PRECONDITIONING TECHNIQUES IN FLUID DYNAMICS

E. Turkel *
School of Mathematical Sciences
Sackler Faculty of Exact Sciences
Tel-Aviv University
Tel-Aviv, Israel
and
ICASE, NASA Langley

*Paper prepared for the Annual Review of Fluid Mechanics, Volume 31, 1999

Contents

1 Elliptic equations

The concept of preconditioning was used for solving elliptic equations with a conjugate gradient iterative procedure. Consider the equation $\mathcal{L}u = f$. After discretizing we have the algebraic system

$$Av = b. \tag{1}$$

where A is positive definite. We define the condition number of A as $\kappa(A) = \max \frac{|\lambda_i|}{|\lambda_j|}$ over all the eigenvalues λ of A. The convergence rate of many schemes depends on the condition number. One way to accelerate the scheme is to replace (1) by

$$\mathbf{P}Av = \mathbf{P}b. \tag{2}$$

As long as $\mathbf{P}$ is nonsingular the solutions of (2) are the same as (1). The preconditioning is introduced only as a mechanism to accelerate the iteration procedure. For the preconditioning to be effective we require that $\kappa(\mathbf{P}A) \ll \kappa(A)$ and that $\mathbf{P}$ be easy to invert.

We consider several classes of preconditioners.

I. $\mathbf{P}$ is a diagonal matrix. In this case the preconditioner only affects variables at a single mesh point.

   If the original set (1) is derived from a system of equations then $\mathbf{P}$ can be a block diagonal matrix with the size of the block corresponding to the number of equations in the system. Thus, the preconditioning changes the equations within the system but does not couple neighboring points. One special choice is $P_{ii} = A_{ii}^{-1}$. This scales the new system so that the identity matrix appears on the diagonal of $\mathbf{P}A$. We shall refer to this as Jacobi preconditioning since it corresponds to introducing a Jacobi iteration.

II. Either $\mathbf{P}$ or $\mathbf{P}^{-1}$ has non-zero elements in the same positions as the non-zero elements of A. This is called a differential preconditioner since we can interpret $\mathbf{P}$ or $\mathbf{P}^{-1}$ as a differential operator.

III. $\mathbf{P}$ does not satisfy either of the above rules. This is referred to as a global preconditioner.

One frequently constructs preconditioners from various viewpoints

- The preconditioner is based on the inverse of a simpler differential equation. Thus, for example, if $\mathcal{L}$ is the Helmholtz operator $\Delta + k^2$ then one could choose $\mathbf{P} = \Delta^{-1}$ which could be accomplished numerically by a multigrid sweep. This approach leads to a diagonal or differential preconditioner.

- The preconditioner is based on strictly algebraic properties of the matrix A without accounting for the connection between A and a differential equation. This approach usually leads to a global preconditioner.

2 Steady-State Conservation Laws

We shall concentrate on solutions to the two dimensional steady state Euler and Navier-Stokes equations. Extension to three dimensions is straightforward. For time dependent problems these methods can be used if the scale of the time dependence is slow compared with the convective speed of the fluid. In that case the time dependent equations can be solved as a sequence of modified steady state problems.

The steady state equations written as a system of conservation laws can be expressed as

$$f_x + g_y = 0 \quad .$$

We approximate the flux derivatives by central differences. It is well known that one needs an artificial viscosity in order to stabilize the scheme. This viscosity consists of a combination of second and fourth differences. The second order difference contains a nonlinear term and prevents overshoots in the neighborhood of shocks. The fourth order viscosity is required to prevent odd-even decoupling which prevents the iteration scheme from reaching a steady state [15]. A similar algorithm results using an upwind scheme based on the Roe linearization. To simplify the presentation we shall only describe a linear second difference viscosity term. The extension to nonlinearities and a fourth order viscosity is straightforward.

In a general coordinate system (ξ, η) the discretization is

$$\frac{F_{i+1/2,j} - F_{i-1/2,j}}{\Delta\xi} + \frac{G_{i,j+1/2} - G_{i,j-1/2}}{\Delta\eta} = d_{i+1/2,j} - d_{i-1/2,j} + d_{i,j+1/2} - d_{i,j-1/2} \tag{3}$$

Let $A = \frac{\partial F}{\partial w}$, $B = \frac{\partial G}{\partial w}$, then we choose

$$\begin{aligned} d_{i+1/2,j} &= \frac{1}{2}\left(|\sigma(A)|\right)_{i+1/2,j}(w_{i+1,j} - w_{i,j}) \\ d_{i,j+1/2} &= \frac{1}{2}\left(|\sigma(B)|\right)_{i,j+1/2}(w_{i,j+1} - w_{i,j}). \end{aligned} \tag{4}$$

$\sigma(A)$ is a function of the matrix A. A scalar viscosity [15] uses for σ the spectral radius of A times the identity matrix. The matrix viscosity [29] is based on the absolute value of the matrix A while the CUSP artificial viscosity [16] is equivalent to a linear function of the matrix A. The physical portion of (3) depends only on the fluxes and not on the state vector w. However, the artificial viscosity does depend on w.

The system (3) is a nonlinear system of partial differential equations of mixed type. Because of difficulties with the direct numerical solution it is common to replace (3) by a time dependent hyperbolic system and to march this to a steady state. So we consider

$$\frac{\partial\psi}{\partial t}+\frac{F_{i+1/2,j}-F_{i-1/2,j}}{\Delta\xi}+\frac{G_{i,j+1/2}-G_{i,j-1/2}}{\Delta\eta}=d_{i+1/2,j}-d_{i-1/2,j}+d_{i,j+1/2}-d_{i,j-1/2},$$

$$d_{i+1/2,j}=\frac{1}{2}\left(|\sigma(\hat{A})|\right)_{i+1/2,j}(\psi_{i+1,j}-\psi_{i,j})$$

$$d_{i,j+1/2}=\frac{1}{2}\left(|\sigma(\hat{B})|\right)_{i,j+1/2}(\psi_{i,j+1}-\psi_{i,j}). \tag{5}$$

$\hat{A}=\frac{\partial F}{\partial\psi}$, $\hat{B}=\frac{\partial G}{\partial\psi}$ should be positive definite matrices. In practice one never solves (5) through a true time dependent path since this would require thousands of time steps. Instead, one uses acceleration techniques such as local time stepping, multigrid etc. to increase the convergence rate. These techniques destroy the time dependence of the marching scheme and replaces it by a pseudo-time path to the steady state.

The choice of the marching variable, ψ, affects not only the iteration procedure but also the artificial viscosity in the steady state. However, ψ is a purely fictitious variable. So we can choose ψ as the marching variable but base the artificial viscosity on Φ, so that

$$d_{i+1/2,j}=\frac{1}{2}|\sigma(\frac{\partial F}{\partial\Phi})|_{i+1/2,j}(\Phi_{i+1,j}-\Phi_{i,j}),$$

$$d_{i,j+1/2}=\frac{1}{2}|\sigma(\frac{\partial G}{\partial\Phi})|_{i,j+1/2}(\Phi_{i,j+1}-\Phi_{i,j}) \quad . \tag{6}$$

Recall that for the original system $\psi=\Phi=w$. Define the matrix $\mathbf{P}^{-1}=\frac{\partial\psi}{\partial w}$. By the chain rule we have $\frac{\partial\psi}{\partial t}=\frac{\partial\psi}{\partial w}\frac{\partial w}{\partial t}=\mathbf{P}^{-1}\frac{\partial w}{\partial t}$. Also, define $\mathbf{Q}^{-1}=\frac{\partial\Phi}{\partial w}$. Since $\frac{\partial F}{\partial\Phi}=\frac{\partial F}{\partial w}\frac{\partial w}{\partial\Phi}$ equation (6) can be expressed as

$$d_{i+1/2,j} = \frac{1}{2}|\sigma(A\mathbf{Q})|_{i+1/2,j}(\Phi_{i+1,j} - \Phi_{i,j}) = \frac{1}{2}(|\sigma(A\mathbf{Q})|\mathbf{Q}^{-1})_{i+1/2,j}(w_{i+1,j} - w_{i,j}),$$
$$d_{i,j+1/2} = \frac{1}{2}|\sigma(B\mathbf{Q})|_{i,j+1/2}(\Phi_{i,j+1} - \Phi_{i,j}) = \frac{1}{2}(|\sigma(B\mathbf{Q})|\mathbf{Q}^{-1})_{i,j+1/2}(w_{i,j+1} - w_{i,j}) \quad . \tag{7}$$

We also consider the following preconditioning [33] for solving the pseudo-time equations:

$$\mathbf{P}^{-1}\frac{\partial w}{\partial t} + \frac{F_{i+1/2,j} - F_{i-1/2,j}}{\Delta\xi} + \frac{G_{i,j+1/2} - G_{i,j-1/2}}{\Delta\eta} = d_{i+1/2,j} - d_{i-1/2,j} + d_{i,j+1/2} - d_{i,j-1/2}, \tag{8}$$

$$d_{i+1/2,j} = \frac{1}{2}\left(\mathbf{Q}_x^{-1}|\sigma(\mathbf{Q}_x A)|\right)_{i+1/2,j}(w_{i+1,j} - w_{i,j}),$$
$$d_{i,j+1/2} = \frac{1}{2}\left(\mathbf{Q}_y^{-1}|\sigma(\mathbf{Q}_y B)|\right)_{i,j+1/2}(w_{i,j+1} - w_{i,j}) \quad . \tag{9}$$

where A and B are the Jacobian matrices $A = \frac{\partial F}{\partial w}, B = \frac{\partial G}{\partial w}$. We require that $\mathbf{Q}^{-1}|\sigma(\mathbf{Q}A)|$ be symmetric and so both forms (7) and (9) are equivalent.

In this algorithm $\mathbf{P}$ is constructed to accelerate the convergence to a steady state while $\mathbf{Q}_x, \mathbf{Q}_y$ are designed to give a "better" artificial viscosity, i.e. to improve the accuracy of the scheme. Originally it was assumed that $\mathbf{P} = \mathbf{Q}_x = \mathbf{Q}_y$. This is not always the best strategy. As an example, if we wish to keep the time accuracy of the equations then one can set $\mathbf{P} = I$ but still retain the preconditioned artificial viscosity for improved accuracy.

3 Limit to Incompressible Equations

We discuss two distinct ways of improving the convergence rate to a steady state. The first approach is to examine low speed flow. The second approach is to consider multigrid methods and the requirements on the damping of high frequency error components. This leads to two different ways of constructing a preconditioning matrix $\mathbf{P}$ which can be combined to achieve both goals. When analyzing the error we shall only consider the low speed problem. We shall demonstrate that a "standard" algorithm for the compressible equations does not converge to a reasonable algorithm for the incompressible equations. However, the introduction of the preconditioning **Q** into the artificial viscosity, or an upwind algorithm, solves this difficulty.

To see what is happening we explicitly display the scheme for the one dimensional Euler equations including a second order artificial viscosity. To simplify the notation we shall use differential notation instead of finite differences. It is understood that all spatial derivatives are to be replaced by central differences and that the time derived is replaced by some marching algorithm. We shall choose the independent variables to be

$$dw_0 = (d\phi, du, dv, \frac{a}{c_p}dS) \tag{10}$$

where $d\phi = \frac{dp}{\rho a}$ and $dS = \frac{c_v}{p}(dp - a^2 d\rho)$ and so $\frac{a}{c_p}dS = \frac{dp - a^2 d\rho}{\rho a}$. This leads to a nonconservative formulation for ease of presentation. The actual computations are based on (8-9) which is a conservative formulation.

We therefore consider

$$\begin{aligned} \phi_t + u\phi_x + au_x &= AV_1 \\ u_t + uu_x + a\phi_x &= AV_2 \\ v_t + uv_x &= AV_3 \\ S_t + uS_x &= AV_4 \quad . \end{aligned}$$

The equations for v and S decouple and so we shall only consider the two acoustic equations

$$\phi_t + u\phi_x + au_x = Q_{11}\phi_{xx} + Q_{12}u_{xx} \tag{11a}$$

$$u_t + uu_x + a\phi_x = Q_{21}\phi_{xx} + Q_{22}u_{xx} \quad . \tag{11b}$$

As $M \to 0$ we know from the asymptotic theory [11] that $\phi = O(M)$, $u = O(1)$ and $c = O(\frac{1}{M})$. Hence, in (11a) the first two terms on the left hand side are $O(M)$ while the third term is $O(\frac{1}{M})$. So, in the limit, on the left hand side we have only the divergence terms which arise from the $O(\frac{1}{M})$ term. We therefore require that the right hand side also have terms that are $O(\frac{1}{M})$. This requires that $Q_{11} = O(\frac{1}{M^2})$. We thus have ensured that the diagonal term of the artificial viscosity is of the correct order of magnitude as $M \to 0$. If we also wish the off-diagonal term to appear in the limit equations we require that $Q_{12} = O(\frac{1}{M})$. If Q_{12} is asymptotically smaller than this then it disappears in the incompressible limit. This is permissible but less desirable. We now repeat this process with (11b). The terms on the left hand side of the equation are now $O(1)$. Hence, we require that $Q_{22} = O(1)$. As before, if we wish the off-diagonal terms to remain in the incompressible limit then $Q_{21} = O(\frac{1}{M})$. Hence, the most important quantity is Q_{11}.

We now consider an artificial viscosity of the form $\mathbf{P}^{-1}|\mathbf{P}A|$. For the w_0 variables

$$A = \begin{pmatrix} u & a \\ a & u \end{pmatrix}.$$

We choose the simplest preconditioner given by

$$\mathbf{P} = \begin{pmatrix} \beta^2 & 0 \\ 0 & 1 \end{pmatrix} \quad , \quad \mathbf{P}^{-1} = \begin{pmatrix} \frac{1}{\beta^2} & 0 \\ 0 & 1 \end{pmatrix}.$$

We shall later see that this is an effective preconditioning when $\beta = O(M)$. On the other hand, when $\beta = 1$ there is no preconditioning. The eigenvalues of $\mathbf{P}A$ are given by

$$\lambda_{\pm} = 0.5 \left(zu \pm \sqrt{z^2 u^2 + 4\beta^2 (a^2 - u^2)} \right), \quad z = 1 + \beta^2.$$

One then finds for the matrix artificial viscosity that for $|u| < a$

$$\mathbf{P}^{-1}|\mathbf{P}A| = \frac{1}{\sqrt{z^2 u^2 + 4\beta^2 (a^2 - u^2)}} \begin{pmatrix} -(1-\beta^2)u^2 + 2a^2 & (1+\beta^2)au \\ (1+\beta^2)au & (1-\beta^2)u^2 + 2\beta^2 a^2 \end{pmatrix}.$$

Note that $\det(\mathbf{P}^{-1}|\mathbf{P}A|) = \frac{(1-\beta^2)^2u^2(a^2-u^2)+4\beta^2a^2}{z^2u^2+4\beta^2(a^2-u^2)} > 0$ and so both eigenvalues are positive for subsonic flow. For $\beta = M \ll 1$ we get

$$\mathbf{P}^{-1}|\mathbf{P}A| \sim \frac{1}{\sqrt{u^2+4\beta^2a^2}}\begin{pmatrix} 2a^2 & au \\ au & u^2+2\beta^2a^2 \end{pmatrix} \quad .$$

We consider the two extremes $\beta = M \ll 1$ and $\beta = 1$. We then have

$$\mathbf{P}^{-1}|\mathbf{P}A| \sim \frac{|u|}{\sqrt{5}}\begin{pmatrix} \frac{2}{\beta^2} & \frac{1}{\beta} \\ \frac{1}{\beta} & 3 \end{pmatrix}, \qquad \mathbf{P}^{-1}|\mathbf{P}A| = \begin{pmatrix} a & |u| \\ |u| & a \end{pmatrix}$$

respectively. Hence, without preconditioning Q_{11}, Q_{12}, Q_{21} are asymptotically too small while Q_{22} is too large. This implies that the momentum equation has too much artificial viscosity and this will swamp the physical terms while the divergence equation has too little viscosity which will cause oscillations. However, $\beta = M \ll 1$ gives the correct asymptotic order. Since $\mathbf{P}^{-1}|\mathbf{P}A|$ is symmetric both formulations (7) and (9) give the same artificial viscosity. Note that increasing β decreases the $(1,1)$ element of the artificial viscosity.

If we use a scalar viscosity instead of the matrix viscosity then for $M \ll 1$ we have

$$\mathbf{P}^{-1}\rho(\mathbf{P}A) = \frac{|u| + \sqrt{u^2+4\beta^2a^2}}{2}\begin{pmatrix} \frac{1}{\beta^2} & 0 \\ 0 & 1 \end{pmatrix} \quad .$$

In one dimension $\beta = \frac{u}{a}$ and this simplifies to

$$\mathbf{P}^{-1}\rho(\mathbf{P}A) = \frac{1+\sqrt{5}}{2}|u|\begin{pmatrix} \frac{1}{\beta^2} & 0 \\ 0 & 1 \end{pmatrix} \quad .$$

The matrix viscosity scales the acoustic waves with $\frac{|u|}{\beta^2}$ and the convective and entropy waves with u. The scalar viscosity scales all waves with $\frac{|u|}{\beta^2}$. If $\beta^2 \sim \frac{u^2+v^2}{a^2}$ then the scalar viscosity is much more viscous, in the x direction, whenever $v \ll u$. However, if we choose $\beta_x = \frac{u}{a}$ and $\beta_y = \frac{v}{a}$ in the artificial viscosity then the difference between the scalar and matrix viscosities is a small constant. However, when β is affected by the lower limit (see section 7) then again the scalar viscosity on the convective waves will be large compared with the matrix viscosity. Thus, from the accuracy viewpoint it is important to localize the effect of β_{min}.

The CUSP scheme reduces to a scalar artificial viscosity for $M < \frac{1}{2}$, [16, 30]. The viscosity of the CUSP scheme scales as u in the x direction and as v in the y direction. We can accomplish a similar scaling by choosing β_x and β_y as before. Hence, the two dimensional viscosities are identical to the one dimensional counterparts. The CUSP scheme was developed as a variation of the AUSM scheme. Preconditioning coupled with the AUSM scheme has been presented in [8].

4 Convergence Acceleration

We consider the Euler and Navier-Stokes equations when the Mach number in the flow is very small and so the convective speed is much smaller than the acoustic speed. Thus, any time step that is appropriate for one wave is inappropriate for the other wave. In particular, for an explicit scheme, the time step is bounded by the inverse of the acoustic speed. Hence, the convective waves move only a fraction of a cell in each time step. One can also show that the use of an ADI implicit method does not overcome these difficulties.

There are many cases in which one is interested in solving the compressible equations in spite of the fact that the local speeds are very small. In some cases there is a sufficient change in the energy so that the incompressible assumption is not valid. There are also many problems in which the flow has a very small Mach number in large regions but other regions of the flow are compressible and the flow may even become supersonic with shocks. Hence, we wish to keep the scheme in conservation form even though much of the flow is incompressible. Typical examples are flows about wings at high angle of attack, flows around bodies with large stagnation regions and internal flow, turbines, pumps, etc.

Hence, we shall introduce a preconditioner whose purpose is to equalize the speeds of all the waves. This can only be accomplished by destroying the time accuracy of the equations. Since we are primarily interested in steady-state flow this is not very important. We demand that at $M = 0$ the condition number of $\mathbf{P}A$ and $\mathbf{P}B$ be bounded independent of the Mach number. In addition we prefer that the condition number (at $M=0$) be close to one.

For the Euler equations, in w_0 variables, we introduce the preconditioner

$$\mathbf{P}^{-1} = \begin{pmatrix} \frac{1}{\beta^2} & 0 & 0 & \delta \\ \frac{\alpha u}{a\beta^2} & 1 & 0 & 0 \\ \frac{\alpha v}{a\beta^2} & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad , \quad \mathbf{P} = \begin{pmatrix} \beta^2 & 0 & 0 & -\beta^2\delta \\ -\frac{\alpha u}{a} & 1 & 0 & \frac{\alpha u}{a}\delta \\ -\frac{\alpha v}{a} & 0 & 1 & \frac{\alpha v}{a}\delta \\ 0 & 0 & 0 & 1 \end{pmatrix} . \tag{12}$$

β, α, δ are parameters to be chosen. δ is introduced mainly so that this preconditioner includes that suggested by Y. Choi and Merkle [3] ($\delta = 1$). In order to achieve a bounded condition number at $M = 0$ it is necessary to choose β proportional to M. We stress that both β and M are local quantities. However, near stagnation points M is very small relative to the rest of the flow and hence β becomes locally small. This implies that the matrix $\mathbf{P}$ is close to being singular. It is therefore necessary to prevent β from becoming too small locally. We shall later describe (section 7) some ways to achieve this.

In addition to difficulties at $M=0$ the condition number of A becomes infinite at $M=1$. If we choose

$$\alpha = 1 + \beta^2 \quad , \quad \beta^2 = \frac{M^2}{1 - M^2}$$

then we get the same eigenvalues as those of the symmetric optimal preconditioner (see next section). Hence, at $M = 1$ the condition number grows as $O(\frac{1}{\sqrt{1-M^2}})$ rather than $O(\frac{1}{1-M^2})$.

At least for external flows, the sonic line $M = 1$ does not seem to create difficulties for the iteration procedure. Hence, we prefer to eliminate the preconditioning for supersonic flows. This can be accomplished by specifying $\beta = 1, \alpha = 0, \delta = 0$ for supersonic flow. This is then connected smoothly to the desired preconditioning at low Mach numbers. We then have

$$\mathbf{P}A = \begin{pmatrix} \beta^2 u & \beta^2 a & 0 & \frac{-\beta^2 u\delta}{\rho a} \\ a\left(1 - \frac{\alpha u^2}{a^2}\right) & (1-\alpha)u & 0 & \frac{\alpha u^2 \delta}{\rho a^2} \\ \frac{-\alpha uv}{a} & -\alpha v & u & \frac{\alpha uv\delta}{\rho a^2} \\ 0 & 0 & 0 & u \end{pmatrix} \quad . \tag{13}$$

The eigenvalues of $\mathbf{P}A$ are

$$\lambda_0 = u$$

and

$$\lambda_\pm = 0.5\left(zu \pm \sqrt{z^2u^2 + 4\beta^2(a^2 - u^2)}\right), \quad z = 1 - \alpha + \beta^2 \quad . \tag{14}$$

Note that the eigenvalues are independent of δ. If $\alpha = 1+\beta^2$ then $z = 0$ and $\lambda_\pm = \pm\beta a\sqrt{1 - M^2}$. The right eigenvectors of $\mathbf{P}A$ are given by the columns of

$$R^{-1} = \begin{pmatrix} \beta^2 a & \beta^2 a & 0 & 0 \\ \lambda_+ - \beta^2 u & \lambda_- - \beta^2 u & 0 & \frac{\delta u}{\rho a^2} \\ \frac{\alpha v\lambda_+}{u-\lambda_+} & \frac{\alpha v\lambda_-}{u-\lambda_-} & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad . \tag{15}$$

For generalized coordinates we replace (14) by

$$U = us_x + vs_y, \qquad |S_\xi|^2 = s_x^2 + s_y^2$$

$$\lambda_0 = U$$

$$\lambda_\pm = \frac{1}{2}\left(zU \pm \sqrt{z^2U^2 + 4\beta^2(a^2|S_I|^2 - U^2)}\right) \tag{16}$$

where $S_\xi = (s_x, s_y)$ is the cell face vector in ξ-direction. The largest eigenvalue is then used to determine the (inviscid) time step. For $M \sim 0$, $\lambda_\pm \sim 0.5U\left((1-\alpha) \pm \sqrt{(1-\alpha)^2+4}\right)$. So for $\alpha = 0$ the condition number is $\frac{1+\sqrt{5}}{1-\sqrt{5}} \sim 2.6 + O(M)$ while for $\alpha = 1$ the condition number is $1 + O(M)$. We also express these matrices in terms of the variables (p, u, v, T). Let

$$g = 1 + (\gamma - 1)\delta \quad = \quad \begin{cases} 1, & \text{if} \quad \delta = 0 \\ \gamma, & \text{if} \quad \delta = 1 \end{cases}$$

Then,

$$A_T = \begin{pmatrix} u & \rho a^2 & 0 & 0 \\ \frac{1}{\rho} & u & 0 & 0 \\ 0 & 0 & u & 0 \\ 0 & (\gamma-1)T & 0 & u \end{pmatrix}, \quad B_T = \begin{pmatrix} v & 0 & \rho a^2 & 0 \\ 0 & v & 0 & 0 \\ \frac{1}{\rho} & 0 & v & 0 \\ 0 & 0 & (\gamma-1)T & v \end{pmatrix},$$

$$\mathbf{P}_T^{-1} = \begin{pmatrix} \frac{1}{\beta^2} - (\gamma-1)\delta & 0 & 0 & \frac{\gamma p}{T}\delta \\ \frac{\alpha u}{\rho \beta^2 a^2} & 1 & 0 & 0 \\ \frac{\alpha v}{\rho \beta^2 a^2} & 0 & 1 & 0 \\ \frac{1}{\rho c_p}\left(\frac{1}{\beta^2} - g\right) & 0 & 0 & g \end{pmatrix}, \mathbf{P}_T = \begin{pmatrix} \beta^2 g & 0 & 0 & -\beta^2 \frac{\gamma p}{T}\delta \\ -\frac{\alpha u g}{\rho a^2} & 1 & 0 & \frac{\alpha u}{\rho a^2}\frac{\gamma p}{T}\delta \\ -\frac{\alpha v g}{\rho a^2} & 0 & 1 & \frac{\alpha v}{\rho a^2}\frac{\gamma p}{T}\delta \\ \frac{1}{\rho c_p}(\beta^2 g - 1) & 0 & 0 & 1 - (\gamma-1)\beta^2\delta \end{pmatrix}. \quad (17)$$

5 Symmetric Preconditioner

Until now we have stressed the importance of preconditioning for low Mach numbers. Van Leer, Lee and Roe [39] noticed that the two dimensional supersonic equations can be diagonalized if an appropriate preconditioner is added to the system. They then constructed a preconditioner for the subsonic case that it some sense is the continuation of the supersonic case. This continuation can never be exact since there must be a singularity at the sonic line. Considering

$$\mathbf{P}^{-1}\frac{\partial w}{\partial t} + A\frac{\partial w}{\partial x} + B\frac{\partial w}{\partial y} = 0$$

they demanded that A, B be symmetric and $\mathbf{P}$ be symmetric positive definite. Define

$$\widehat{\beta} = \begin{cases} \sqrt{1-M^2} & M < 1, \\ \sqrt{M^2-1} & M \geq 1 \end{cases}$$

$$\tau = \begin{cases} \sqrt{1-M^2} = \widehat{\beta} & M < 1, \\ \sqrt{1-M^{-2}} = \frac{\widehat{\beta}}{M} & M \geq 1 \end{cases} \quad .$$

The van Leer-Lee-Roe preconditioning for general nonaligned flow in w_0 variables is

$$\mathbf{P}_{VL} = \begin{pmatrix} \frac{\tau}{\beta^2} M^2 & \frac{-\tau}{\beta^2} u/a & \frac{-\tau}{\beta^2} v/a & 0 \\ \frac{-\tau}{\beta^2} u/a & (\frac{\tau}{\beta^2}+1)\frac{u^2}{u^2+v^2} + \tau\frac{v^2}{u^2+v^2} & (\frac{\tau}{\beta^2}+1)\frac{uv}{u^2+v^2} & 0 \\ \frac{-\tau}{\beta^2} v/a & (\frac{\tau}{\beta^2}+1)\frac{uv}{u^2+v^2} & (\frac{\tau}{\beta^2}+1)\frac{v^2}{u^2+v^2} + \tau\frac{u^2}{u^2+v^2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

$\widehat{\beta}$ is introduced to change the envelope of the wave form while τ changes the amplitude of the waves. Define $D = A\omega_1 + B\omega_2$. The eigenvalues of $\mathbf{P}_{VL}D$, for subsonic flow, are

$$\begin{aligned} \lambda_0 &= u\omega_1 \quad \text{twice} \\ \lambda_{\pm} &= \pm\tau Mc\sqrt{\omega_1^2 + \frac{\omega_2^2}{1-M^2}} \quad . \end{aligned}$$

The symmetric preconditioner has the nice property that it always leads to a well posed problem, at least for Mach numbers bounded away from zero (note the preconditioner of the previous section is symmetric in w_0 variables only when $\alpha = \delta = 0$, however it can sometimes be symmetrized in other variables). It is also useful in constructing canonical forms that separate the Euler equations into its elliptic and hyperbolic constitute parts [22, 26]. However, this preconditioner seems to have difficulties near stagnation points. One possible reason is that the preconditioner depends on the flow angle even near stagnation points. However, for very small Mach numbers the angle the flow makes with the x-axis is not well defined.

This leads to a numerical indeterminacy near stagnation points. Possible cures for this are discussed in [40]. These cures depend on modifications that remove this dependency on the flow angle at the cost of a loss of being optimal. Similar to the preconditioner (12) one may limit the value of M in the (1,1) element of $\mathbf{P}_{VL}$ near stagnation points.

6 Artificial Viscosity

As we previously demonstrated the "standard" methods do not converge to the correct solution as the Mach number goes to zero. In particular this is a consequence of an incorrect scaling, as $M \to 0$, of the artificial viscosity.

For the general preconditioner in two dimensions expressed in terms of w_0 variables a necessary condition for convergence as $M \to 0$ is that [33]

$$\mathbf{Q}_x^{-1}|\mathbf{Q}_x A| \sim \begin{pmatrix} O(\frac{1}{M^2}) & O(\frac{1}{M}) & O(\frac{1}{M}) & O(1) \\ O(\frac{1}{M}) & O(1) & O(1) & O(1) \\ O(\frac{1}{M}) & O(1) & O(1) & O(1) \\ O(1) & O(1) & O(1) & O(1) \end{pmatrix} . \tag{18}$$

and similarly in the other directions.

For simplicity of presentation we consider the one dimensional example of (9):

$$\mathbf{P}^{-1}w_t + f_x = \mathbf{P}^{-1}w_t + Aw_x = h(\mathbf{P}^{-1}|\mathbf{P}A|w_x)_x \quad . \tag{19}$$

A sufficient condition for this to be well posed is that

I. A is symmetric .

II. $\mathbf{P}$ be positive definite and hence $\mathbf{P}^{1/2}$ exists and is symmetric.

It is sufficient that this be true with respect to some set of variables. Equivalently these conditions hold after a similarity transformation. A consequence of these two conditions is that $\mathbf{P}A$ is similar to a symmetric matrix and $\mathbf{P}^{-1}|\mathbf{P}A|$ is symmetric positive definite. However, the converse is not true and so it is not sufficient to show that $\mathbf{P}A$ is similar to a symmetric matrix. When both $\mathbf{P}$ and A are symmetric it follows that $\mathbf{P}A$ is similar to a symmetric matrix since $\mathbf{P}A = \mathbf{P}^{1/2}(\mathbf{P}^{1/2}A\mathbf{P}^{1/2})\mathbf{P}^{-1/2}$. $\mathbf{P}^{-1}|\mathbf{P}A|$ is symmetric as a consequence of the Cayley-Hamilton theorem. By this theorem the absolute value can be written as a finite polynomial in the matrices. Hence

$$|\mathbf{P}A| = \alpha_0 I + \alpha_1 \mathbf{P}A + \alpha_2(\mathbf{P}A)^2 + \dots$$

and so

$$\mathbf{P}^{-1}|\mathbf{P}A| = \alpha_0\mathbf{P}^{-1} + \alpha_1 A + \alpha_2 A\mathbf{P}A + \alpha_3 A\mathbf{P}A\mathbf{P}A \dots \qquad .$$

By inspection this is symmetric whenever both $\mathbf{P}$ and A are symmetric. Furthermore,

$$\mathbf{P}^{-1}|\mathbf{P}A| = \mathbf{P}^{-1/2}\left(\mathbf{P}^{-1/2}|\mathbf{P}A|\mathbf{P}^{-1/2}\right)\mathbf{P}^{1/2}$$

and so the sign of the eigenvalues of $\mathbf{P}^{-1}|\mathbf{P}A|$ is the same as $|\mathbf{P}A|$ and so it is positive definite. Multiplying (19) by an arbitrary matrix Q we can rewrite it as

$$Qw_t + Q\mathbf{P}Aw_x = \frac{h}{2}Q(\mathbf{P}^{-1}|\mathbf{P}A|w_x)_x \qquad .$$

We wish to choose Q positive definite and $Q\mathbf{P}A$ as a symmetric matrix S. The system is then well-posed in a norm weighted by the matrix Q. It then follows that $Q^{1/2}\mathbf{P}AQ^{-1/2} = Q^{-1/2}SQ^{-1/2}$ and so $\mathbf{P}A$ is symmetrizable. So these are equivalent approaches. In practice, the matrix Q is singular at $M=0$ and so this approach does not give more information about the limit process.

It can be shown, [31], for the preconditioner (12) that $\mathbf{P}A$ can be symmetrized if and only if $\beta^2 > \alpha M^2$. $\alpha = 1$ is optimal in terms of the condition number at $M=0$. In practice, choosing $\alpha = 1$ results in a less robust scheme. One possible reason for this is that the preconditioning affects not only the time step but also the artificial viscosity. An optimal preconditioning lowers the artificial viscosity which improves the accuracy but deteriorates the robustness of the scheme. Most successful computations to-date have been with the original preconditioner, $\alpha = 0$ and both $\delta = 0$ and $\delta = 1$ (see, for example, [17, 23, 34, 35, 37, 38, 41]).

We have previously shown the details of the artificial viscosity for one dimensional flow and $\alpha = \delta = 0$. We showed that the artificial viscosity scales correctly only with preconditioning. Guillard and Viozat [10] have further shown that if one considers the two dimensional Euler equations with an upwind scheme then the pressure jumps scale correctly as $M \to 0$ only if preconditioning is used.

In two dimensions we have (see (9))

$$\mathbf{P}^{-1} w_t + f_x + g_y = \frac{h}{2}\left[(Q_x^{-1}|Q_x A| w_x)_x + (Q_y^{-1}|Q_y B| w_y)_y \right] \quad .$$

As previously stated we need not choose $Q_x = Q_y = \mathbf{P}$. One particular choice is to choose them all of the same form but with different a β. Thus, for $\mathbf{P}$ we can choose $\beta^2 = \frac{u^2+v^2}{a^2}$ while for Q_x we have $\beta^2 = \frac{u^2}{a^2}$ and for Q_y we have $\beta^2 = \frac{v^2}{a^2}$. For time dependent flow we can choose $\mathbf{P} = I$ while keeping the preconditioned Q_x and Q_y.

7 Definition of β

As previously stated all the preconditioners become singular at $M = 0$. To avoid inverting a singular matrix it is necessary to prevent β from becoming too small. This rationale indicates that we can cutoff β at some very low level. However, experience indicates that one cannot allow β to get very small relative to the general flow. The reasons for this are still unclear. Darmofal and Schmid [4] present one possible explanation based on the non-orthogonality of the eigenvectors since the matrices are not normal. However, even this explanation is qualitative rather than quantitative. A complete theory for this is quite difficult. In particular it would have to allow for variable coefficients. When the global Mach number approaches zero the preconditioners described work quite nicely. The difficulties arise only when the Mach number locally approaches zero near a stagnation point. The cutoff suggested by Turkel [31] is

$$\beta^2 = \max\left(\min(K_1 M^2, \beta^2_{min}), 1\right) \qquad \beta^2_{min} = K_2 M^2_{ref} \quad K_1 \sim 1 \quad . \tag{20}$$

This removes the preconditioning for supersonic flow. M_{ref} is usually taken as the inflow Mach number for external flow. For internal flow one can use an average or maximum Mach number over the flow. This cutoff has several disadvantages. The cutoff near the stagnation point depends on global properties of the flow. In addition the constant K_2 seems to be very problem dependent. For easy cases one can choose $K_2 \sim .2$, more typically $K_2 \sim 1$ while for hard cases $K_2 \sim 3 - 5$. When $K_2 \geq 1$ we do not allow β to be less than the inflow Mach number. This usually implies that β is constant over much of the flow.

As an alternative to this Darmofal and Siu [5] suggest letting β_{min} depend on local gradients of the pressure and velocity components. Consider a cell-centered scheme. At each face we have right and left (or up and down) states. We set

$$\beta^2_{min} = K_2 \left(\frac{|p_r - p_l|}{\rho a^2} + \frac{\sqrt{(u_r - u_l)^2 + (v_r - v_l)^2}}{a} \right) \qquad .$$

Choose β_{min} at the center node as the maximum of β_{min} over all the faces. This definition is now a strictly local definition and avoids many of the deficiencies of the original cutoff.

Contour plots show that indeed β is now limited only in the neighborhood of stagnation points. For easy cases one can choose $K_2 = 1$ however for harder problems one must change K_2. Hence, the robustness of this cutoff also needs improvement. Unfortunately, the convergence rate to the steady state does not seem to improve with this newer limiter compared with the original. In some cases this local cutoff gives a slower convergence rate than the global cutoff.

As previously stressed the preconditioner affects not only the convergence rate but also the accuracy of the steady state solution. It therefore is possible that the local cutoff for β may improve the accuracy even when it does not improve the convergence rate.

To account for aspect ratio we consider $\frac{A}{\Delta x}$ and $\frac{B}{\Delta y}$. Then we replace M in (20) by

$$M^2 \to \frac{(\frac{u}{a\Delta x})^2 + (\frac{v}{a\Delta y})^2}{(\frac{1}{\Delta x})^2 + (\frac{1}{\Delta y})^2} \quad .$$

In generalized coordinates we have

$$M_\xi = \frac{u\xi_x + v\xi_y}{a\sqrt{\xi_x^2 + \xi_y^2}}, \qquad M_\eta = \frac{u\eta_x + v\eta_y}{a\sqrt{\eta_x^2 + \eta_y^2}}, \qquad \beta^2 = M_\xi^2 + M_\eta^2 \quad .$$

8 Boundary Conditions

In many codes the boundary conditions in the far field are based on characteristic variables, even for viscous flow. Thus, at inflow the incoming variables, corresponding to positive eigenvalues, are specified while the outgoing variables, corresponding to negative eigenvalues, are extrapolated. Once we change the time dependent equations we also change the characteristics of the system (though not the signs of the eigenvalues). Hence, it may be necessary to also modify the boundary conditions for the preconditioned system. The original characteristic variables do not scale correctly as the Mach goes to zero and so we do not obtain well-posed boundary conditions in the limit.

Since the eigenvalues and eigenvectors are invariant under a similarity transform we shall use the (p, u, v, S) set of variables to determine the new eigenvalues and characteristic variables. We shall consider the flow normal to the boundary which we take to be a line $x = const.$ Thus, we need to use the characteristic variables corresponding to the left eigenvectors of $P_0 A_0$. Flow variables at a boundary are then obtained by taking the values for R_1, R_2 and R_+ from upstream flow whereas downstream values are used for R_-. When we diagonalize $P_0 A_0 = R D R^{-1}$ the rows of R are the left eigenvectors of $P_0 A_0$ while the columns of R^{-1} are the right eigenvectors of $P_0 A_0$, (15). Let variables with a subscript "b" denote the variables to be computed at the boundary. Variables with a subscript "a" denote the flow state approaching the boundary, i.e. specified at inflow, whereas variables with a subscript "l" are used for the flow leaving the boundary, e.g. extrapolated at inflow. This notation can be consistently used for inflow and outflow boundaries. The characteristic boundary conditions are expressed in (p, u, v, T) variables as:

$$
\begin{aligned}
p_b = \frac{\lambda_+ p_a - \lambda_- p_l}{\lambda_+ - \lambda_-} - \frac{1}{\lambda_+ - \lambda_-} \Big[& u\Big(1 - \alpha - \delta \frac{(\gamma-1)\beta M_r^2}{a^2}\Big)(p_a - p_l) \\
& - \rho\beta M_r^2 (u_a - u_l) + \frac{\delta \rho \beta M_r^2 u}{T}(T_a - T_l) \Big]
\end{aligned}
$$

$$
\begin{aligned}
u_b &= u_a - \frac{\lambda_+ - u(1-\alpha)}{\rho\beta M^2}(p_b - p_a) \\
T_b &= T_a + \frac{(\gamma-1)T}{\gamma p}(p_b - p_a) \\
v_b &= v_a - \frac{\alpha v}{\beta M^2 - \alpha u^2}\left[\frac{p_b - p_a}{\rho} + u(u_b - u_a)\right] \quad .
\end{aligned} \tag{21}
$$

In generalized coordinates u and v are replaced by the normal and tangential components of the velocity relative to the boundary.

Computationally we observe that a simplified set of conditions yields almost the same results as those of the characteristic boundary conditions. These are given by

$$v_b = v_a, \quad p_b = p_l, \quad u_b = u_a,$$

and either $T_b = T_a$ or $\rho_b = \rho_a + \frac{p_b - p_a}{a^2}$.

9 Implementation

We have described the preconditioner mainly in terms of the symmetric variables w_0. While these variables are useful for analysis they are not the most appropriate for computations. Since we wish to also capture shocks it is important that the flux terms be evaluated in terms of the conservative variables. Hence, one frequently retains the conservative variables $w_c = (\rho, \rho u, \rho v, E)$ both for the time stepping and preconditioning and also for the artificial viscosity. On the other hand the density is not very appropriate for almost incompressible flow. In particular, space derivatives of the density are almost zero and so ill-conditioned for use in the artificial viscosity. Many researchers prefer the variables $w_T = (p, u, v, T)$ or (p, u, v, h). Also some researchers use a reduced pressure and reduced enthalpy to reduce roundoff errors [12]. Hauke and Hughes [14] have compared many posisble sets of variables and found that the w_T variables is the best all-around. For simplicity of notation we shall describe all the formulas in one space dimension and in differential notation. However, the extension to multidimensions and finite differences is immediate. We first consider the non-conservative quasi-linear system and only later consider the conservative flux form. Hence, we consider the preconditioned system

$$\mathbf{P}_0^{-1}\frac{\partial w_0}{\partial t} + A_0\frac{\partial w_0}{\partial x} = \frac{h}{2}\mathbf{P}_0^{-1}|\sigma(\mathbf{P}_0 A_0)|\frac{\partial^2 w_0}{\partial x^2} \quad . \tag{22}$$

If we wish to work with another set of variables w_i this is replaced by

$$\mathbf{P}_i^{-1}\frac{\partial w_i}{\partial t} + A_i\frac{\partial w_i}{\partial x} = \frac{h}{2}\mathbf{P}_i^{-1}|\sigma(\mathbf{P}_i A_i)|\frac{\partial^2 w_i}{\partial x^2} \tag{23}$$

where $\mathbf{P}_i = \frac{\partial w_i}{\partial w_0}\mathbf{P}_0\frac{\partial w_0}{\partial w_i}$ and similarly for A_i. Because we wish the flux differencing to be conservative we introduce the conservative variables w_c with $\frac{\partial f}{\partial x} = A_c\frac{\partial w_c}{\partial x}$. We rewrite (23) for w_c as

$$\mathbf{P}_c^{-1}\frac{\partial w_c}{\partial t} + A_c\frac{\partial w_c}{\partial x} = \mathbf{P}_c^{-1}\frac{\partial w_c}{\partial t} + \frac{\partial f}{\partial x} = \frac{h}{2}\mathbf{P}_c^{-1}|\sigma(\mathbf{P}_c A_c)|\frac{\partial^2 w_c}{\partial x^2} \quad .$$

Changing the time derivative to w_T variables we get

$$\frac{\partial w_c}{\partial w_T}\mathbf{P}_T^{-1}\frac{\partial w_T}{\partial t} + \frac{\partial f}{\partial x} = \frac{h}{2}\mathbf{P}_c^{-1}|\sigma(\mathbf{P}_c A_c)|\frac{\partial^2 w_c}{\partial x^2} \quad .$$

Define

$$\Gamma^{-1} = \frac{\partial w_c}{\partial w_T}\mathbf{P}_T^{-1} = \mathbf{P}_c^{-1}\frac{\partial w_c}{\partial w_T} = \frac{\partial w_c}{\partial w_0}\mathbf{P}_0^{-1}\frac{\partial w_0}{\partial w_T} \quad .$$

It then follows that

$$\Gamma^{-1}\frac{\partial w_T}{\partial t} + \frac{\partial f}{\partial x} = \frac{h}{2}\mathbf{P}_c^{-1}|\sigma(\mathbf{P}_c A_c)|\frac{\partial^2 w_c}{\partial x^2} \quad . \tag{24}$$

In two dimensional curvilinear coordinates we have

$$\frac{\Delta(w_T)_{i,j}}{\Delta t} = -\Gamma_{i,j}\left\{\frac{F_{i+1/2,j} - F_{i-1/2,j}}{\Delta\xi} + \frac{G_{i,j+1/2} - G_{i,j-1/2}}{\Delta\eta}\right. \tag{25}$$

$$\left. - \left[d_{i+1/2,j} - d_{i-1/2,j} + d_{i,j+1/2} - d_{i,j-1/2}\right]\right\} \tag{26}$$

where

$$d_{i+1/2,j} = \Gamma^{-1}_{i+1/2,j} \left(|\sigma(P_T A_T)| \right)_{i+1/2,j} \left[(w_T)_{i+1,j} - (w_T)_{i,j} \right].$$

$$d_{i,j+1/2} = \Gamma^{-1}_{i,j+1/2} \left(|\sigma(P_T B_T)| \right)_{i,j+1/2} \left[(w_T)_{i,j+1} - (w_T)_{i,j} \right].$$

Hence, given $\mathbf{P}_0^{-1}$ we can determine Γ^{-1} and vice-versa. There have been several suggestion for Γ^{-1} presented in the literature based on other considerations and not beginning with $\mathbf{P_0}$. We present several of these options (note these papers usually use Γ for what we denote as Γ^{-1}).

Let $h = \frac{E+p}{\rho}$. The suggestion of Choi and Merkle [3] is

$$\Gamma^{-1}_{CM} = \begin{pmatrix} \frac{1}{a^2\beta^2} & 0 & 0 & 0 \\ \frac{u}{a^2\beta^2} & \rho & 0 & 0 \\ \frac{v}{a^2\beta^2} & 0 & \rho & 0 \\ \frac{h}{a^2\beta^2} - 1 & \rho u & \rho v & \frac{\gamma p}{(\gamma-1)T} \end{pmatrix}.$$

This corresponds to the preconditioner (12) with $\alpha = 0$ and $\delta = 1$. This has the disadvantage that it does not yield the nonpreconditioned case when $\beta = 1$. Hence, there is no easy way to transition from the preconditioned problem to the original equations across the sonic line.

Weiss and Smith [41] suggest

$$\Gamma^{-1}_{WS} = \begin{pmatrix} \Theta & 0 & 0 & -\frac{\rho}{T} \\ \Theta u & \rho & 0 & -\frac{\rho u}{T} \\ \Theta v & 0 & \rho & -\frac{\rho v}{T} \\ \Theta h - 1 & \rho u & \rho v & -\rho(\frac{h}{T} - C_p) \end{pmatrix} ,$$

$\Theta = \frac{1}{a^2}(\frac{1}{\beta^2} + (\gamma - 1))$. This corresponds to the preconditioner (12) with $\alpha = 0$ and $\delta = 0$ (see also [6]). Instead of computing in (p, u, v, T) variables one can work directly with the conservative variables. One then multiplies with $\mathbf{P}_c$ where

$$\mathbf{P}_c = I - \frac{(1-\beta^2)(\gamma-1)}{a^2} \begin{pmatrix} \frac{q^2}{2} & -u & -v & 1 \\ \frac{uq^2}{2} & -u^2 & -uv & u \\ \frac{vq^2}{2} & -uv & v^2 & v \\ \frac{hq^2}{2} & -uh & vh & h \end{pmatrix}.$$

Because of the special structure of the matrix we have

$$\mathbf{P}_c \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} - \frac{(1-\beta^2)(\gamma-1)}{a^2} \left(\frac{q^2}{2} x_1 - u x_2 - v x_3 + x_4 \right) \begin{pmatrix} 1 \\ u \\ v \\ h \end{pmatrix}$$

One can realize this formula also through a Sherman Morrison formula (Jack Edwards - private communication - see also [17]).

In a later paper Merkle [20] suggests a different matrix that depends on the cell Reynolds number, R_Δ and the Prandtl number Pr,

$$\Gamma_M^{-1} = \begin{pmatrix} \rho'_p & 0 & 0 & \rho'_T \\ u\rho'_p & \rho & 0 & u\rho'_T \\ v\rho'_p & 0 & \rho & v\rho'_T \\ h\rho'_p & \rho u & \rho v & h\rho'_T + \rho C_p \end{pmatrix}. \tag{27}$$

where $\rho'_p = \frac{k_p \min(1, R_\Delta^2)}{a^2\beta^2} = \frac{k_p}{a^2\beta^2 \max(1, 1/R_\Delta^2)}$, $\rho'_T = \frac{k_T \min(1, \frac{R_\Delta}{Pr})\rho}{T}$. If we wish the preconditioned matrix to transition to the unpreconditioned matrix at the sonic line we must choose $k_p = \gamma$. This does not correspond to any δ in the formulation of (12). Γ^{-1} corresponding to a general δ is given in the appendix of [35].

The major difficulty in implementing a preconditioning is the modification of the artificial viscosity. There are two choices in terms of this implementation. One possibility is to remain with conservative variables. This requires converting the matrices $\mathbf{P}$ and $\mathbf{P}|\sigma(\mathbf{P}A)|$ to conservative variables. Another possibility is to base the artificial viscosity on primitive variables, for example, (p, u, v, T) (one can also use enthalpy instead of temperature [37]). The conversion of the preconditioners to these set of variables is relatively simple since it is close to the $(d\phi, du, dv, dp)$ set of variables. The advantage of the primitive variables is that they are more appropriate for incompressible flow than conservative variables since the density is a constant in the incompressible limit. Then

$$\mathbf{P_c}^{-1}|\mathbf{P}_c A_c|\frac{\partial w_c}{\partial x} = \Gamma^{-1}|\mathbf{P}_T A_T|\ \frac{\partial w_T}{\partial x} \quad .$$

The explicit formula for $\Gamma^{-1}|\mathbf{P}_T A_T|$ is given in [34].

We present the steps used to introduce preconditioning into a pre-existing code. For simplicity we assume an explicit iteration scheme (e.g. Runge-Kutta) possibly augmented by implicit residual smoothing and multigrid.

I. Precondition the residual

- Select a reasonable set of variables for the problem e.g. $w_T = (p, u, v, T)$. Other sets of variables are possible. This set is used for illustration only.
- Viscosity includes Γ^{-1} and so even a scalar viscosity requires a matrix multiply.
- Add inviscid, viscous and artificial viscosity fluxes to obtain the total residual.
- Multiply this total residual by Γ, (26).

II. In each Runge-Kutta cycle perform residual smoothing and restrict the residual to the next coarser grid. These are performed on the residuals in w_T variables.

III. Calculate the total preconditioned residual in w_T variables. Then add the residuals to the old values to get w_T at the next Runge-Kutta stage. Calculate the conservative variables as nonlinear functions of w_T.

IV. Choose a new time step based on the preconditioning.

V. Modify the far field boundaries based on the preconditioned characteristics. If characteristic data is used at physical boundaries (e.g. inflow to a nozzle) these should also be modified. Alternatively, specify/extrapolate (p, u, v, T) at the boundaries.

10 Computational Results

Consider inviscid flow over a NACA 0012 airfoil on a C-type grid with 224×40 cells clustered near the leading and trailing edges. The results are compared with a panel method. As seen in table 1 the code using the preconditioning ($\alpha = 0$ and $\alpha = 1$) gives the lift and drag quite close to the panel method. In order to get the case $\alpha = 1$ to converge we used a coefficient of artificial viscosity that was twice as large as that used for $\alpha = 0$. This cancels out the effect, in the artificial viscosity, of using a smaller eigenvalue. Hence, the accuracy of the two cases are approximately equal. The non-preconditioned code has large variations as the Mach number goes to zero and does not converge to the panel code results. In particular the drag, which should be zero for inviscid flow, is very large without preconditioning but close to zero when using preconditioning [35].

	Mach	panel	no precond	α=0	α=1
C_L	0.1	.241	.2448	.2434	.2432
	0.01		.2172	.2421	.2419
	0.001		.1097	.2421	.2419
C_D	0.1	0.0	.0008	.00027	.00025
	0.01		.0086	.00027	.00025
	0.001		.0731	.00027	.00025

Table 1: Comparison of inviscid lift and drag with/without preconditioning.

Further cases demonstrating both the improved accuracy of the preconditioning and the improved convergence rates are given in [4, 9, 17, 21, 23, 27, 34, 35, 37, 38].

11 Jacobi Preconditioning

In this section we shall consider the matrix artificial viscosity. Consider the semi-discrete approximation

$$w_t + \frac{f_{i+1/2,j} - f_{i-1/2,j}}{\Delta x} + \frac{g_{i,j+1/2} - g_{i,j-1/2}}{\Delta y} = \frac{1}{2}\left(\Delta_x(|A|\Delta_x w) + \Delta_y(|B|\Delta_y w)\right) \tag{28}$$

where Δ are appropriate finite difference operators. A and B are the Jacobians of f and g with respect to w respectively. Linearizing and taking Fourier transforms we get

$$G(\theta, \phi) = i\left(A\sin(\theta) + B\sin(\phi)\right) - |A|(1 - \cos(\theta)) - |B|(1 - \cos(\phi))$$

and so

$$G(\pi, \pi) = -2|A| - 2|B| \quad . \tag{29}$$

It is difficult to choose coefficients in a Runge-Kutta scheme that will damp all the eigenvalues of the matrix $G(\pi, \pi)$. We note that with a scalar artificial viscosity this difficulty does not arise. The CUSP artificial viscosity is also a scalar viscosity when $M < \frac{1}{2}$. For $M > \frac{1}{2}$, CUSP is equivalent to a viscosity that is a linear function in A (in the x direction). Hence, we have a similar difficulty to that encountered with the matrix viscosity. Allmaras [2] and later Pierce and Giles [24] suggest replacing (28) by

$$\frac{|A| + |B|}{\rho(A) + \rho(B)} w_t + \frac{f_{i+1/2,j} - f_{i-1/2,j}}{\Delta x} + \frac{g_{i,j+1/2} - g_{i,j-1/2}}{\Delta y} = \frac{1}{2}\left(\Delta_x(|A|\Delta_x w) + \Delta_y(|B|\Delta_y w)\right) \tag{30}$$

where $\rho(A)$ is the spectral radius of A. (29) is now replaced by

$$G(\pi, \pi) = -2\rho(A) - 2\rho(B) \quad .$$

Hence, at the highest frequency G is essentially a scalar and so one can choose Runge-Kutta coefficients to damp this scalar quantity. As described in the previous section damping the highest frequencies is only one necessary property of the smoother. It is also necessary to convect low frequency components of the error out of the domain. One is required to use a smaller time step than the standard method with residual smoothing. Low speed preconditioning on the other hand can substantially change the eigenvalues of the system and allow much large time steps. Equation (30) is a preconditioner that does not change the steady state. It accelerates the convergence to a steady state by altering the high frequencies of the finite difference approximation.

12 Preconditioning-Squared

The Jacobi preconditioning (30) does nothing to improve the condition number of the system for low Mach numbers. In particular it does not satisfy the condition (18). On the other hand the low Mach number preconditioner (12) was not constructed to treat high frequency errors. Hence, we combine the two approaches. We first consider the low Mach number preconditioner to both accelerate the convergence to a steady state and to give the correct scaling for the artificial viscosity for low speed flow. On top of this we apply the Jacobi preconditioning. The Jacobi preconditioning is only applied as a convergence acceleration technique. The artificial viscosity terms are affected only by the low speed preconditioner. We thus have

$$\mathbf{P}^{-1}\frac{|\mathbf{P}A|+|\mathbf{P}B|}{\rho(\mathbf{P}A)+\rho(\mathbf{P}B)}w_t+\frac{f_{i+1/2,j}-f_{i-1/2,j}}{\Delta x}+\frac{g_{i,j+1/2}-g_{i,j-1/2}}{\Delta y}$$
$$=\frac{1}{2}\left(\Delta_x(\mathbf{P}^{-1}|\mathbf{P}A|\Delta_x w)+\Delta_y(\mathbf{P}^{-1}|\mathbf{P}B|\Delta_y w)\right) \quad . \quad (31)$$

The preconditioner (12) was based on the partial differential equation. This is advantageous because it can then be applied together with any numerical scheme. In the past it has been used with both central difference, [34], and upwind schemes [9, 41], structured and unstructured meshes [41, 19], explicit and implicit [3, 20, 23] time marching methods for both external and internal [37] flows without any alteration of the algorithm. On the other hand this approach is valid only when dealing with low frequencies. The high frequency component in the error does not approximate the differential equation. Hence, any treatment of the high frequencies must account for the details of each scheme, in particular the details of the artificial viscosity or the upwinding.

The Jacobi preconditioning, on the other hand, only deals with smoothing the high frequencies. Thus, it is mainly applicable to codes that use multigrid to accelerate the convergence. For a multigrid code it is essential that the high frequencies of the error be damped before one transfers the residual to a coarser grid. This prevents aliasing of the errors on the coarser grid. The Jacobi preconditioner allows for an improved damping of high frequencies. However, this only helps for the high frequencies in both the x and y Fourier directions. For a standard multigrid method it is necessary to also damp modes that are high frequency in one direction and low frequency in the other direction. Hence, Jacobi preconditioning is most useful when coupled with semi-coarsening.

Combining the low speed preconditioner with the Jacobi preconditioner (preconditioning-squared) automatically treats the matrix viscosity similar to the Jacobi method. In addition, calculating the condition number numerically verifies that it also satisfies the conditions for improved convergence at low Mach numbers. Hence this combination offers improved convergence rates compared with the original scheme over the entire subsonic flow range.

However, Jacobi preconditioning has several disadvantages. First the inverse of the preconditioning is not known analytically and so a 4×4 matrix must be numerically inverted at each grid point. In three dimensions this becomes a 5×5 matrix. In addition Jacobi preconditioning is most useful in conjunction with semi-coarsening. However, since semi-coarsening coarsens the grid in only one direction it has more grid points on the coarser grids than full-coarsening. Thus, in both two and three dimensions the coarse grid has half the points of the fine grid instead of a quarter (eighth) of the points in two (three) dimensions. This dramatically increases both the computing time and storage of the code. In particular this implies that only a V cycle and not a W-cycle can be used for the multigrid iteration. Semi-coarsening has the additional difficulty that it requires a choice of direction. This gives difficulties near junctures. It is also difficult to implement in a multi-block code where neighboring blocks may be semi-coarsened in different directions.

An extension to a point Jacobi preconditioner is to consider a line Jacobi method in the direction of the finest meshes [19]. This eliminates the need for semi-coarsening as the line-implicit method accomplishes the same purpose as the semi-coarsening. When one uses a Runge-Kutta time stepping method this is equivalent to using a one dimensional sweep from an ADI method every stage of the Runge-Kutta method. Hence, a block tridiagonal matrix needs to be inverted at each stage. This is costly both in terms of storage and computation time. For the standard ADI methods this can be eliminated by diagonalizing the sub-blocks.

Calculations using preconditioned-squared have been performed in [4, 19].

13 Viscous Preconditioning

We now consider preconditioners that are also appropriate for the Navier-Stokes equations. A non-symmetric candidate was presented before, (27). We now consider a symmetric positive definite preconditioner and so we need to write the (constant coefficient) Navier-Stokes equations as a symmetric system. We express the equations in terms of the variables $dw_N = (\frac{a}{\sqrt{\gamma}\rho}d\rho, du, dv, \frac{R\sqrt{\frac{\gamma}{\gamma-1}}}{a}dT)$. Using the Stokes hypothesis, $\lambda + \frac{2\mu}{3} = 0$ and ignoring the dissipation function Φ, we have

$$\mathbf{P}^{-1}(w_T)_t + \begin{pmatrix} u & \frac{a}{\sqrt{\gamma}} & 0 & 0 \\ \frac{a}{\sqrt{\gamma}} & u & 0 & \sqrt{\frac{\gamma-1}{\gamma}}a \\ 0 & 0 & u & 0 \\ 0 & \sqrt{\frac{\gamma-1}{\gamma}}a & 0 & u \end{pmatrix} (w_N)_x + \begin{pmatrix} v & 0 & \frac{a}{\sqrt{\gamma}} & 0 \\ 0 & v & 0 & 0 \\ \frac{a}{\sqrt{\gamma}} & 0 & v & \sqrt{\frac{\gamma-1}{\gamma}}a \\ 0 & 0 & \sqrt{\frac{\gamma-1}{\gamma}}a & v \end{pmatrix} (w_N)_y =$$

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \frac{4\mu}{3\rho} & 0 & 0 \\ 0 & 0 & \frac{\mu}{\rho} & 0 \\ 0 & 0 & 0 & \frac{\gamma\mu}{Pr\rho} \end{pmatrix} (w_N)_{xx} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \frac{\mu}{\rho} & 0 & 0 \\ 0 & 0 & \frac{4\mu}{3\rho} & 0 \\ 0 & 0 & 0 & \frac{\gamma\mu}{Pr\rho} \end{pmatrix} (w_N)_{yy} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{\mu}{3\rho} & 0 \\ 0 & \frac{\mu}{3\rho} & 0 & 0 \\ 0 & 0 & 0 & u \end{pmatrix} (w_N)_{xy} \quad .$$

Connecting to the previous set of variables dw_0 we have

$$\frac{\partial w_0}{\partial w_N} = \begin{pmatrix} \frac{1}{\sqrt{\gamma}} & 0 & 0 & \sqrt{\frac{\gamma-1}{\gamma}} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\frac{\gamma-1}{\sqrt{\gamma}} & 0 & 0 & \sqrt{\frac{\gamma-1}{\gamma}} \end{pmatrix} \qquad \frac{\partial w_N}{\partial w_0} = \begin{pmatrix} \frac{1}{\sqrt{\gamma}} & 0 & 0 & -\frac{1}{\sqrt{\gamma}} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \sqrt{\frac{\gamma-1}{\gamma}} & 0 & 0 & \frac{1}{\sqrt{\gamma(\gamma-1)}} \end{pmatrix} \quad .$$

The preconditioner (12) is given in w_N variables by

$$\mathbf{P}_N^{-1} = \begin{pmatrix} \frac{1+\beta^2[1+(\delta-1)(\gamma-1)]}{\gamma\beta^2} & 0 & 0 & -\frac{\sqrt{\gamma-1}\beta^2(\delta-1)+1}{\gamma\beta^2} \\ \frac{\alpha u}{\sqrt{\gamma}a\beta^2} & 1 & 0 & -\frac{\alpha u}{a\beta^2}\sqrt{\frac{\gamma-1}{\gamma}} \\ \frac{\alpha v}{\sqrt{\gamma}a\beta^2} & 0 & 1 & -\frac{\alpha v}{a\beta^2}\sqrt{\frac{\gamma-1}{\gamma}} \\ -\frac{\sqrt{\gamma-1}}{\beta^2[1+\delta(\gamma-1)]}\gamma\beta^2 & 0 & 0 & \frac{\beta^2[1+\delta(\gamma-1)]+\gamma-1}{\gamma\beta^2} \end{pmatrix} \quad .$$

For $\alpha = \delta = 0$ only the four corner elements are nontrivial. Extending this we can consider symmetric preconditioners with nontrivial entries in the four corners. So

$$\mathbf{P}_N = \begin{pmatrix} A & 0 & 0 & B \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ B & 0 & 0 & D \end{pmatrix} \quad .$$

We find that the symmetric preconditioner in the w_0 variables, (12), $\alpha = \delta = 0$, remains symmetric in the w_N variables with

$$A = 1 - \frac{1-\beta^2}{\gamma}, \quad B = -\frac{\sqrt{\gamma-1}(1-\beta^2}{\gamma}, \quad D = 1 - \frac{(\gamma-1)(1-\beta^2)}{\gamma} \quad .$$

Hence, this preconditioner is well-posed also for the semi-discrete Navier-Stokes equations as it is a symmetric incomplete parabolic system. One can also show stability when the time is discretized by a Runge-Kutta scheme using energy methods [18]. Since, the determinant is not changed by a similarity transform the determinant remains, $\det = \beta^2$ and thus the preconditioner becomes singular as $M \to 0$ unless β is limited as previously described.

We expect that β should depend not only on the Mach number but also on the viscous parameters. The main dependence is on the Reynolds number with a weaker dependence the Prandtl number.

Merkle [3, 20, 21] suggests replacing β^2 by $\beta^2 \min(1, R_\Delta^2)$. To generalize this we need to determine a cell-Reynolds number for the compressible equations in generalized coordinates. An extra complication is that R_Δ is itself affected by the preconditioning since it affects the inviscid speeds. Thus, we have an implicit definition in which β depends on R_Δ which itself might depend on β. A direct definition of β for Cartesian meshes is given in [3]. Another approach is to define R_Δ as the ratio of the inviscid time step to the viscous time step in each direction. This gives a directional cell-Reynolds number. As previously noted this leads to a recursive definition since the inviscid time step depends on β and the final β depend on the ratio of the inviscid to viscous time steps. However, one step of the recursive formula is sufficient to determine β.

We now consider the general preconditioner (12) for the one dimensional case and Fourier transform with Fourier variable k and $\theta = k\Delta x$. Define

$$A = \begin{pmatrix} u & \frac{a}{\sqrt{\gamma}} & 0 \\ \frac{a}{\sqrt{\gamma}} & u & \sqrt{\frac{\gamma-1}{\gamma}}a \\ 0 & \sqrt{\frac{\gamma-1}{\gamma}}a & u \end{pmatrix} \quad , \quad C = \frac{\mu}{\rho}\begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{4}{3} & 0 \\ 0 & 0 & \frac{\gamma}{Pr} \end{pmatrix} \quad .$$

We consider both the continuous and central difference cases simultaneously. Let

$$G = iAg_I - Cg_r, \qquad g_I = \begin{cases} k, \\ \frac{\sin(\theta)}{\Delta x} \end{cases} \quad , \qquad g_r = \begin{cases} k^2, & \text{continuous} \\ \frac{2(1-\cos(\theta))}{(\Delta x)^2}, & \text{discrete} \end{cases} \quad .$$

Since G is not a normal matrix the eigenvalues are not sufficient for stability. Nevertheless, they indicate the properties of the system. One obtains a cubic equation for the eigenvalues of the G. If $Pr = \frac{3}{4}$ we can factor the equation explicitly and find that one eigenvalue is given by

$$\lambda_0 = iug_I - \frac{4\mu}{3\rho}g_r \quad .$$

Let $\alpha_2 = \frac{4\mu}{3\rho}$, $\alpha_3 = \frac{\gamma\mu}{\rho Pr} = \gamma\alpha_2$. We then obtain a quadratic equation for the remaining eigenvalues $\lambda_\pm$:

$$\lambda^2 - \lambda\left\{\left[\beta^2(\gamma-1)(\delta-1)-1\right]\alpha_2 g_r + iug_I(1-\alpha+\beta^2)\right\} + \beta^2\left((a^2-u^2)g_I^2 - iug_I\alpha_3 g_r\right) = 0. \quad (32)$$

Note that now δ affects the eigenvalues. However, if β is small then the effect is only $O(\beta^2)$.

Solving (32) we get

$$2\lambda_\pm = \left[\beta^2(\gamma-1)(\delta-1)-1\right]\alpha_2 g_r + iug_I(\beta^2-\alpha+1) \pm \sqrt{s}$$

$$s = -u^2\left[(1-\alpha+\beta^2)^2 - 4\beta^2\right]g_I^2 - 4\beta^2 a^2 g_I^2 + \left[\beta^2(\gamma-1)(\delta-1)-1\right]^2\alpha_2^2 g_r^2$$

$$+ 2i\left\{2\gamma\beta^2 + (1-\alpha+\beta^2)\left[\beta^2(\gamma-1)(\delta-1)-1\right]\right\}\alpha_2 g_I g_r u \quad . \tag{33}$$

Define

$$R_\Delta = \frac{u\Delta x}{\alpha_2}, \quad \delta_1 = 1+\beta^2(\gamma-1)(1-\delta) \quad \hat{g}_r = \frac{2(1-\cos(\theta))}{\Delta x} \quad .$$

Then we can rewrite (33) as

$$\frac{2\Delta x}{\alpha_2}\lambda_\pm = -\delta_1\hat{g}_r + iR_\Delta g_I(1-\alpha+\beta^2) \pm \sqrt{s} \tag{34}$$

$$s = \delta_1^2\hat{g}_r{}^2 - R_\Delta^2 g_I^2\left[(1-\alpha+\beta^2)^2 + 4\frac{\beta^2}{M^2}(1-M^2)\right] + 2iR_\Delta\left[2\gamma\beta^2 - (1-\alpha+\beta^2)\delta_1\right]g_I\hat{g}_r \quad .$$

The real part of s can be forced to zero for some frequency. Thus, $\lambda_+ = \lambda_-$ if the imaginary part of s is zero, i.e. $2\gamma\beta^2 + (1-\alpha+\beta^2)\delta_1 = 0$. This involves only the preconditioning parameters and not the frequency. If $\delta = 1$ then $s = 0$ when $\beta^2 = \frac{1-\alpha}{2\gamma-1} \sim \frac{5}{9}(1-\alpha)$. When also $\alpha = 0$ (Choi-Merkle) there is a double eigenvalue for $\beta^2 = 5/9$. When $\delta = 0$ and $\alpha = 1$ then s cannot be zero and if $\delta = 0$ and $\alpha = 1$ s can be zero only when $\beta = 1$ (no preconditioning) and $\beta = 2.5$ (supersonic). Summarizing, for $\delta = 1$ there are frequencies for which we have a double eigenvalue. For $\delta = 0, \beta < 1$ we have distinct eigenvalues and so a strictly hyperbolic system which can be diagonalized. See [25] for more details.

To get a better idea of what happens in the viscous limit we expand (34) in powers of R_Δ through quadratic terms. One difficulty with the analysis is that $\hat{g}_r \ll g_I$ for small θ

while $g_I = 0$ at $\theta = \pi$ and so one cannot uniformly bound $\frac{\hat{g}_r}{g_I}$ or its inverse. Demanding that $|\lambda_-| = |\lambda_+|$ for an intermediate frequency shows that β should be proportional to $1/R_\Delta^2$ for R_Δ small. Merkle et. al. [21], based on a perturbation analysis [20], suggest that

$$\beta_V^2 = \beta^2 \max(1, 1/R_\Delta^2) \quad .$$

where β_V is the β appropriate for the viscous problem. With this choice one may be able to use a higher cutoff for β (see section 7). For a general mesh one needs to define R_Δ. Hence, it may be easier to use this formula separately in each direction, e.g. $\beta_{Vx}^2 = \beta_x^2 \min(1, K/R_{\Delta x}^2)$, and $\beta_V^2 = \max(\beta_{Vx}^2, \beta_{Vy}^2)$. Unrau and Zingg [38] chose

$$\beta_V^2 = \max(\beta^2, \frac{K}{R_\Delta}) \quad , \qquad R_\Delta = \frac{\rho u_\infty}{\mu_\infty \sqrt{\max(\xi_x^2 + \xi_y^2, \eta_x^2 + \eta_y^2)}} \quad .$$

This scales differently then the previous formula.

Hirsch and Hakimi [12] computed the solution for the driven cavity problem. They found that for Reynolds numbers below 1000 one had to increase the value of β from the inviscid formula. They present a figure to show the ratio of the optimal β to the inviscid β as a function of the reference Reynolds number. For $Re = 1$ they find that the optimal β^2 was 10,000 times larger than that given by $u^2 + v^2$. An approximate reading from their graph yields that β^2 depends on $Re^{4/3}$ for this specific example. For more general problems β should depend on the local cell Reynolds-number R_Δ and not the reference Reynolds number.

14 Differential Preconditioners

The preconditioners that we have described are all local preconditioners as they affect the variables only at a single node. As we have described, it is common for elliptic equations to consider preconditioners that couple neighboring nodes. We consider the linearized Euler equations with constant coefficients and the symmetrizing variables $w_0 = (\frac{dp}{\rho c}, u, v, S)$.

$$\text{Let,} \qquad Q = u\partial_x + v\partial_y, \qquad D = Q^2 - c^2(\partial_x^2 + \partial_y^2).$$

Then

$$\mathcal{L} = \begin{pmatrix} Q & c\partial_x & c\partial_y & 0 \\ c\partial_x & Q & 0 & 0 \\ c\partial y & 0 & Q & 0 \\ 0 & 0 & 0 & Q \end{pmatrix} \quad .$$

Define:

$$\mathbf{P_E} = \begin{pmatrix} Q & -c\partial_x & -c\partial_y & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad \mathbf{P_E^t} = \begin{pmatrix} Q & 0 & 0 & 0 \\ -c\partial_x & 1 & 0 & 0 \\ -c\partial_y & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad .$$

Then

$$\mathbf{P_E}\mathcal{L}\mathbf{P_E^t} = \begin{pmatrix} DQ & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & Q & 0 \\ 0 & 0 & 0 & Q \end{pmatrix} \quad ,$$

and so we have diagonalized $\mathcal{L}$ by a congruent transformation. But,

$$\mathbf{P_E}\mathcal{L}\mathbf{P_E}^t = \mathbf{P_E^{t^{-1}}}(\mathbf{P_E^t}\mathbf{P_E}\mathcal{L})\mathbf{P_E}.$$

So the congruent transform is similar to a preconditioning with a positive definite matrix $\mathbf{P_E^t}\mathbf{P_E}$. Alternatively, $(\mathbf{P_E^t}\mathbf{P_E})\mathcal{L}$ is similar to a diagonal matrix.

$$\mathbf{P_E^t}\mathbf{P_E} = \begin{pmatrix} Q & -c\partial_x & -c\partial_y & 0 \\ -c\partial_x & 1 + c^2\partial_x^2 & c^2\partial_x\partial_y & 0 \\ -c\partial_y & c^2\partial_x\partial_y & 1 + c^2\partial_y^2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

These transformations are independent of the flow regime as long as the preconditioner is nonsingular. As usual the entropy wave decouples with eigenvalue D and eigenvector

(0,0,0,1). The shear wave also has an eigenvalue D and the eigenvector is $(v_1, v_2, v_3, 0)$ where

$$Dv_1 = 0$$

$$D(\frac{\partial v_2}{\partial x} + \frac{\partial v_3}{\partial y}) = 0.$$

The other two 'acoustic' eigenvalues of $\mathbf{P}$ are $Q^2 \pm cQ\sqrt{\partial_x^2 + \partial_y^2}$ and the eigenvectors satisfy the pseudo-differential equation

$$\left[c(\partial_x^2 + \partial_y^2) \pm Q\sqrt{\partial_x^2 + \partial_y^2} \quad\right]\left(\frac{\partial v_3}{\partial x} - \frac{\partial v_2}{\partial y}\right) = 0$$
$$\mp\sqrt{\partial_x^2 + \partial_y^2} = \rho c\left(\frac{\partial v_2}{\partial x} + \frac{\partial v_3}{\partial y}\right).$$

We therefore have to show that the eigenvalues are all nonzero so that $\mathbf{P}$ is nonsingular. The operator D is just the potential operator i.e. for any variable w

$$Dw = (u^2 - c^2)w_{xx} + 2uvw_{xy} + (v^2 - c^2)w_{yy}.$$

For subsonic flow this is an elliptic operator and so invertible. For supersonic flow D is a hyperbolic operator. Q is a hyperbolic operator denoting convection along a streamline. Thus, given appropriate boundary conditions this too should be invertible. At a stagnation point Q is singular and so it is necessary to limit the values of u and v in the definition of Q so that they do not become too small in a neighborhood of the stagnation point. A similar smoothing is needed near the sonic line.

With both residual smoothing and $\mathbf{P_D}$ we increase the order of the system and so change the number of boundary conditions needed for the equation to be well posed. Instead these preconditioners should be used as a post-processor for the Euler or Navier-Stokes equations.

Thus, at each time step a residual is calculated based on one's favorite scheme. This gives a predicted value of the change in time, $\Delta w_{predicted}$. Then the boundary conditions are updated. We then operate on Δw with $\mathbf{P}$ using the boundary condition that $\Delta w_{corrected} = 0$, i.e. the boundary values calculated by the predictor are not altered.

If the partial derivatives are all replaced by central differences then everything still works since only formal multiplications were performed which use no properties of the derivative. Hence, a central difference approximation to the constant coefficient multidimensional Euler equations can be preconditioned so that it reduces to a set of scalar equations. Hence, high aspect ratio problems need to be considered only for a scalar equation.

References

[1] Abarbanel S, Gottlieb D, 1981. Optimal Time Splitting for Two- and Three-Dimensional Navier-Stokes Equations with Mixed Derivatives, *J. Comput. Phys.* 41:1-33.

[2] Allmaras S, 1993. Analysis of a Local Matrix Preconditioner for the 2-D Navier-Stokes Equations. *AIAA 11th Computational Fluid Dynamics Conference, AIAA pap* 93-3330.

[3] Choi YH, Merkle CL, 1993. The Application of Preconditioning to Viscous Flows, *J. Comput. Phys.* 105:207-223.

[4] Darmofal DL, Schmid PJ, 1996. The Importance of Eigenvectors for Local Preconditioners of the Euler Equations. *J. Comput. Phys.* 127:346-362.

[5] Darmofal DL, Siu K, 1998. A Robust Locally Preconditioned Multigrid Algorithm for the Euler Equations, *AIAA pap* 98-2428.

[6] Darmofal DL, van Leer B, 1998. Local Preconditioning: Manipulating Mother Nature to Fool Father Time. *Computing the Future II: Computational Fluid Dynamics and Transonic Flow* ed. D. Caughey and M. Hafez.

[7] Edwards JR, Roy CJ, 1998. Preconditioned Multigrid Methods for Two-Dimensional Combustion Calculations at All Speeds, *AIAA Journal* 36:185-192.

[8] Edwards JR, Liou MS, 1998. Low Diffusion Flux-Splitting Schemes for Flows at All Speeds, *AIAA Journal* to appear.

[9] Godfrey AG, 1994. Steps Toward a Robust Preconditioning, *AIAA pap* 94-0520.

[10] Guillard H, Viozat C, 1998. On the Behavior of Upwind Schemes in the Low Mach Number Limit, *Computers and Fluids* to appear.

[11] Gustafsson B, 1987. Unsymmetric Hyperbolic Systems and the Euler Equations at Low Mach Numbers. *J. Scientific Computing* 2:123-136.

[12] Hirsch Ch, Hakimi N, 1995. Preconditioning Methods for Time-marching Navier-Stokes Solvers, *Solution Techniques for Large-Scale CFD Problems* ed. WG Habashi, John Wiley & Sons, 334-353.

[13] Hirsch Ch, Hakimi N, 1995. Preconditioning Methods for Very Low Speed Computations with Time-Marching Navier-Stokes Solvers, *First Asian Computational Fluid Dynamics Conference.*

[14] Hauke G, Hughes TJR, 1998. A Comparative Study of Different Sets of variables for Solving Compressible and Incompressible Flows, *Comput. Methods Appl. Mech. Engng* 153:1-44.

[15] Jameson A, Schmidt W, Turkel E, 1981. Numerical Solutions of the Euler Equations by a Finite Volume Method using Runge-Kutta Time-Stepping Schemes, *AIAA pap 81-1259.*

[16] Jameson A, 1995. Analysis and Design of Numerical Schemes for Gas Dynamics I: Artificial Diffusion, Upwind Biasing, Limiters and their Effect on Accuracy and Multigrid Convergence, *International Journal of Computational Fluid Dynamics* 4:171-218.

[17] Jespersen D, Pulliam T, Buning P, 1997. Recent Enhancements to OVERFLOW, *AIAA pap* 97-0644.

[18] Levy D, Tadmor E, 1998. From Semidiscrete to Fully Discrete: Stability of Runge-Kutta Schemes by the Energy Method, *SIAM Review* 40:40-73.

[19] Mavriplis D, 1997. Multigrid Strategies for Viscous Flow Solvers on Anisotropic Unstructured Meshes, *AIAA pap* 97-1952.

[20] Merkle CL, 1995. Preconditioning Methods for Viscous Flow Calculations, *Computational Fluid Dynamics Review 1995* 419-436, ed. Hafez M, Oshima K, Wiley, Chichester, England.

[21] Merkle CL, Sullivan JY, Buelow PEO, Venkatewaran S, 1998. Computation of Flows with Arbitrary Equations of State, *AIAA J.* 36:515-521.

[22] Mesaros L, Roe P, 1995. Multidimensional Fluctuation Splitting Schemes Based on Decomposition Methods, *AIAA 12th Computational Fluid Dynamics Conf.* 582-591.

[23] Muradoglu M, Caughey DA, 1998. Implicit Multigrid Solution of the Preconditioned Multi-dimensional Euler Equations, *AIAA pap* 98-0114.

[24] Pierce NA, Giles MB, 1997. Preconditioned Multigrid Methods for Compressible Flow Calculations on Stretched Meshes, *J. Comput. Phys.* 136:425-445.

[25] Rico D, 1997. Numerical Methods to Accelerate the Computation of Steady State Solutions of the Navier-Stokes Equations. Masters thesis, Dept. Math., Tel Aviv University.

[26] Roe P, Turkel E, 1996. The Quest for Diagonalization of Differential Systems, *Workshop on Barriers and Challenges in Computational Fluid Dynamics*, ed. V. Venkatakrishnan, MD Salas, SR Chakaravarthy, Kluwer Publisher 351-370.

[27] Saint Requier C, Duquesne N, 1998. Application of a Preconditioning Method for Aerodynamic Problem, *Proc. 38th Israel Annual Conference on Aerospace Sciences*, Omanuth Press Ltd. 45.

[28] Swanson, RC, Turkel E, 1985. A Multistage Time-stepping Scheme for the Navier-Stokes Equations, *AIAA Pap* 85-0035.

[29] Swanson, RC, Turkel E, 1987. On Central Difference and Upwind Schemes, *J. Comput. Phys.* 101:292-306.

[30] Swanson, RC, Radespiel R, Turkel E, 1997. On Some Numerical Dissipation Schemes *13th AIAA Computational Fluid Dynamics Conference, AIAA pap* 97-1945.

[31] Turkel E, 1987. Preconditioned Methods for Solving the Incompressible and Low Speed Compressible Equations, *J. Comput. Phys.* 72:277-298.

[32] Turkel E, 1993. A Review of Preconditioning Methods for Fluid Dynamics, *Applied Numerical Mathematics*, 12:257-284.

[33] Turkel E, Fiterman A, van Leer B, 1994. Preconditioning and the Limit to the Incompressible Flow Equations, *Computing the Future: Frontiers of Computational Fluid Dynamics 1994* ed. Caughey DA, Hafez MM, Wiley Publishing, 215-234.

[34] Turkel E, Vatsa VN, Radespiel R, 1996. Preconditioning Methods for Low Speed Flow, *AIAA Pap* 96-2460.

[35] Turkel E, Radespiel R, Kroll N, 1997. Assessment of Two Preconditioning Methods for Aerodynamic Problems, *Computers and Fluids* 613-634.

[36] Turkel E, 1997. Preconditioning-Squared Methods for Multidimensional Aerodynamics, *13th AIAA Computational Fluid Dynamics Conference, AIAA pap* 97-2025.

[37] Tweedt DL, Chima RV, Turkel E, 1997. Preconditioning for Numerical Simulation of Low Mach Number Three-Dimensional Viscous Turbomachinery Flows, *AIAA pap* 97-1828.

[38] Unrau D, Zingg DW, 1997. Viscous Airfoil Computations Using Local Preconditioning, *AIAA pap* 97-2027.

[39] van Leer B, Lee WT, Roe PL, 1991. Characteristic Time-Stepping or Local Preconditioning of the Euler Equations, *AIAA Pap* 91-1552.

[40] van Leer B, Mesaros L, Tai CH, Turkel E, 1995. Local Preconditioning in a Stagnation Point, *AIAA 12th Computational Fluid Dynamics Conference*, 88-101.

[41] Weiss JM, Smith WA, 1995. Preconditioning Applied to Variable and Constant Density Flows, *AIAA J.* 33:2050-2057.