crucial in determining the potential for transient amplification of perturbations. Several existing local preconditioners are shown to possess a highly non-normal structure for low Mach numbers. This non-normality leads to significant robustness problems at stagnation points. A modification to these preconditioners which eliminates the non-normality is suggested, and numerical results are presented showing the marked improvement in robustness.

1. INTRODUCTION

Recently, several authors have investigated the possibility of locally preconditioning the Euler equations to accelerate convergence to a steady state [1–7]. A review of the current state of preconditioning was given by Turkel [8]. The essential idea of local preconditioning is to premultiply the spatial operator by a preconditioning matrix, **P**. Specifically, we wish to solve the two-dimensional Euler equations which may be written

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{A}\frac{\partial \mathbf{U}}{\partial x} + \mathbf{B}\frac{\partial \mathbf{U}}{\partial y} = 0, \tag{1}$$

where **U** is the state vector and **A** and **B** are matrices which depend on the local flow state. The preconditioned form of these equations is

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{PA}\frac{\partial \mathbf{U}}{\partial x} + \mathbf{PB}\frac{\partial \mathbf{U}}{\partial y} = 0. \tag{2}$$

In the steady state, the solution remains unchanged as long as the preconditioner **P** is invertible.

Most previous preconditioning efforts have focused on manipulating the eigenvalues of the spatial operator. For example, Turkel [2] derives a family of preconditioners which reduces the spread of the wave speeds for pseudo-compressible and low Mach number compressible flows. In [3], van Leer *et al.* derive a symmetric preconditioner for the two-dimensional Euler equations which reduces the spread of the characteristic speeds across from $(M + 1)/\min(M, |M - 1|)$ to $1/\sqrt{1 - \min(M^2, M^{-2})}$, where M is the Mach number. Lee [9] shows that this is the lowest ratio of characteristic speeds attainable using a symmetric preconditioner. For grid-aligned upwind schemes, Allmaras [6] finds that a block Jacobi preconditioner clusters the high frequency eigenmodes of the two-dimensional Euler and Navier–Stokes discretized operator allowing the formulation of an effective smoother for multigrid algorithms. Although the block Jacobi preconditioner does cluster the discrete high frequency modes, it does very little to reduce the spread of characteristic speeds for low frequency modes [9].

While these previous investigations have focused on altering the eigenvalues of the preconditioned system, recent evidence suggests that the eigenvectors must also play some role in determining the effectiveness of a preconditioner [4, 8, 10]. The goal of this paper is to clarify the importance of eigenvectors in effective preconditioner design. We concentrate on one particular problem, transient energy growth, that results from poorly conditioned eigenvectors. Due to the lack of eigenvector orthogonality, we show—theoretically as well as numerically—that small perturbations in a linearized evolution problem can be significantly amplified over short time scales, while the long time or asymptotic behavior of the linearized system is governed by the eigenvalues. However, for practical applications to a nonlinear problem, this short time non-normal growth may completely alter the mean state such that the predicted long time asymptotic behavior is never observed. As we demonstrate using nonlinear preconditioned Euler simulations, non-normal amplification does occur and, in practice, results in a significant lack of robustness.

0021-9991/96 $18.00

Our analysis of transient growth in non-normal preconditioned systems is based on the experiences with non-normal analysis in hydrodynamic stability [11–14]. Specifically, these efforts have shown that substantial perturbation energy growth can occur over short time scales for incompressible shear flows in which all eigenmodes are damped (i.e., the flow is subcritical). Similar techniques have also been applied to numerical algorithms to derive timestep restrictions which limit error amplification [15–18]. A recent review of the subject has been given by Van Dorsselaer *et al.* [19]. All of these results underscore the importance of considering more than just the eigenvalues in determining the behavior of any evolution process.

A fundamental issue in the analysis of eigenvector orthogonality is the choice of dependent variables. This choice changes the eigenvector basis and, therefore, changes the results of any analysis. In this paper, we have chosen to work with the Euler equations in symmetrizing variables. The symmetrizing variables are those variables for which the **A** and **B** matrices are symmetric. As we show, the symmetrizing variables are convenient for at least two reasons. First, in the symmetrizing variables, the unpreconditioned Euler equations do not have any transient growth because the eigenvectors are orthogonal. Thus, a preconditioner in the symmetrizing variables which preserves this property (or at least does not introduce significant non-orthogonality) should possess eigenvectors which are as well conditioned as the unpreconditioned Euler equations. Our second reason for choosing the symmetrizing variables is that non-dimensionalization (in particular by the reference velocity) does not alter the eigenvectors in this basis. This is not true for most other bases (such as the conserved variables) and therefore would significantly complicate the analysis.

In this paper, we consider only subsonic preconditioners. We limit our discussion to the subsonic regime because preconditioning in supersonic flow is generally more straightforward and because significant problems are known to still exist in current subsonic multi-dimensional preconditioners. In particular, many existing preconditioners exhibit severe robustness problems at stagnation points [20]. We show that the stagnation point problems are a result of eigenstructure non-orthogonality as the Mach number approaches zero; furthermore, by correcting the non-normality, the robustness problems can be eliminated.

The remainder of the paper is organized as follows. First, the basic theory for analyzing non-normal effects of a general, linear evolution problem is described. Then, this theory is applied to the linearized preconditioned Euler equations for several different preconditioners. Next, the results of this analysis are verified using a two-dimensional linearized preconditioned Euler solver. Finally, techniques to bound non-normal effects in conjunction with local preconditioners are presented and the resulting modified preconditioners are used in fully nonlinear preconditioned Euler calculations. These nonlinear simulations show a significant improvement in robustness when the preconditioners are modified to bound non-normal amplification.

2. THEORY

Although the Euler equations provide the specific examples in later sections, the basic concepts can be formulated in fairly general terms. Let us consider the initial value problem,

$$\frac{d\mathbf{u}}{dt} + i\mathbf{L}\mathbf{u} = 0, \tag{3}$$

where $\mathbf{L}$ is an $N \times N$ matrix. The solution to this equation can be written very compactly using the matrix exponential as

$$\mathbf{u}(t) = \exp(-it\mathbf{L})\mathbf{u}_0, \tag{4}$$

where $\exp(t\mathbf{L}) = \mathbf{I} + t\mathbf{L} + (t^2/2)\mathbf{L}^2 + \cdots$ and $\mathbf{u}_0$ is the initial condition.

Assuming the spatial operator can be diagonalized, we can decompose $\mathbf{L}$ into

$$\mathbf{L} = \mathbf{R}\boldsymbol{\Omega}\mathbf{R}^{-1}, \tag{5}$$

where $\mathbf{R}$ is the eigenvector matrix with unit norm eigenvectors and $\boldsymbol{\Omega} = \mathrm{diag}(\omega_1, \ldots, \omega_N)$ is a diagonal matrix with the eigenvalues as its entries. The eigenmodes can be summed to give the solution to Eq. (3) as

$$\mathbf{u}(t) = \sum_{n=1}^{N} (\mathbf{l}_n^H \mathbf{u}_0)\mathbf{r}_n e^{-i\omega_n t}, \tag{6}$$

where $\mathbf{r}_n$ are the right eigenvectors from the columns of $\mathbf{R}$, and $\mathbf{l}_n^H$ are the left eigenvectors from the rows of $\mathbf{R}^{-1}$. In the long time limit, a system is asymptotically stable if all eigenvalues have negative imaginary parts. In this case, all perturbations decay exponentially as $t \to \infty$. In our analysis of the preconditioned Euler equations, we are interested in neutrally stable systems for which all eigenvalues are purely real.

Although a system may be asymptotically (or neutrally) stable as $t \to \infty$, the possibility of short time or transient amplification still exists. To demonstrate this mechanism, we need to describe the evolution of the system's total energy. We define the total energy of a system using the vector L_2-norm as

$$E = \|\mathbf{u}\|^2 = \|\mathbf{u}\|_2^2 = \mathbf{u}^{\mathrm{H}}\mathbf{u},$$

and the matrix L_2-norm is defined as

$$\|\mathbf{L}\| = \|\mathbf{L}\|_2 = \max_{\mathbf{u}} \|\mathbf{L}\mathbf{u}\| \quad \text{for } \|\mathbf{u}\| = 1,$$

where the superscript ^{H} denotes the Hermitian of the vector or matrix. The maximum amplification, $G(t)$, at a given instant in time for a unit norm initial condition is simply the norm of the matrix exponential,

$$\begin{aligned} G(t) &= \sup_{\|\mathbf{u}_0\|=1} \|\mathbf{u}\| = \|\exp(-it\mathbf{L})\|, \\ &= \|\mathbf{R}\exp(-it\mathbf{\Omega})\mathbf{R}^{-1}\|. \end{aligned}$$

Even for a stable system, this amplification can be greater than one. In fact, the ability to amplify perturbations can be directly connected to the eigenvector conditioning. The maximum amplification for all time may be defined as

$$G_{\max} \equiv \sup_{t\geq 0} G(t).$$

For asymptotically (and neutrally) stable systems, the value of $G_{\max}$ can be bounded by

$$1 \leq G_{\max} \leq \kappa(\mathbf{R}), \tag{7}$$

where $\kappa(\mathbf{R})$ is the condition number of $\mathbf{R}$ defined as $\kappa(\mathbf{R}) = \|\mathbf{R}\|\,\|\mathbf{R}^{-1}\|$. The lower bound of $G_{\max}$ occurs at $t = 0$ for which $G(t) = 1$ by definition. For the upper bound, the condition number of the eigenvector matrix scales the maximum possible energy growth. For an asymptotically stable, perfectly conditioned system, that is, $\kappa(\mathbf{R}) = 1$, the maximum amplification $G_{\max}$ is identically one. A matrix, $\mathbf{L}$, whose eigenvector matrix $\mathbf{R}$ satisfies $\kappa(\mathbf{R}) = 1$, is termed normal and has orthogonal eigenvectors. Examples of normal matrices include Hermitian matrices (of which a subcategory is real symmetric matrices). For non-normal systems with $\kappa(\mathbf{R}) > 1$, the larger the departure from eigenvector orthogonality, the larger the possible amplification. Thus, although a system may be asymptotically stable, short time or transient energy amplification is still possible when the system is non-normal.

Next, the evolution of the energy can be found,

$$\begin{aligned} \frac{d}{dt}\|\mathbf{u}\|^2 &= \frac{d}{dt}\mathbf{u}^{\mathrm{H}}\mathbf{u}, \\ &= -i\mathbf{u}^{\mathrm{H}}(\mathbf{L} - \mathbf{L}^{\mathrm{H}})\mathbf{u}. \end{aligned} \tag{8}$$

A fundamental result from this equation is that if $\mathbf{L} = \mathbf{L}^{\mathrm{H}}$, i.e., the matrix is Hermitian, the total energy remains unchanged. The right-hand side of Eq. (8) can be used to compute the maximum possible amplification rate over all unit norm initial conditions,

$$\left.\frac{dG}{dt}\right|_{t=0^+} = \frac{1}{2}\lambda_{\max}[-i(\mathbf{L} - \mathbf{L}^{\mathrm{H}})],$$

where $\lambda_{\max}$ is the maximum eigenvalue of a matrix. Thus, even with a non-normal matrix, the total energy will decay when $\lambda_{\max}[-i(\mathbf{L} - \mathbf{L}^{\mathrm{H}})] < 0$.

To illustrate the basic mechanism of the transient energy growth, it is instructive to expand the energy in terms of the eigenmodes of $\mathbf{L}$,

$$\|\mathbf{u}\|^2 = \sum_{j,k=1}^{N} v_j^* v_k \mathbf{r}_j^{\mathrm{H}}\mathbf{r}_k \exp[-i(\omega_j - \omega_k)t], \tag{9}$$

where $v_j = \mathbf{l}_j^{\mathrm{H}}\mathbf{u}_0$ and the superscript * indicates the complex conjugate. If the eigenvectors are orthogonal, then $\mathbf{r}_j^{\mathrm{H}}\mathbf{r}_k = \delta_{jk}$, otherwise, the eigenmodes couple and produce a transient change in the energy. A simple demonstration of the eigenmode coupling and its relation to transient energy growth is given in the Appendix.

As shown above, the condition number gives a first estimate of transient growth. Sharper estimates of the maximum transient amplification of energy can be derived using the resolvent norm defined as

$$R(z) = \|(z\mathbf{I} + i\mathbf{L})^{-1}\|.$$

For $z \in \{-i\omega_n\}$ the matrix $z\mathbf{I} + i\mathbf{L}$ is singular and the resolvent norm is assumed to be infinite. An upper bound on $G(t)$ can be derived using the matrix equivalent of Cauchy's integral formula,

$$\begin{aligned} 2\pi\|\exp(-it\mathbf{L})\| &= \left\| \oint_\Gamma \exp(-itz)(z\mathbf{I} + i\mathbf{L})^{-1}\, dz \right\|, \\ &\leq \oint_\Gamma |\exp(-itz)|\,\|(z\mathbf{I} + i\mathbf{L})^{-1}\|\,|dz|, \end{aligned}$$

where Γ stands for an integration contour including the spectrum of $-i\mathbf{L}$. A lower bound for the maximum possible amplification can also be derived based on the resolvent norm. It is given as

$$\max_{t\geq 0} \|\exp(-it\mathbf{L})\| \geq C',$$

where C' is defined as

$$C' = \sup_{\mathrm{Re}(z)>0} \mathrm{Re}(z)\|z\mathbf{I} + i\mathbf{L})^{-1}\|.$$

This estimate can be derived using the Hille–Yosida theorem for linear operators and the interested reader is referred to Kato [21].

As a result of this connection between the resolvent norm and the maximum amplification, plots of the resolvent norm are extremely useful in quantifying transient growth. In addition to the upper and lower bounds mentioned above, plots of the resolvent norm can be used more directly to detect the potential for transient growth and to estimate the maximum possible amplification. We will use the relation between the resolvent norm at a point z in the complex plane and the distance of this point to the spectrum to draw conclusions about the non-normality of the underlying operator. We can derive

$$R(z) = \|(z\mathbf{I} + i\mathbf{L})^{-1}\| = \|\mathbf{R}(z\mathbf{I} + i\mathbf{\Omega})^{-1}\mathbf{R}^{-1}\| \leq \frac{\kappa(\mathbf{R})}{\min_j|z + i\omega_j|}.$$

In the case of normal matrices, i.e., $\kappa(\mathbf{R}) = 1$, contours of the resolvent norm in the complex z-plane represent the inverse of the distance from the closest eigenvalue. For non-normal matrices, i.e., $\kappa(\mathbf{R})$, the resolvent norm at z can be considerably larger than the inverse distance from the closest eigenvalue. Thus, any deviation of the resolvent norm contours from a plot of the inverse distance function from the spectrum indicates that the operator is non-normal and accommodates transient behavior. A simple example of how to interpret contour plots of the resolvent norm is given in the Appendix.

3. ANALYSIS

In this section, we apply the analysis tools derived above to the two-dimensional preconditioned linearized Euler equations. After a Fourier transform in space, the preconditioned Euler equations can be expressed as an initial value problem (Eq. (3)) with

$$\mathbf{L} = k\mathbf{P}(\mathbf{A}\cos\phi + \mathbf{B}\sin\phi),$$

where k is the wave number and ϕ is the wave angle. The vector $\mathbf{u}$ is based on stream-aligned symmetrized variables, $(\tilde{p}/\rho a, \tilde{u}, \tilde{v}, \tilde{s})$, where $\tilde{p}$, $\tilde{u}$, $\tilde{v}$, and $\tilde{s}$ are the perturbation pressure, velocity components, and the entropy, respectively. The mean flow density, velocity, and speed of sound are ρ, u, and a. With this choice of variables, the matrices $\mathbf{A}$ and $\mathbf{B}$ are

$$\mathbf{A} = \begin{pmatrix} u & a & 0 & 0 \\ a & u & 0 & 0 \\ 0 & 0 & u & 0 \\ 0 & 0 & 0 & u \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 0 & 0 & a & 0 \\ 0 & 0 & 0 & 0 \\ a & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

and the mean flow Mach number is defined as $M = u/a$. The specific preconditioners, $\mathbf{P}$, which we have analyzed are given below. The wavenumber, k, only serves to scale the time derivative for the Euler equations; thus, without loss of generality, we let $k = 1$ for the rest of the analysis.

3.1. *Choice of Basis*

As we previously described, the choice of basis alters the orthogonality of the eigenvectors. Suppose that the desired state variables are represented by the vector $\mathbf{v}$ and that they are related to the state vector $\mathbf{u}$ by the transformation

$$\mathbf{v} = \mathbf{T}\mathbf{u}.$$

The governing equation for the Fourier transformed $\mathbf{v}$ variables follows easily from Eq. (3),

$$\frac{d\mathbf{v}}{dt} + i\mathbf{L}_v\mathbf{v} = 0,$$

where $\mathbf{L}_v = \mathbf{T}\mathbf{L}\mathbf{T}^{-1}$. From this, it is easy to show that the eigenvector matrix of $\mathbf{L}_v$ is related to the eigenvector matrix of $\mathbf{L}$ by

$$\mathbf{R}_v = \mathbf{T}\mathbf{R}.$$

Furthermore, the condition number of $\mathbf{R}_v$ can be bounded by

$$\max\left[1, \frac{\kappa(\mathbf{R})}{\kappa(\mathbf{T})}\right] \leq \kappa(\mathbf{R}_v) \leq \kappa(\mathbf{R})\kappa(\mathbf{T}). \tag{10}$$

Of course, the lower bound must be at least one since the condition number of any matrix cannot be less than one. Thus, the conditioning of the eigenvectors in a different basis could be different by a factor of $\kappa(\mathbf{T})$ or $1/\kappa(\mathbf{T})$.

In principle, it may be possible to choose a set of dependent variables so that the resulting preconditioned system does not exhibit transient behavior. For example, this could be done by choosing an eigenvector basis for the dependent variables (in this case, the transformation matrix is $\mathbf{T} = \mathbf{R}^{-1}$). However, the energy in this basis will generally depend on the wave angle and thus lead to an amplification measure which unequally weights perturbations of different wave angles. In the case where the transformation matrix is wave angle independent, the resulting basis variables may still be difficult to interpret. In either case, this choice of dependent variables may be unsuited for conclusions about physically interesting variables like pressure, density, and velocity.

In this work, we will analyze the non-normality of the preconditioned system in the symmetrizing variables. This set of dependent variables is clearly connected to physically interesting quantities and has other advantages as well.

For the unpreconditioned equations, the symmetrizing variables lead to symmetric **A** and **B** matrices. Therefore, **R** is an orthogonal basis and $\kappa(\mathbf{R}) = 1$. Thus, from Eq. (10), any poor conditioning of the eigenvectors in another basis is due solely to the poor conditioning of the transformation matrix, **T**. In principle, preconditioning can account for **A** and **B** matrices with poorly conditioned eigenvectors by requiring the preconditioned system to be symmetric—that is, $\mathbf{PA} = (\mathbf{PA})^{\mathrm{T}}$ and $\mathbf{PB} = (\mathbf{PB})^{\mathrm{T}}$. However, we take a slightly different approach in this paper. Namely, we would like to develop a preconditioner for the symmetrizing variables of the original unpreconditioned equations which optimally clusters the eigenvalues while minimizing the departure from eigenvector orthogonality. Since the unpreconditioned system was symmetric, all non-orthogonality in this basis is purely a result of the preconditioning. Also, if the preconditioned eigenvectors of this originally symmetric system are well-behaved, any difficulties with non-orthogonality in a different basis, such as the conservation variables, must be due to the transformation matrix, not the preconditioning.

Unfortunately, Eq. (10) also implies that non-dimensionalization can alter the eigenvector conditioning and therefore affect any eigenvector analysis. Typical non-dimensionalizations can be written as a diagonal transform matrix and the condition number of a diagonal matrix is the ratio of the largest to smallest diagonal element. Therefore, **T** must be a scalar multiple of the identity matrix to have $\kappa(\mathbf{T}) = 1$ or else the eigenvector conditioning will be altered. In order to non-dimensionalize a set of variables by a scalar implies that the variables must all have the same dimensions and that the matrices, **A** and **B**, should have entries with the same dimensions. It is easy to show that the **A** and **B** matrices of the symmetrizing variables (of any system of equations) must have this property. Thus, another reason for our selection of the symmetrizing variables is that the eigenvectors are not affected by non-dimensionalization. In the remainder of the paper, the symmetrized variables have been non-dimensionalized by the mean speed of sound.

3.2. *Results*

We consider three preconditioners for the subsonic Euler equations. The original van Leer preconditioner is defined as

$$\mathbf{P}_{\mathrm{vl}} = \frac{1}{M}\begin{pmatrix} \frac{M^2}{\beta} & -\frac{M}{\beta} & 0 & 0 \\ -\frac{M}{\beta} & \frac{1+\beta}{\beta} & 0 & 0 \\ 0 & 0 & \beta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

The van Leer preconditioned system has been shown by Lee [9] to have the optimal eigenvalues for a symmetric preconditioner. In the same work, Lee has also derived a subsonic preconditioner using the same basic form as Turkel's incompressible preconditioner [2] but again enforcing the optimal eigenvalue distribution. This gives

$$\mathbf{P}_{\mathrm{t}} = \frac{1}{M}\begin{pmatrix} M^2 & 0 & 0 & 0 \\ -M(1+M^2) & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Note that, using the transpose of a preconditioner in the symmetrized variable, the same eigenvalues for the preconditioned system arise. Thus, the transpose of the Turkel preconditioner, $\mathbf{P}_{\mathrm{t}}^{\mathrm{T}}$, is also analyzed. Finally, the fourth preconditioner we consider is the continuous form of the block Jacobi preconditioner as analyzed by Allmaras [6]. Specifically, we analyze $\mathbf{P}_{\mathrm{j}} = a(|\mathbf{A}| + |\mathbf{B}|)^{-1}$, where a is the local speed of sound. The matrix, $|\mathbf{A}|$, is defined as $\mathbf{R}|\mathbf{\Lambda}|\mathbf{R}^{-1}$, where $|\mathbf{\Lambda}|$ is a diagonal matrix containing the absolute values of the eigenvalues of **A**. In terms of the Mach number, $\mathbf{P}_{\mathrm{j}}$ is

$$\mathbf{P}_{\mathrm{j}} = \frac{1}{M}\begin{pmatrix} \frac{M}{1+\beta^2} & -\frac{M^2}{1+\beta^2} & 0 & 0 \\ -\frac{M^2}{1+\beta^2} & \frac{2M}{1+\beta^2} & 0 & 0 \\ 0 & 0 & \frac{M}{M+1} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

In Figs. 1 and 2, contours of the logarithm of the maximum resolvent norm are plotted for all four preconditioners at mean flow Mach numbers of 0.01, 0.5, and 0.99. These resolvent contours were constructed by discretizing the complex z-plane and then determining the largest resolvent norm over all wave angles for a given z-location,

$$R_{\max}(z) = \sup_{\phi} R(z) = \sup_{\phi} \|(z\mathbf{I} + i\mathbf{L})^{-1}\|.$$

The eigenvalues are located on the imaginary axis and, as a result of the non-dimensionalization by the speed of sound, lie between ± 1. A simple indicator of possible transient growth is the value of the $R_{\max}$ at $z = (1, 0)$. For a preconditioned system resulting in a normal evolution matrix, **L**, $R_{\max}$ should be one at the point $z = (1, 0)$ since the nearest eigenvalue at (0, 0) is a distance of one away.

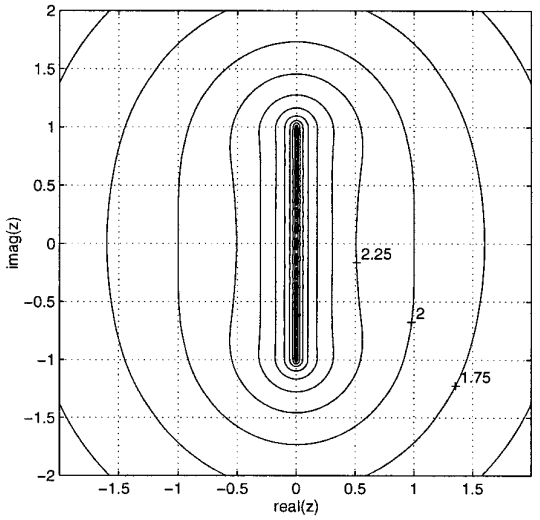

(a) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.01$

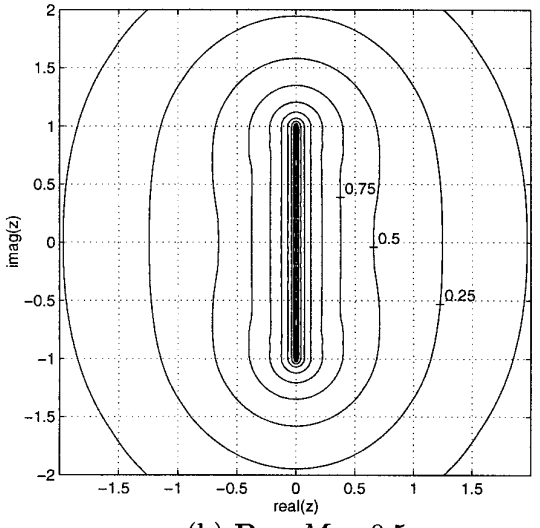

(b) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.5$

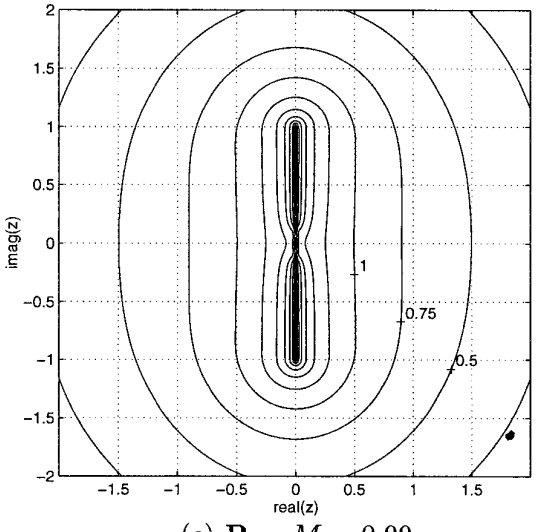

(c) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.99$

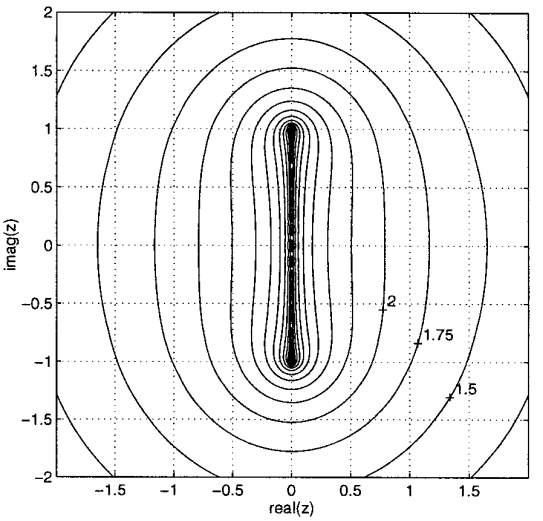

(d) $\mathbf{P}_{\mathrm{t}}$, $M = 0.01$

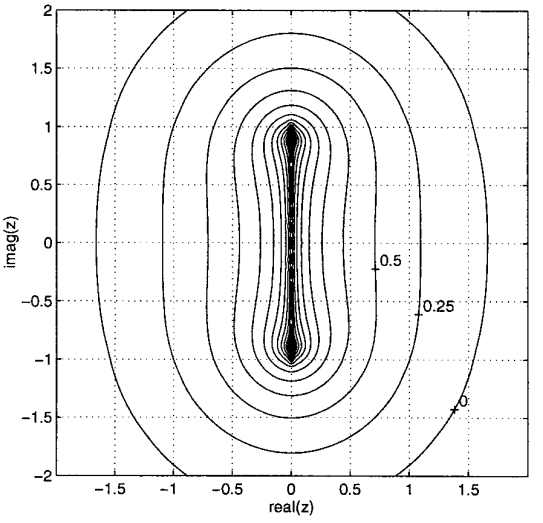

(e) $\mathbf{P}_{\mathrm{t}}$, $M = 0.5$

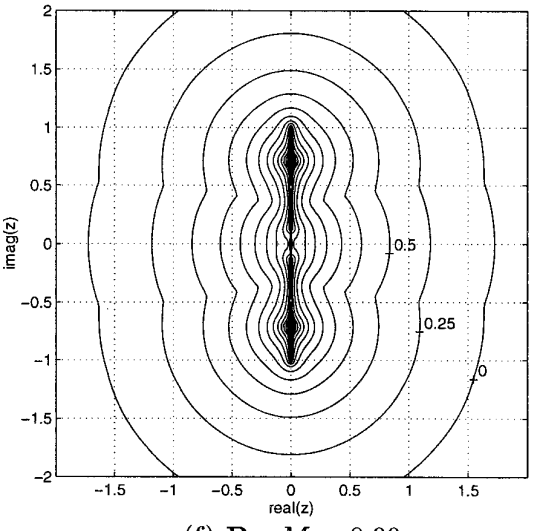

(f) $\mathbf{P}_{\mathrm{t}}$, $M = 0.99$

FIG. 1. Contours of $\log_{10} R_{\max}$, for van Leer and Turkel preconditioners.

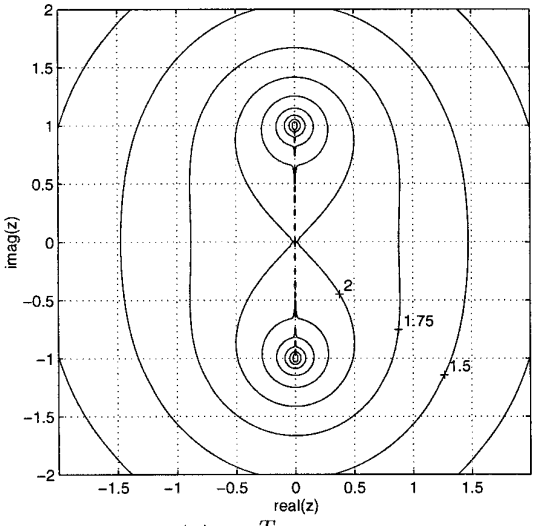

(a) $\mathbf{P}_t{}^T$, $M = 0.01$

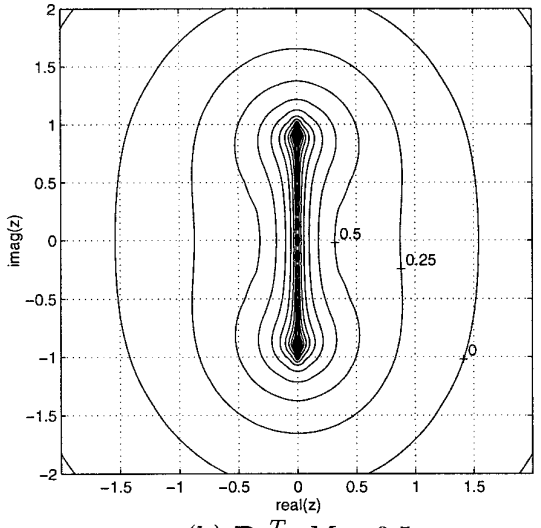

(b) $\mathbf{P}_t{}^T$, $M = 0.5$

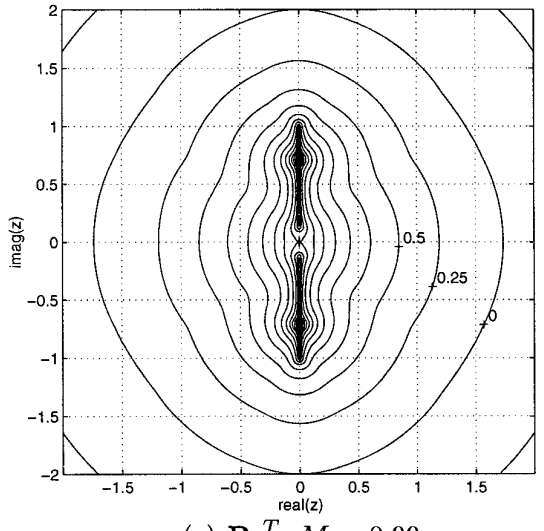

(c) $\mathbf{P}_t{}^T$, $M = 0.99$

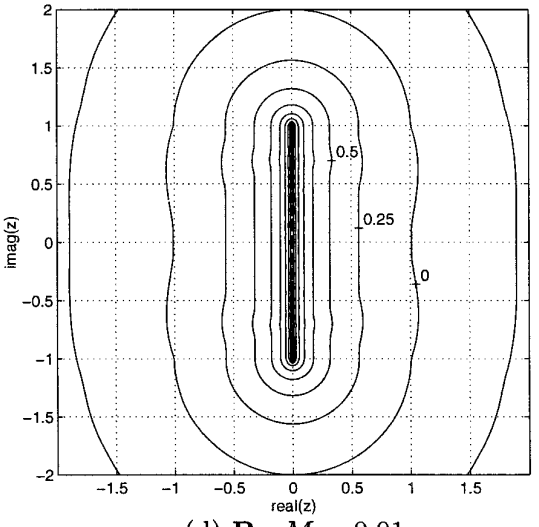

(d) $\mathbf{P}_j$, $M = 0.01$

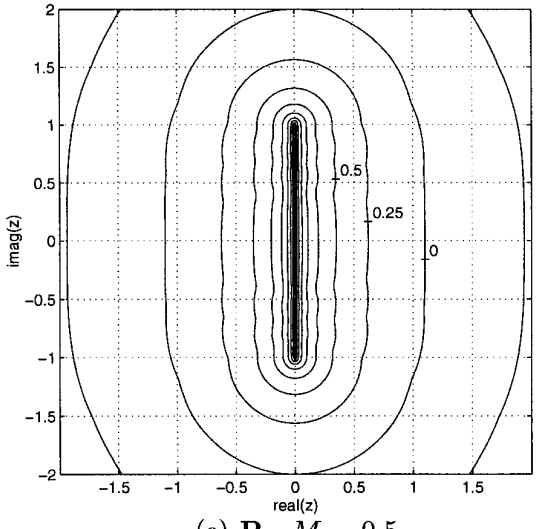

(e) $\mathbf{P}_j$, $M = 0.5$

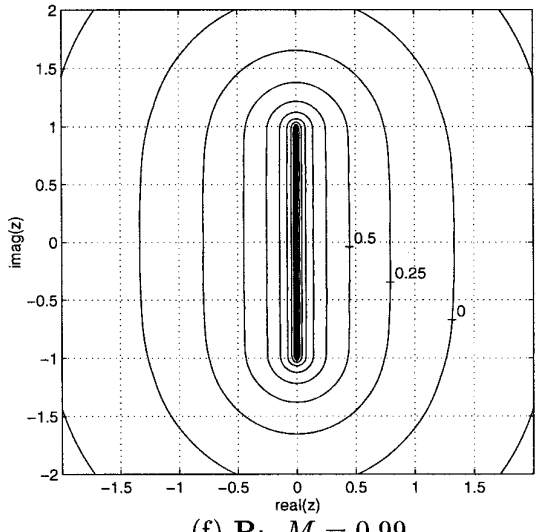

(f) $\mathbf{P}_j$, $M = 0.99$

FIG. 2. Contours of $\log_{10} R_{\max}$ for Turkel transpose and block Jacobi preconditioners.

TABLE I

Initial Growth Bounds, $dG/dt|_{t=0^+}$, for 2D Preconditioned Euler Equations

	M		
	0.01	0.5	0.99
$\mathbf{P}_{\text{vl}}$	100.0	1.9	4.9
$\mathbf{P}_{\text{t}}$	50.0	1.0	1.0
$\mathbf{P}_{\text{t}}^{\text{T}}$	50.0	1.1	1.0
$\mathbf{P}_{\text{j}}$	0.3	0.3	0.5

However, for $M = 0.01$, only the block Jacobi preconditioner is well conditioned. By comparison, the van Leer preconditioner has $R_{\max}(z) = 10^2$ at $z = (1, 0)$. This indicates the potential for transient growth of 100 times the initial perturbation energy. At this point, the Turkel preconditioner and its transpose have resolvent norms of approximately $10^{1.8}$ and $10^{1.7}$, respectively, while the block Jacobi resolvent norm is approximately, $10^{0.0}$. For $M = 0.5$, all of the preconditioners appear fairly well behaved. However, for $M = 0.99$, the van Leer preconditioned system appears to be deviating further from normality with the resolvent norm at $z = (1, 0)$ approximately $10^{0.7}$.

These general trends can also be seen in the maximum rate of energy amplification, dG/dt, at $t = 0^+$ which is tabulated in Table I. At low Mach numbers, the van Leer and Turkel preconditioners allow substantial growth rates. In fact, our results indicate that the maximum growth rate for the van Leer preconditioned system behaves as $1/M$, while the two Turkel preconditioned systems behave as $1/(2M)$ for $M \to 0$. Regardless, significant transient growth is likely to exist as $M \to 0$ for both the van Leer and Turkel preconditioners. However, as we discuss in Section 5, this potential for transient growth is more likely to impact robustness for stagnation points rather than for nearly incompressible flow calculations.

4. NUMERICAL RESULTS

In this section, we report numerical results which support the previous transient growth analysis. Specifically, we use a first-order grid-aligned upwind scheme to solve the linearized preconditioned equations on a Cartesian grid with constant Δx and Δy. For example, the x-derivative terms at cell (j, k) are approximated as

$$\left(\mathbf{A}\frac{\partial \mathbf{U}}{\partial x}\right)_{j,k} = \frac{1}{\Delta x}[\mathbf{F}_{j+1/2,k} - \mathbf{F}_{j-1/2,k}],$$

where the fluxes are evaluated using Roe's flux-difference splitting [22],

$$\mathbf{F}_{j+1/2,k} = \tfrac{1}{2}(\mathbf{F}_{j,k} + \mathbf{F}_{j+1,k}) - \tfrac{1}{2}|\mathbf{A}|(\mathbf{U}_{j+1,k} - \mathbf{U}_{j,k}). \quad (11)$$

For this linear problem, the flux vector, $\mathbf{F}_{j,k} = \mathbf{A}\mathbf{U}_{j,k}$. In conjunction with the van Leer and Turkel preconditioners, a modification of the dissipation matrix is required [3, 9]. The modified form of Eq. (11) is

$$\mathbf{F}_{j+1/2,k} = \tfrac{1}{2}(\mathbf{F}_{j,k} + \mathbf{F}_{j+1,k}) - \tfrac{1}{2}\mathbf{P}^{-1}|\mathbf{P}\mathbf{A}|(\mathbf{U}_{j+1,k} - \mathbf{U}_{j,k}). \quad (12)$$

For the block Jacobi preconditioner, a straightforward Jacobi-like algorithm is implemented as suggested by Allmaras [6] and does not require any modification of the dissipation matrix. In either case, the maximum timestep is defined as $\Delta t_{\max}^{-1} = c_x/\Delta x + c_y/\Delta y$, where c_x and c_y are the magnitudes of the largest characteristic speeds in the x and y directions. For the Jacobi scheme, the characteristic speeds are still based on the original Euler equations; for the other schemes, the characteristic speeds are based on the preconditioned equations. Then, the actual timestep for the calculations is defined as $\Delta t = \nu\, \Delta t_{\max}$, where ν is the CFL number. For these linear results, the time integration is approximated using forward Euler:

$$\mathbf{U}_{j,k}^{n+1} = \mathbf{U}_{j,k}^{n} - \Delta t\mathbf{P}\left[\left(\mathbf{A}\frac{\partial \mathbf{U}}{\partial x}\right)_{j,k}^{n} + \left(\mathbf{B}\frac{\partial \mathbf{U}}{\partial y}\right)_{j,k}^{n}\right].$$

The CFL number is set to $\nu = 0.5$ which provides good high-frequency damping properties.

As with the resolvent contour plots, three mean flow Mach numbers are run, $M = 0.01$, 0.5, and 0.99. The grid has 17 points in both coordinate directions with unit cell aspect ratio ($\Delta x = \Delta y$). The initial flow was composed of random perturbations containing both high and low frequencies. The flow angle was also randomly set. 100 different initial conditions were run for 250 iterations using each preconditioner. For boundary nodes, a ghost cell approach was used with the ghost state set to zero (i.e., no farfield perturbations).

The convergence histories for all combinations are shown in Figs. 3 and 4. These plots show the average, minimum, and maximum norm of the symmetrized perturbation states for the 100 different random initial conditions which were tested. First, we compare the results for all preconditioners at $M = 0.01$. The transient growth for the van Leer and Turkel preconditioned systems is clear; within five iterations, these preconditioned systems all amplify the perturbations over an order of magnitude. The striking result is that the minimum, average, and maximum norms of the perturbations are nearly the same for each preconditioner; thus, essentially all of the random conditions tested suffered the same transient growth. By comparison, the block Jacobi preconditioner does not have any transient growth effects. However, although the van

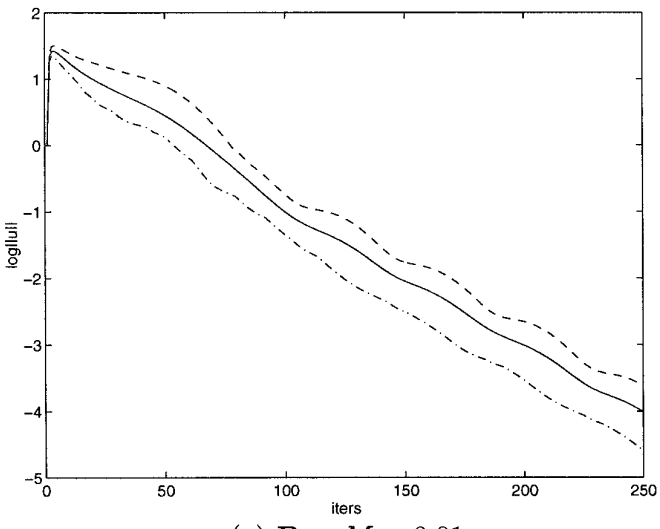

(a) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.01$

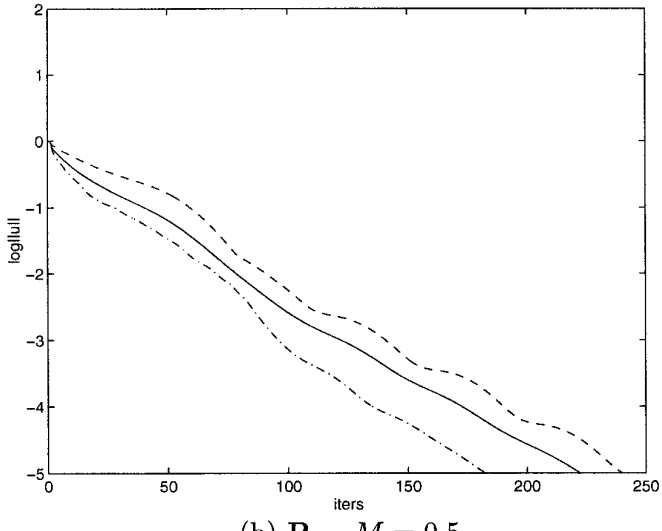

(b) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.5$

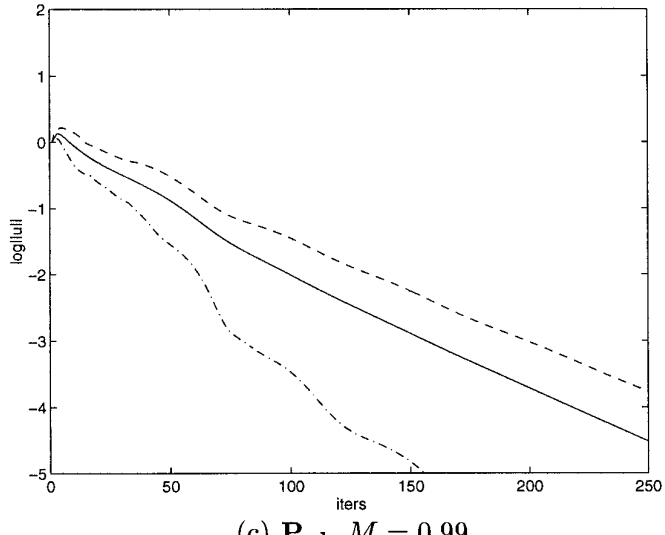

(c) $\mathbf{P}_{\mathrm{vl}}$, $M = 0.99$

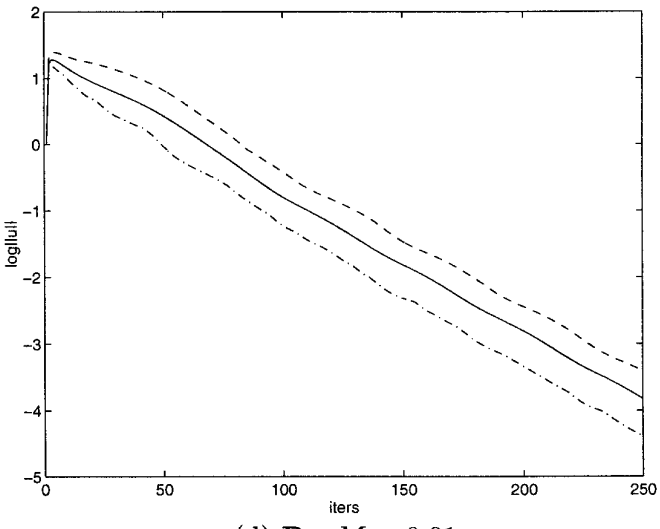

(d) $\mathbf{P}_{\mathrm{t}}$, $M = 0.01$

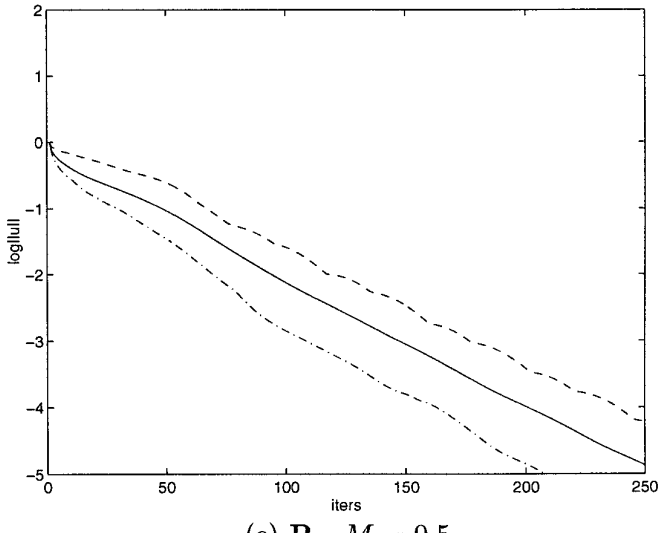

(e) $\mathbf{P}_{\mathrm{t}}$, $M = 0.5$

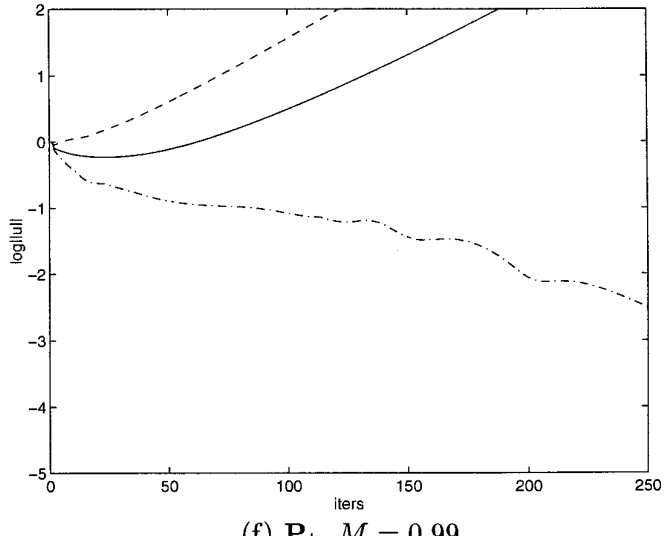

(f) $\mathbf{P}_{\mathrm{t}}$, $M = 0.99$

FIG. 3. Average, minimum, and maximum residual history for linearized preconditioned Euler equations, van Leer, and Turkel preconditioners. First-order upwind; 100 random initial conditions; farfield boundary condition; 17 by 17 grid.

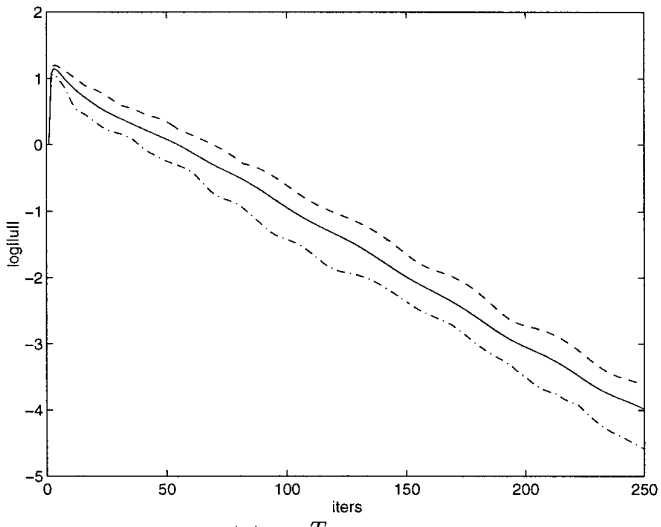

(a) $\mathbf{P}_{\mathrm{t}}{}^{T}$, $M = 0.01$

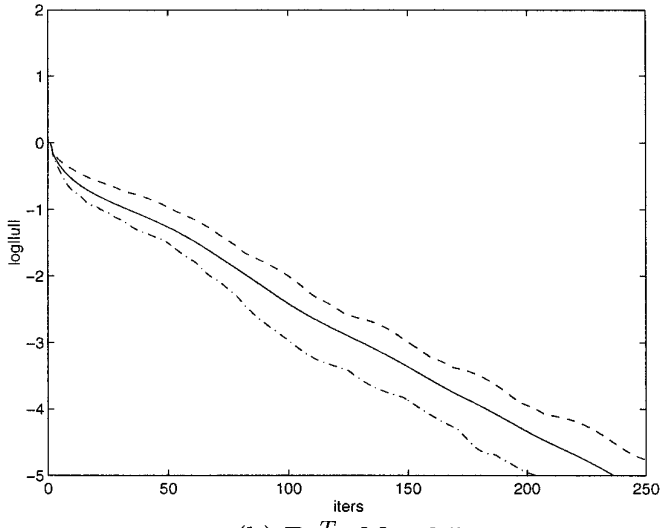

(b) $\mathbf{P}_{\mathrm{t}}{}^{T}$, $M = 0.5$

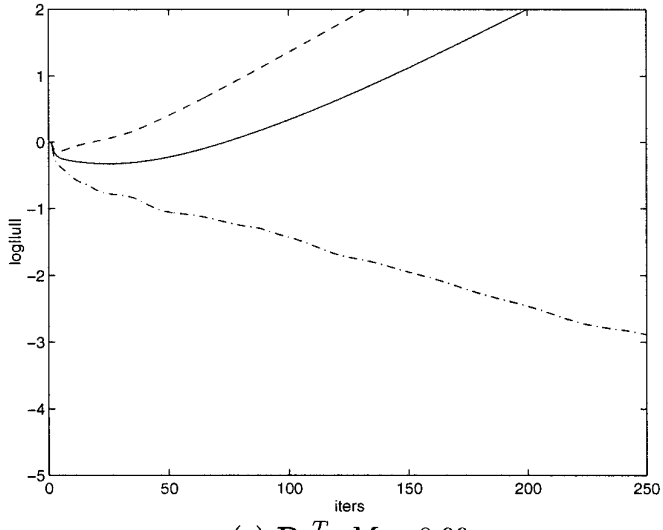

(c) $\mathbf{P}_{\mathrm{t}}{}^{T}$, $M = 0.99$

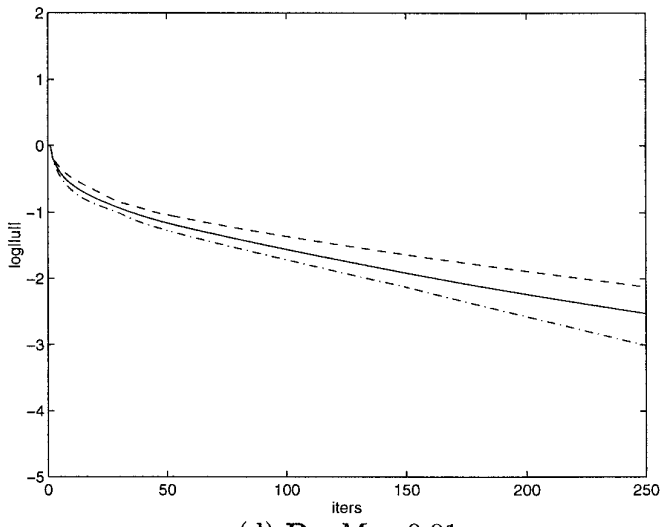

(d) $\mathbf{P}_{\mathrm{j}}$, $M = 0.01$

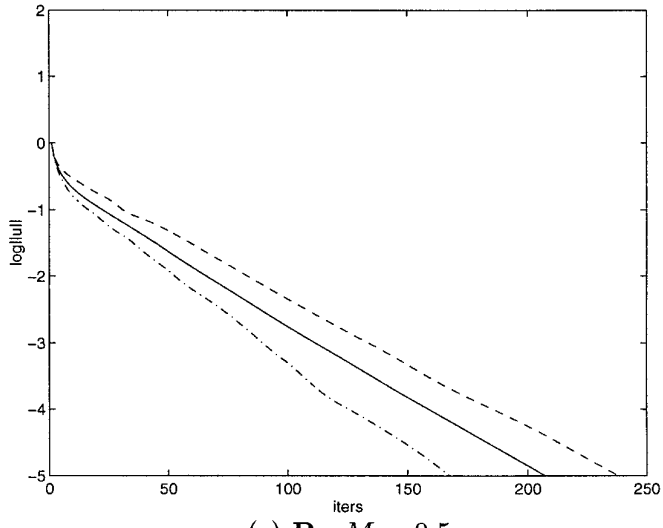

(e) $\mathbf{P}_{\mathrm{j}}$, $M = 0.5$

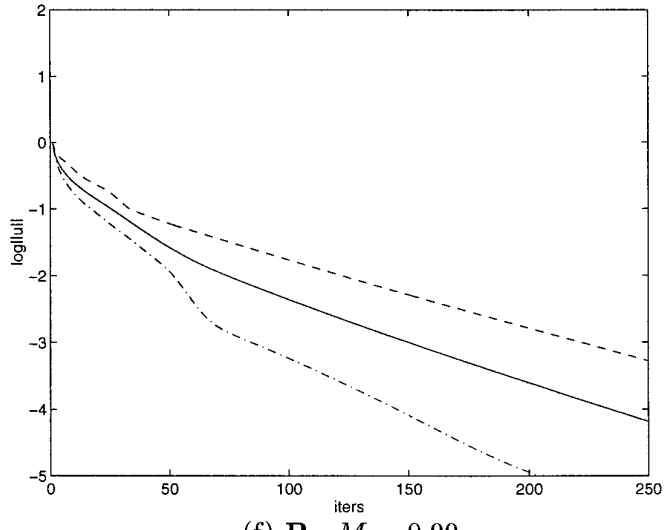

(f) $\mathbf{P}_{\mathrm{j}}$, $M = 0.99$

FIG. 4. Average, minimum, and maximum residual history for linearized preconditioned Euler equations, Turkel transpose, and block Jacobi preconditioners. First-order upwind; 100 random initial conditions; farfield boundary condition; 17 by 17 grid.

Leer and Turkel preconditioners have significant initial growth rates, these preconditioners have better asymptotic convergence rates than the block Jacobi. The asymptotic rates are better for the van Leer and Turkel preconditioners because they were specifically designed to accelerate low frequency disturbances; however, the block Jacobi preconditioner, which was designed for effective high frequency damping, has little beneficial effect on the propagation speeds of low frequency disturbances. In fact, after 250 iterations, the van Leer and Turkel preconditioned schemes still converge further than the block Jacobi scheme, even though the former preconditioners have large initial transients at $M = 0.01$. These results demonstrate how the eigenvectors play a significant role in determining the initial stages of a calculation while the eigenvalues are significant in the asymptotic portion of the convergence.

The $M = 0.5$ convergence histories for all of the preconditioners show a similar behavior. No transient growth is observed for any of the results at this Mach number. For $M = 0.99$, however, some small transient effects can be seen for the van Leer preconditioner as predicted from the previous analysis. Another interesting point is that the upwind scheme applied to the Turkel preconditioners results in eigenvalues of the discretized spatial operator which lie in the unstable, right half plane for $M = 0.99$ (note, the discrete eigenvalues are not shown here). Thus, the long-time behavior of these solutions should exhibit an instability. This instability can be seen in Figs. 3 and 4 for the Turkel preconditioners at $M = 0.99$. We have also observed this instability for nonlinear simulations of the Turkel preconditioned Euler equations for transonic flow when employing an upwind difference method.

5. ROBUST PRECONDITIONING

As we have shown, poor eigenvector conditioning can lead to transient growth in both the continuous and discretized linearized equations. For linear equations, this poses no particular problem since the magnitude of the perturbation quantities does not alter the local mean flow state. However, this transient growth can cause significant problems for nonlinear calculations where the local mean flow is altered by the perturbations. If the potential for transient amplification is significant, the calculation may slowly converge as transient effects are continuously stimulated by incoming waves, or worse, the calculation may abort if the perturbations result in a non-physical mean state (i.e., negative pressures or densities).

If one wishes to avoid transient effects, a possible solution would be to use the block Jacobi preconditioner. However, as noted in Section 4, the block Jacobi preconditioner does not accelerate the long wavelength modes as $M \to 0$. Furthermore, the block Jacobi preconditioner does not improve the solution quality for nearly incompressible flows since the dissipation is not modified with this approach. Thus, if solutions for low freestream Mach numbers are desired, the block Jacobi preconditioner is not an adequate choice. This suggests modifying existing preconditioners or designing new preconditioners to avoid transient amplification effects while maintaining other beneficial properties. In future work, we will investigate this design process in detail; however, in this paper, we illustrate one possible solution for the Turkel preconditioner and briefly describe how we arrive at this solution.

For symmetric preconditioners of the symmetrized variables, a bound on the condition number of the eigenvector matrix can be formulated. Any symmetric, positive definite matrix can be written as

$$\mathbf{P}_{\text{sym}} = \mathbf{Q}\mathbf{Q}^{\mathrm{T}},$$

giving the preconditioned governing equations

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{Q}\mathbf{Q}^{\mathrm{T}}\mathbf{A}\frac{\partial \mathbf{U}}{\partial x} + \mathbf{Q}\mathbf{Q}^{\mathrm{T}}\mathbf{B}\frac{\partial \mathbf{U}}{\partial y} = 0. \tag{13}$$

Applying the similarity transformation, $\mathbf{V} = \mathbf{Q}^{-1}\mathbf{U}$, gives

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{Q}^{\mathrm{T}}\mathbf{A}\mathbf{Q}\frac{\partial \mathbf{V}}{\partial x} + \mathbf{Q}^{\mathrm{T}}\mathbf{B}\mathbf{Q}\frac{\partial \mathbf{V}}{\partial y} = 0,$$

which is a symmetric system (and therefore has orthogonal eigenvectors) since $\mathbf{A}$ and $\mathbf{B}$ are symmetric. Thus, using Eq. (10), the eigenvector condition number, $\kappa(\mathbf{R})$, of Eq. (13) can be bound by the condition number of the similarity transformation, $\kappa(\mathbf{Q})$. Finally, using the definition of the condition number, it is easy to show that

$$\kappa(\mathbf{R}) \le \kappa(\mathbf{Q}) = \sqrt{\lambda_{\max}(\mathbf{P}_{\text{sym}})/\lambda_{\min}(\mathbf{P}_{\text{sym}})}, \tag{14}$$

where $\lambda_{\max}$ and $\lambda_{\min}$ are the maximum and minimum eigenvalues of $\mathbf{P}_{\text{sym}}$. Thus, one method for limiting transient amplification is to bound the condition number of the preconditioning matrix. For a diagonal preconditioner of the symmetrized variables, the eigenvector condition number is bound by the square root of the ratio of the largest to smallest elements of the preconditioner. Note that, although Eq. (14) indicates that a singular preconditioning matrix (i.e., where $\lambda_{\min}/\lambda_{\max} \to 0$) could lead to poorly conditioned eigenvectors, this result is only an upper bound on $\kappa(\mathbf{R})$. For example, the block Jacobi preconditioner, $\mathbf{P}_{\text{j}}$, becomes singular as $M \to 0$; however, as shown in the previous sections, no transient growth occurs when using $\mathbf{P}_{\text{j}}$ at $M = 0.01$. Thus, forcing $\mathbf{P}_{\text{sym}}$ to be well-conditioned is likely to be an overly severe method to limit transient growth.

Applying these results to the Turkel preconditioner, we first modify the preconditioner definition to

$$\mathbf{P}_t = \frac{1}{M}\begin{pmatrix} \beta_t^2 & 0 & 0 & 0 \\ -\alpha_t M & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

We let $\alpha_t = 1 + \beta_t^2$ which results in symmetric acoustic wavefronts [2]. In particular, the eigenvalues for this preconditioned system are

$$(\omega/k)_{1,2} = a\cos\phi,$$

$$(\omega/k)_{3,4} = \pm a\frac{\beta_t}{M}\sqrt{1 - M^2\cos^2\phi}.$$

The first pair of eigenvalues corresponds to the convective waves of entropy and vorticity. The second pair represents the acoustic modes. If β_t is proportional to the Mach number, then all of the eigenvalues scale with the speed of sound. However, as $M \to 0$, the Turkel preconditioner approaches a diagonal matrix; thus, from the previous result, we must bound β_t away from zero to limit the transient growth. One possible option is to define

$$\beta_t = \max(M, \eta M_\infty),$$

where M_∞ is the freestream Mach number and η is a free parameter. Unfortunately, the acoustic eigenvalues are now proportional to M_∞/M which becomes unbounded at stagnation points. Thus, while the Mach number cutoff bounds the transient growth, it has the negative effect of reducing the eigenvalue conditioning. We note that this modification to the Turkel preconditioner is not unique and that a very similar suggestion is used by Turkel *et al.* [10] to improve the robustness of a central difference code. The same type of limiting is also performed by Choi and Merkle [1, 7]. However, in these references, the exact connection of transient growth and eigenvector conditioning was not pursued.

With this suggested modification, the transient growth of the symmetrized variables at $M = 0$ is bound by $(\eta M_\infty)^{-1}$. From this bound, one might expect that low Mach number calculations would still experience significant robustness problems unless η scales as M_∞^{-1}. A more thorough analysis (as performed in the Appendix for the 1D van Leer preconditioner) shows that the transient growth at $M = 0$ is due to the creation of velocity disturbances from pressure disturbances which are amplified by a factor of $(\eta M_\infty)^{-1}$. However, for low Mach number flows with no large incoming pressure perturbations, one can generally expect very small pressure perturbations such that $\tilde{p}/\rho_\infty a_\infty^2$ is at most of the order M_∞^2 as $M_\infty \to 0$ [23]. Thus, although large transient growth is still possible at low freestream Mach numbers, the pressure perturbations are small such that the resulting velocity perturbations, $\tilde{u}/U_\infty$, are bounded by η^{-1} as $M \to 0$.

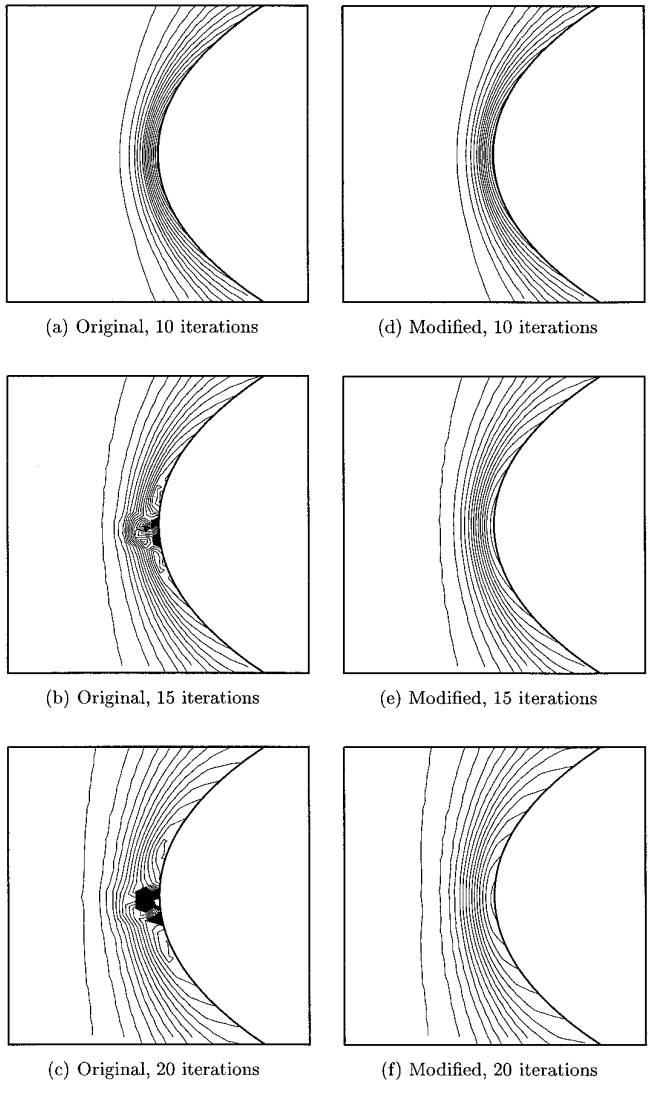

FIG. 5. Comparison of Mach number contours for original and modified Turkel preconditioners at leading edge after 10, 15, and 20 iterations. $M_\infty = 0.01$, 2 degrees angle of attack, NACA 0012. 31 equally spaced contours from $M = 0$ to 0.02.

Finally, we illustrate the effect of this modification with results from an unstructured Euler solver. The solver is a locally preconditioned version of the Barth and Jespersen node-based upwind algorithm [24]. Only first-order accurate results are presented in this paper. The time marching scheme is an optimally smoothing multi-stage scheme for unstructured grids which has been recently developed by

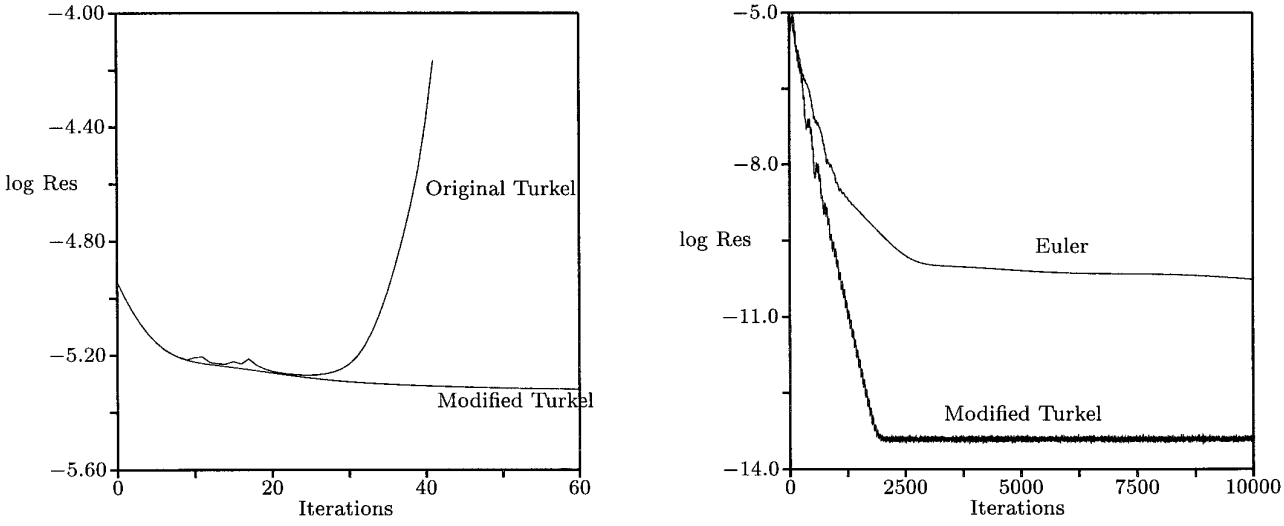

FIG. 6. Convergence history for NACA 0012, $M = 0.01$, 2 degrees angle of attack at two CFL numbers.

the first author in conjunction with John Lynn [25]. Specifically,

$$\mathbf{U}_i^{(1)} = \mathbf{U}_i^n - \alpha_1 \Delta t_i \mathbf{P}_i \mathbf{R}_i^{(0)},$$
$$\mathbf{U}_i^{(2)} = \mathbf{U}_i^n - \alpha_2 \Delta t_i \mathbf{P}_i \mathbf{R}_i^{(1)},$$
$$\mathbf{U}_i^{(3)} = \mathbf{U}_i^n - \alpha_3 \Delta t_i \mathbf{P}_i \mathbf{R}_i^{(2)},$$
$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n - \alpha_4 \Delta t_i \mathbf{P}_i \mathbf{R}_i^{(3)},$$

where $\alpha_1 = 0.07051$, $\alpha_2 = 0.1803$, $\alpha_3 = 0.3854$, $\alpha_4 = 1.0$, and $\Delta t_i = \nu \, \Delta t_{\max i}$. Note that, to accelerate convergence further, the timestep varies with the node. For this time integration, the optimal CFL is $\nu_{\text{opt}} = 3.1$ with $\Delta t_{\max i}$ defined as

$$\frac{A_i}{\Delta t_{\max i}} = \sum_j^{N_i} \Delta s_j \lambda_j^+,$$

where A_i is the area of the control volume for node i, N_i is the number of nodes neighboring i, Δs_j is the length of the dual edge associated with nodes i and j, and λ_j^+ is the largest positive eigenvalue normal to the dual edge. This timestep was derived by Barth and Jespersen by enforcing monotonicity [24]. We solve the flow over a NACA 0012 airfoil at 2 degrees angle of attack and a freestream Mach number, $M_\infty = 0.01$. The grid has 3021 nodes with 130 nodes on the airfoil surface. For this calculation, we set $\eta = 0.5$. Note, we have used $\eta = 0.5$ for a variety of freestream Mach numbers, flow angles, and grids with the first-order scheme and never experienced any lack of robustness with all calculations converging to well-behaved steady solutions.

Solutions were attempted for two CFL numbers, $\nu = \nu_{\text{opt}}/5$ and ν_{opt}. Mach number contours for the original and modified Turkel preconditioners after 10, 15, and 20 iterations using the smaller CFL are shown in Fig. 5. The Mach number contours from Fig. 5 show significant perturbation growth at the stagnation point for the original Turkel preconditioner. From our experience with this unstructured code, this type of stagnation point disturbance is typical of not only the original Turkel but also the original van Leer preconditioner. As expected, the Mach contours for the modified Turkel scheme show no evidence of the stagnation point problem. Convergence histories for both CFL numbers are shown in Fig. 6. At the smaller CFL, the original Turkel preconditioner eventually aborts after 42 iterations while the modified preconditioner is smoothly varying. For the optimal CFL, the original Turkel preconditioner aborts after a single iteration from negative pressure and density at the leading edge while the modified preconditioner converges to machine zero in about 2000 iterations. For comparison, the unpreconditioned Euler convergence history is also shown in Fig. 6b and after 10000 iterations is still far from machine zero. Mach number contours for the preconditioned results at convergence and the unpreconditioned Euler results at 10000 iterations are shown in Fig. 7. As observed in previous efforts, the improved solution accuracy for preconditioned low Mach number calculations is very impressive [3, 9, 5]. Thus, although the unpreconditioned Euler residual drops four

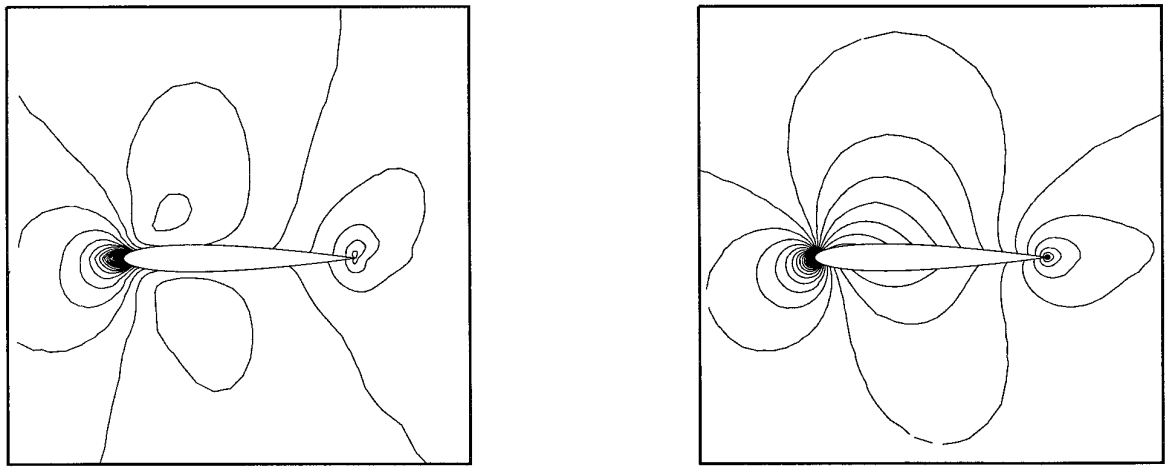

(a) Unpreconditioned after 10000 iterations

(b) Converged preconditioned results

FIG. 7. Comparison of M contours for modified Turkel preconditioned results and unpreconditioned Euler results. $M_\infty = 0.01$, 2 degrees angle of attack, NACA 0012. 51 equally spaced contours from $M = 0$ to 0.015.

orders of magnitude, the quality of the solution is still very poor.

6. CONCLUSIONS

In this paper we have demonstrated the influence of eigenvectors on the convergence dynamics of the locally preconditioned Euler equations. The non-orthogonality of the eigenvectors of the preconditioned system has been found responsible for a transient amplification of the residual and a subsequent loss of robustness. A number of mathematical tools to analyze preconditioners have been introduced, and this mathematical framework has been applied to a number of standard local preconditioners. We found that the van Leer and Turkel preconditioned systems are highly non-normal for low Mach numbers. This non-normality was shown numerically to result in a large amplification of the residual norm as $M \to 0$. In particular, this non-normality resulted in a loss of robustness near stagnation points for non-linear simulations. A modification to the van Leer and Turkel preconditioners was suggested which limits the departure of the preconditioned system from normality. This modification was demonstrated to result in a significant improvement in robustness for an unstructured, upwind finite volume code.

The only preconditioner analyzed which did not suffer from non-normality at low Mach numbers was the block Jacobi preconditioner. However, this preconditioner does not effectively equalize the characteristic speeds of low frequency waves; thus, although the block Jacobi did not suffer from any initial transient growth, its asymptotic convergence rate was generally less than that of the van Leer and Turkel preconditioners.

APPENDIX

As an aid to the reader unfamiliar with the mathematical tools employed in this paper, a simple 2×2 model problem will be discussed that illustrates the need for non-modal analysis and demonstrates the mathematical framework to quantify transient effects.

We will consider the one-dimensional linearized Euler equations in symmetric variables $\mathbf{U} = (\tilde{p}/\rho a, \tilde{u})$, where $\tilde{p}$ and $\tilde{u}$ are the perturbation pressure and velocity, and ρ and a are the mean density and the speed of sound. For simplicity, we have not included the entropy in these variables because it fully decouples from the pressure and velocity perturbations. After non-dimensionalization by the mean speed of sound, the linearized Euler equations read

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{A}\frac{\partial \mathbf{U}}{\partial x} = 0 \quad \text{with } \mathbf{A} = \begin{pmatrix} M & 1 \\ 1 & M \end{pmatrix},$$

where M is the Mach number. We will use the one-dimensional version of the van Leer preconditioner given by

$$\mathbf{P}_{\text{vl}} = \frac{1}{M}\begin{pmatrix} \dfrac{M^2}{\beta^2} & -\dfrac{M}{\beta^2} \\ -\dfrac{M}{\beta^2} & \dfrac{1+\beta^2}{\beta^2} \end{pmatrix}, \quad \beta^2 = 1 - M^2,$$

to explore the occurrence of transient effects in the preconditioned system,

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{P}_{\mathrm{vl}}\mathbf{A}\frac{\partial \mathbf{U}}{\partial x} = 0.$$

After Fourier transforming in the streamwise direction, we obtain the evolution equation,

$$\frac{d\mathbf{u}}{dt} + i\mathbf{L}\mathbf{u} = 0 \quad \text{with } \mathbf{L} = k\begin{pmatrix} -1 & 0 \\ \dfrac{2}{M} & 1 \end{pmatrix},$$

where k is the wavenumber of the perturbation. Note, as in the two-dimensional analysis, the wave number only scales the time derivative; thus, we assume $k = 1$ for the rest of the analysis. It is straightforward to determine the eigenvalues and normalized eigenvectors of the evolution matrix $\mathbf{L}$ as

$$\omega_1 = 1, \qquad \mathbf{r}_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

$$\omega_2 = -1, \quad \mathbf{r}_2 = \frac{1}{\sqrt{1+M^2}}\begin{pmatrix} -M \\ 1 \end{pmatrix}$$

and the complete solution expanded in an eigenvector basis can thus be written as

$$\mathbf{u}(t) = v_1\begin{pmatrix} 0 \\ 1 \end{pmatrix}e^{it} + \frac{v_2}{\sqrt{1+M^2}}\begin{pmatrix} -M \\ 1 \end{pmatrix}e^{-it}.$$

The constants v_1 and v_2 are determined from the initial conditions.

These eigenvectors are clearly non-orthogonal as $M \to 0$. The angle ϕ between the eigenvectors can be determined as

$$\phi = \arccos\left(\frac{1}{\sqrt{1+M^2}}\right)$$

which rapidly approaches zero as the Mach number decreases, thus leading to an increasingly ill-conditioned expansion basis. The possibility for large transients as $M \to 0$ can be shown if one considers an initial condition with a finite pressure perturbation but zero velocity perturbation, for example, $\mathbf{u}_0 = (\tilde{p}_0/\rho a, 0)$. In order to match the initial pressure perturbation, the strength of the second wave must be $v_2 = (-\sqrt{1+M^2}/M)(\tilde{p}_0/\rho a)$. For zero initial velocity perturbation, the strength of the first wave is then $v_1 = (1/M)(\tilde{p}_0/\rho a)$. Thus, for low Mach numbers, the strengths of the eigenmodes can become extremely large. At $t = 0$, the large wave strengths are not observed because of cancellation. However, as time progresses, the waves propagate in different directions and the fragile cancellation no longer occurs, resulting in a significant growth in the velocity.

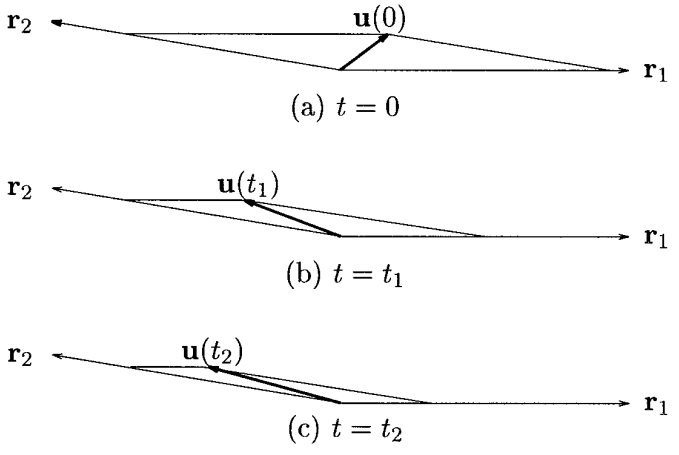

FIG. 8. Example of a 2×2 system with transient growth resulting from non-orthogonal eigenvectors for $t_2 > t_1 > 0$.

This effect can be visualized considering a 2×2 system with highly non-orthogonal eigenvectors as depicted in Fig. 8. The initial condition, $\mathbf{u}_0$, although of small magnitude, requires large expansion coefficients v_j due to the non-orthogonality of the eigenvectors. At the initial instant, the large amplitude eigenvectors are not evident because of their mutual cancellation. Suppose that the eigenvalues of the associated eigenvectors are distinct, negative imaginary numbers, then, as time advances, the amplitude of the eigenmodes will decrease at different rates. As a result of the different decay rates, the initial cancellation of the eigenvectors does not remain and large transient growth will be observed in the state phase space. This growth can be seen in Fig. 8 as the increase in length of the state vector, $\mathbf{u}$. For purely real eigenvalues (as for the Euler equations), the state space diagram is more difficult to conceptualize. In this case, the initial eigenmode cancellation disappears because of different propagation speeds, not different decay rates. Regardless, the result is the same; for non-orthogonal eigenvectors, a system with real eigenvalues can experience transient energy amplification.

Another way of expressing the solution is by using the matrix exponential, $\mathbf{u}(t) = \exp(-it\mathbf{L})\mathbf{u}_0$, with $\mathbf{u}_0$ as the initial state vector. As described in Section 2, the L_2-norm of the matrix exponential is the maximum growth, $G(t)$, as a function of time and optimized over all unit norm initial conditions. For the 1D van Leer preconditioned equations, $G(t)$ is plotted in Fig. 9 for selected values of the Mach number M. As can be seen, substantial transient growth can result as $M \to 0$.

It is also instructive to consider the evolution equation for the disturbance norm, Eq. (8). From this, the initial amplification rate optimized over all unit norm initial conditions is

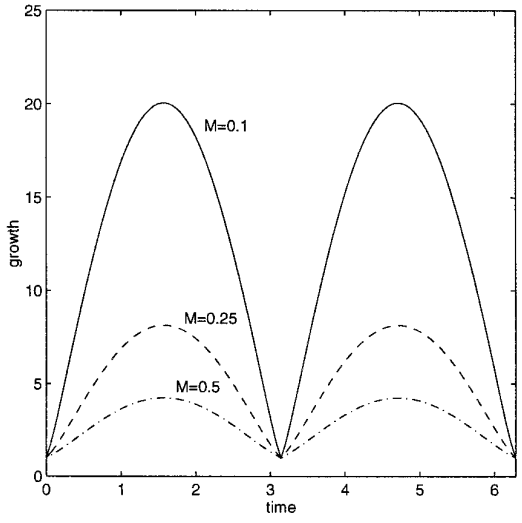

FIG. 9. $G(t)$ for 1D van Leer preconditioned Euler equations.

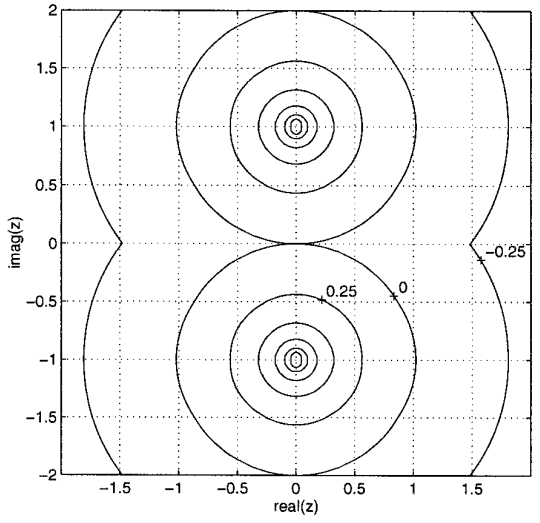

FIG. 11. Contours of $\log_{10} R(z)$ for 1D block Jacobi preconditioned Euler equations at $M = 0.01$.

$$\left.\frac{dG}{dt}\right|_{t=0+} = \frac{1}{M}.$$

Thus, arbitrarily high amplification rates can be achieved for low Mach numbers.

As pointed out in the main text, the resolvent norm, $R(z)$, can be used to estimate the maximum transient growth. For normal evolution matrices, the resolvent norm constitutes an inverse distance function from the spectrum. For highly non-normal matrices, the resolvent norm can be considerably larger than the reciprocal distance from the spectrum. Thus, contours of the resolvent norm in the complex z-plane can be used as an indication of non-normality. As a simple illustration, we plot the resolvent norm for the 1D van Leer and block Jacobi preconditioned systems at $M = 0.01$ in Figs. 10 and 11. Note, the block Jacobi preconditioner is

$$\mathbf{P}_{\mathrm{j}} = \frac{1}{M}\begin{pmatrix} \dfrac{M}{\beta^2} & -\dfrac{M^2}{\beta^2} \\ -\dfrac{M^2}{\beta^2} & \dfrac{M}{\beta^2} \end{pmatrix}.$$

The block Jacobi preconditioned system has not only the same eigenvalues as the van Leer system but also orthogonal eigenvectors. It is easily seen that the resolvent norm for the van Leer system is two orders of magnitude larger than the block Jacobi system. In contrast, the block Jacobi resolvent contours are circles centered at the eigenvalues and equal to the inverse distance from the closest eigenvalue. This resolvent behavior is typical of normal operators. Thus, we expect transient amplification for the van Leer system of approximately two orders of magnitudes while the block Jacobi should not suffer any transient amplification.

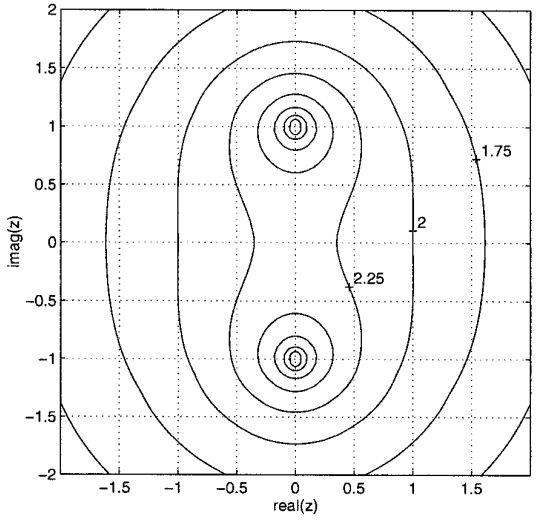

FIG. 10. Contours of $\log_{10} R(z)$ for 1D van Leer preconditioned Euler equations at $M = 0.01$.

ACKNOWLEDGMENTS

Our sincere appreciation to Steve Allmaras, Mike Giles, Bram van Leer, Lisa Mesaros, and Phil Roe, who have provided very helpful suggestions, comments, and information during the course of this work. The majority of this work was completed while the first author was at the W. M. Keck Foundation Computational Fluid Dynamics Laboratory at the University of Michigan funded by an NSF PostDoctoral Research Associateship in Computational Science and Engineering and AFOSR

Augmentation Grant No. F49620-93-1-0417. The second author gratefully acknowledges support from the National Science Foundation under Grant DMS-9406636.

REFERENCES

1. Y. H. Choi and C. L. Merkle, *AIAA J.* **23**(10), 1518 (1985).
2. E. Turkel, *J. Comput. Phys.* **72,** 277 (1987).
3. B. van Leer, W. T. Lee, and P. L. Roe, AIAA Paper 91-1552, 1991 (unpublished).
4. D. Lee and B. van Leer, AIAA Paper 93-3328, 1993 (unpublished).
5. A. C. Godfrey, R. W. Walters, and B. van Leer, AIAA Paper 93-0535, 1993 (unpublished).
6. S. R. Allmaras, AIAA Paper 93-3330, 1993 (unpublished).
7. Y. H. Choi and C. L. Merkle, *J. Comput. Phys.* **105,** 203 (1993).
8. E. Turkel, ICASE Report No. 92-47, 1992 (unpublished).
9. W. T. Lee, Ph.D. thesis, University of Michigan, 1991 (unpublished).
10. E. Turkel, A. Fiterman, and B. van Leer, ''Preconditioning and the Limit to the Incompressible Flow Equations for Finite Difference Schemes,'' in *Computing the Future: Advances and Prospects for Computational Aerodynamics*, edited by M. Hafez and D. A. Caughey, p. 215. (Wiley, New York, 1994).
11. K. M. Butler and B. F. Farrell, *Phys. Fluids A* **4**(8), 1637 (1992).
12. L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. Driscoll, *Science* **261,** 578 (1993).
13. P. J. Schmid, Ph.D. thesis, M.I.T., 1993 unpublished.
14. S. C. Reddy and D. S. Henningson, *J. Fluid Mech.*, in press.
15. S. C. Reddy and L. N. Trefethen, *Numer. Math.* **62,** 235 (1992).
16. H. W. J. Lenferink and M. N. Spijker, *Math. Comput.* **57,** 221 (1991).
17. C. Lubich and O. Nevanlinna, *BIT* **31,** 293 (1991).
18. M. B. Giles, AIAA Paper No. 95-1753, 1995 (unpublished).
19. J. L. M. van Dorsselaer, J. F. B. Kraaijevanger, and M. N. Spijker, *Acta Numer.*, p. 199 (1993).
20. B. van Leer, L. Mesaros, C. H. Tai, and E. Turkel, AIAA Paper 95-1654, 1995 (unpublished).
21. T. Kato, *Perturbation Theory for Linear Operators* (Springer-Verlag, 1976).
22. P. L. Roe, *J. Comput. Phys.* **43,** 357 (1981).
23. B. Gustafsson and H. Stoor, *SIAM J. Numer. Anal.* **28**(6), 1523 (1991).
24. T. J. Barth and D. C. Jespersen, AIAA Paper 89-0366, 1989 (unpublished).
25. J. Lynn, Ph.D. thesis, University of Michigan, 1995 (unpublished).