



AIAA 2002-2951

Optimal Multigrid Convergence by Elliptic/Hyperbolic Splitting

Hiroaki Nishikawa and Bram van Leer

W. M. Keck Foundation Laboratory for Computational Fluid Dynamics

Department of Aerospace Engineering

The University of Michigan, Ann Arbor, MI 48109

**32nd AIAA Fluid Dynamics Conference
June 24-27 2002, St. Louis. MO**

For permission to copy or republish, contact the copyright owner named on the first page. For AIAA-held copyright, write to AIAA, Permissions Department, 1801 Alexander Bell Drive, Suite 500, Reston, VA 20191-4344

Optimal Multigrid Convergence by Elliptic/Hyperbolic Splitting

Hiroaki Nishikawa* and Bram van Leer†

W. M. Keck Foundation Laboratory for Computational Fluid Dynamics

Department of Aerospace Engineering

The University of Michigan, Ann Arbor, MI 48109

Abstract

We describe a multigrid method for solving the steady Euler equations in $O(N)$ operations, where N is the number of unknowns, based on an elliptic/hyperbolic decomposition achieved by local preconditioning. The splitting allows the embedded advection equations to be treated with streamwise semicoarsening rather than full coarsening, which would not be effective. A simple 2-D numerical computation is presented as proof of concept. A convergence study indicates the split method has complexity $N^{0.97}$ over a wide range of grid spacings and Mach numbers, while the use of full coarsening for all equations makes the complexity deteriorate to $N^{1.44}$.

Introduction

We describe a multigrid method for solving the steady Euler equations in $O(N)$ operations, where N is the number of unknowns, based on an elliptic/hyperbolic decomposition achieved by local preconditioning. This method means a significant simplification and computational savings compared to the method of Darmofal and Siu,¹ who also used local preconditioning in order to reach $O(N)$ complexity. Their multigrid method required semicoarsening in all directions, which is computationally expensive, especially for three-dimensional problems. Isotropic semicoarsening, though robust, was understood by those and the current authors to be “overkill”:¹⁰ *for optimal convergence, semicoarsening need only be applied to the embedded advection equations, and only in the flow direction.* Such “streamwise” semicoarsening requires the splitting of the residual into advective and acoustic components, which happens to be a special benefit of the Van Leer-Lee-Roe¹¹ local preconditioning. Below we use local preconditioning for removing local stiffness as well as for equation decomposition, and we compare the rate of convergence to subsonic flow solutions for the new

“elliptic/hyperbolic multigrid” and standard (full-coarsening) multigrid methods.

Toward $O(N)$ Complexity

In a long sequence of papers running from 1989⁹ to 1999¹ Van Leer and collaborators, later joined by Darmofal and collaborators, demonstrated that it is possible to achieve convergence to steady Euler solutions in $O(N)$ operations using an explicit single-grid scheme augmented by multigrid relaxation, provided:

1. local preconditioning is used to remove the local stiffness caused by the spread of the characteristic speeds;
2. the single-grid scheme is designed to efficiently remove all removable high-frequency errors (some errors are not removable when the flow is grid-aligned);
3. a multigrid method more powerful than standard full-coarsening multigrid is used to overcome the alignment problem.

An overview of this research effort has been given by Van Leer and Darmofal;¹⁰ in this paper we shall only present supplemental information needed to understand the current method. This includes a discussion and comparison of the benefits of isotropic and streamwise semicoarsening, and a discussion of elliptic/hyperbolic splitting.

In this paper we restrict ourselves to the 2-D Euler equations and their numerical solution. There is a complication when going to 3-D Euler, caused by the inability to separate the advection equation for streamwise vorticity from the acoustic equations by algebraic preconditioning;^{7,11} we will return to this subject in the last section.

Baseline Euler Solver

Our aim is the efficient computation of steady-state solutions to the Euler equations in two dimensions,

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} + \frac{\partial \mathbf{G}}{\partial x} = \mathbf{0}, \quad (1)$$

*Postdoctoral Researcher, Member AIAA.

†Professor, Fellow AIAA.

Copyright © 2002 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho u H \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho v H \end{bmatrix}, \quad (2)$$

where ρ is the density, u and v are the velocity components in the x and y direction, respectively, and p is the static pressure. The specific total energy and enthalpy are given by

$$E = \frac{1}{\gamma - 1} \frac{p}{\rho} + \frac{1}{2}(u^2 + v^2) \quad (3)$$

$$H = \frac{\gamma}{\gamma - 1} \frac{p}{\rho} + \frac{1}{2}(u^2 + v^2). \quad (4)$$

We discretize this system with a cell-centered finite-volume algorithm on a structured quadrilateral grid. The cell-residual \mathbf{Res} is defined as a numerical approximation of the integral over the cell of the spatial operator, and can be written as

$$\mathbf{Res} = - \sum_{k=1}^4 \Phi_k \Delta s_k, \quad (5)$$

where Δs_k is the length of cell face k and Φ_k is the outward numerical flux normal to the face. In this work, we use Roe's approximate Riemann solver, with a modification required for stability in the presence of preconditioning, and for preserving solution accuracy at low Mach numbers:

$$\Phi_k = \frac{1}{2} [\Phi(\mathbf{U}_L) + \Phi(\mathbf{U}_R)] - \frac{1}{2} \hat{\mathbf{P}}^{-1} |\hat{\mathbf{P}} \hat{\mathbf{A}}| (\mathbf{U}_R - \mathbf{U}_L). \quad (6)$$

Here $\hat{\mathbf{P}}$ is a preconditioning matrix and $\hat{\mathbf{A}}$ is the flux Jacobian, each evaluated at the cell face using Roe's average state. To achieve second-order accuracy, Van Leer's κ -scheme⁸ is employed with $\kappa = 0$ and without limiting. Note that solutions will not be monotonic near shocks for this choice of κ ; to the test problem used below this does not apply.

The residual defined as above drives the evolution of the cell-average $\bar{\mathbf{U}}$ of \mathbf{U} according to the preconditioned equation

$$\frac{\partial \bar{\mathbf{U}}}{\partial t} = \frac{1}{S} \mathbf{P}(\bar{\mathbf{U}}) \mathbf{Res}, \quad (7)$$

where S is the area of the cell.

To reach a steady state we integrate this system in time using a multistage time-stepping optimized for high-frequency damping, as required for use in multigrid relaxation. The particular scheme employed here is the 4-stage optimized smoothing scheme developed for $\kappa = 0$ by Lynn and Van Leer;^{2,3} it has the form

$$\bar{\mathbf{U}}^{(1)} = \bar{\mathbf{U}}^n + \alpha^{(1)} \mathbf{P}^n \mathbf{Res}^n, \quad (8)$$

$$\bar{\mathbf{U}}^{(2)} = \bar{\mathbf{U}}^n + \alpha^{(2)} \mathbf{P}^{(1)} \mathbf{Res}^{(1)}, \quad (9)$$

$$\bar{\mathbf{U}}^{(3)} = \bar{\mathbf{U}}^n + \alpha^{(3)} \mathbf{P}^{(2)} \mathbf{Res}^{(2)}, \quad (10)$$

$$\bar{\mathbf{U}}^{n+1} = \bar{\mathbf{U}}^n + \alpha^{(4)} \mathbf{P}^{(3)} \mathbf{Res}^{(3)}. \quad (11)$$

For $i = 1, 2, 3, 4$ we have

$$\alpha^{(i)} = \beta^{(i)} \Delta t / S, \quad (12)$$

with

$$\Delta t = 2\nu S / \sum_{k=1}^4 q_k ds_k; \quad (13)$$

here q_k is the flow-velocity component normal to cell face k , the CFL-number ν has the value 1.4008, and $\beta^{(1)} = 0.1299$, $\beta^{(2)} = 0.2940$, $\beta^{(3)} = 0.5604$, $\beta^{(4)} = 1.0000$. In the present work, the preconditioning matrix \mathbf{P} will be a part of the full preconditioning matrix: either the part that produces the elliptic component of the residual or the part that produces the hyperbolic component of the residual*.

Elliptic/Hyperbolic Splitting Approach

For subsonic flows, the steady Euler equations are known to have a mixed behavior: elliptic for the acoustic components, hyperbolic for the advective components. Conventional multigrid methods based on full coarsening are known to work pre-eminently for isotropic elliptic problems, being capable of achieving $O(N)$ -type convergence, but not for hyperbolic, advection-dominated, problems. In the latter case a remedy is to coarsen the grid only in the advection direction, while retaining the original resolution in the cross-stream direction, i.e. semi-coarsening.

Our approach is then to separate the two problem-types at the residual level, and apply the optimal strategy to each residual component: full coarsening to the elliptic part, streamwise semicoarsening to the hyperbolic part (see Figure 3). Note that this greatly simplifies the grid structure and therefore the multigrid coding compared with the isotropic semicoarsening of Mulder (see Figure 2), which will be discussed in detail in the next section.

To decompose the residual, we make use of the Van Leer-Lee Roe (VLR)¹¹ preconditioning matrix corrected for the effect of the cell-aspect ratio \mathcal{AR} ;² its generic form, based on the symmetrized Euler equations (with state vector differential $(dp/(\rho a), du, dv, ds)^T$, s denoting entropy), is

$$\mathbf{P}_{VLR} = \begin{bmatrix} \frac{\tau}{\beta^2} M^2 & -\frac{\tau}{\beta^2} M & 0 & 0 \\ -\frac{\tau}{\beta^2} M & \frac{\tau}{\beta^2} + 1 & 0 & 0 \\ 0 & 0 & \tau & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (14)$$

where M is the local Mach number, $\beta = \sqrt{|M^2 - 1|}$, and

$$\tau = \frac{\sqrt{1 - \min(1/M^2, M^2)}}{\beta + \mathcal{AR}}. \quad (15)$$

*This is true only in the update step. In the formulation of the numerical flux (6), the full matrix is used.



Fig. 1 Full coarsening.




Fig. 2 Mulder's isotropic semi-coarsening.

The elliptic/hyperbolic decomposition is given by

$$\mathbf{P}_{VLR} = \mathbf{P}_{VLR}^{ell} + \mathbf{P}_{VLR}^{hyp}, \quad (16)$$

where

$$\mathbf{P}_{VLR}^{ell} = \begin{bmatrix} \frac{\tau}{\beta^2} M^2 & -\frac{\tau}{\beta^2} M & 0 & 0 \\ -\frac{\tau}{\beta^2} M & \frac{\tau}{\beta^2} & 0 & 0 \\ 0 & 0 & \tau & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (17)$$

$$\mathbf{P}_{VLR}^{hyp} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (18)$$

Applying \mathbf{P}_{VLR}^{ell} and \mathbf{P}_{VLR}^{hyp} to the residual gives the elliptic and hyperbolic part of the residual, respectively. Decomposed on the finest grid, each residual is sent to the corresponding sequence of coarse grids: the elliptic part to the fully coarsened sequence, the hyperbolic part to the semicoarsened sequence (cf. Figure 3).

For the simple test problem we consider in this study we used the simplest possible restriction and prolongation strategies: restriction by conservation over four/two cells for the elliptic/hyperbolic part, and prolongation by bilinear/linear interpolation. The acoustic and advective residuals were treated completely independently throughout the multigrid cycle; only at the finest grid level, where the elliptic and hyperbolic grids are identical, solution corrections for the different components were combined into a solution update. The convergence results show that this worked; for more complex problems we would expect the need for tighter coupling between the operators acting on the elliptic and hyperbolic residual components.




Fig. 3 The grid sequences for the new approach: full/semi-coarsening for the elliptic/hyperbolic residual part.

Semicoarsening

Multigrid relaxation on a collection of semicoarsened grids was developed by Mulder⁶ for increased robustness, in particular, when the flow is aligned over appreciable portions of the grid. In that case it is neither possible nor desirable for a single-grid scheme to damp error modes with high frequencies in the stream-normal direction for advected quantities such as entropy. The standard full-coarsening multigrid process then suffers because streamwise low-frequency errors coupled to those modes get averaged out during full coarsening and are not addressed at all. This slows down overall convergence and makes $O(N)$ complexity elusive. Streamwise semicoarsening will make the multigrid process effective in removing streamwise low-frequency error modes from the advected quantities. Since it may not be known in advance if, where, and how the flow aligns with the grid, Mulder reasoned that a robust, general-purpose strategy would be to use contributions from a matrix of grids semicoarsened in all

possible directions (see Figure 2). This will catch all high-frequency advective errors, while also being a powerful strategy for removing acoustic error modes. Thus, there is no need to decompose the residual into acoustic (elliptic) and advective (hyperbolic) components. Furthermore, because of the powerful one-dimensional mode-removal in the corrections on semi-coarsened grids, the demands on the single-grid scheme regarding smoothing are significantly reduced: only error modes with high frequencies in *all* directions need to be damped on each grid. Multistage marching schemes can be designed to do so very effectively, provided the residuals first are treated with local preconditioning. This clusters the eigenvalues of the spatial operator and makes it possible to eliminate the corresponding error modes regardless of the Mach number.

Although powerful, such isotropic semicoarsening is computationally too expensive. Compared to single-grid relaxation, the isotropic semicoarsening adds a factor 4 to the work in 2D and 8 in 3D which are significantly larger than the corresponding factors $\frac{4}{3}$ and $\frac{8}{7}$ for full coarsening. On the other hand, the proposed multigrid method adds no more than a factor $\frac{5}{3}$ to the work in 2D and $\frac{58}{35}$ in 3D (ignoring the effort of finding the flow-aligned cell strings). Coding complexity as compared to isotropic semicoarsening is also significantly reduced; in fact, such isotropic semicoarsening seems virtually unfeasible on unstructured grids, while the streamwise semicoarsening is feasible as demonstrated by Mavriplis,⁵ who selects strings of cells in order to apply line solvers[†].

Numerical Results

Our test case is flow over a smooth 5%-high bump in the interval $0 \leq x \leq 1$ described by $y = 0.05 \sin(\pi x)$ (no stagnation points), for which many results are available in the literature.^{1,4} The flow is almost perfectly aligned with the grid and conventional multigrid methods fail to achieve $O(N)$ -type convergence.

The domain is defined by $-1 \leq x \leq 3.5$ and $0 \leq y \leq 2$; the boundary at $y = 0$ is treated as solid and the other ones as open. At the boundaries, a ghost-cell approach is used: reflection of the normal velocity while keeping all other state variables the same across the solid wall; prescribing freestream values across open boundaries. The boundary fluxes are then determined by the same Riemann solver as used in the interior. The grid is a structured quadrilateral grid with modest compression toward the wall. A sample grid and solution are shown in Figures 4 and 5.

[†]A streamwise line solver definitely is an alternative to a multistage advection scheme.




Fig. 4 A 64x32 Grid




Fig. 5 Mach contours on the 128x64 grid; $M_\infty = 0.1$.




Fig. 6 Convergence histories for full coarsening ($M_\infty = 0.1$). Solid: 64x32; dash: 128x64; dash-dot: 256x128; dot: 512x256.

A numerical solution is taken to be converged when all L_1 norms of the components of the residual, viz. mass, momentum and energy, are reduced five orders of magnitude from their initial values.

For the multigrid relaxation an F-cycle is chosen with a coarsest grid of 8×4 cells; note that the streamwise semicoarsening (essentially in the x -direction) keeps the number of cells in the y direction the same for all grids. A work unit is defined as the amount of work required to evaluate the full residual on the finest grid. The total number of work units presented in the following results includes four initial relaxations on the finest grid prior to the multigrid cycles.

The convergence histories for various grids (64x32, 128x64, 256x128, 512x256) are shown in Figures 6 and 7, comparing the results of full-coarsening F-cycle multigrid with a preconditioned four-stage

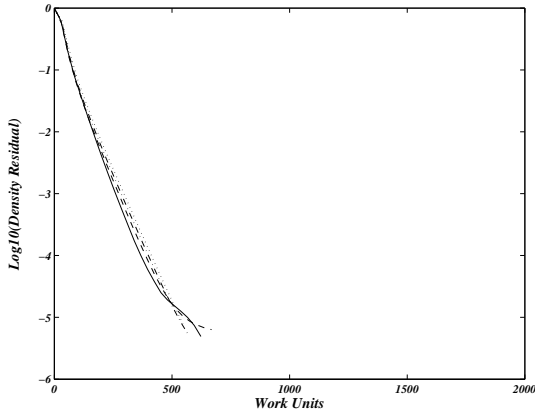


Fig. 7 Convergence histories for the new decomposition approach ($M_\infty = 0.1$). Solid: 64x32; dash: 128x64; dash-dot: 256x128; dot: 512x256.




Fig. 9 Convergence histories for the new approach on 512x256 grid for different Mach numbers. Solid: $M_\infty = 0.1$; dash: $M_\infty = 0.3$; dash-dot: $M_\infty = 0.5$.




Fig. 8 Convergence histories for full coarsening on 512x256 grid for different Mach numbers. Solid: $M_\infty = 0.1$; dash: $M_\infty = 0.3$; dash-dot: $M_\infty = 0.5$.

Jacobi relaxation and the new elliptic/hyperbolic decomposition method that uses the same F-cycle and four-stage Jacobi relaxation for each subproblem. It is evident that the convergence is independent of the grid size in the new method, whereas the conventional method requires an increasing amount of work to converge as the grid gets finer.

In Figures 8 and 9 convergence histories for the two methods are plotted for different Mach numbers. For both methods the plots coincide, showing independence of Mach number. This is a well-known property of the VLR preconditioner. Without preconditioning, this is not possible, as reported by many researches.¹

Lastly, actual CPU times for the two methods are compared in Figure 10, measured when the programs are run on a single Pentium III (500MHz) processor. The curves were determined by least-squares linear fits in log-log space. The plots show that the convergence rate is $O(N)$ for the new method, while nearly $O(N^{1.5})$ for full coarsening. Thus, for sufficiently fine grids the new method outperforms the

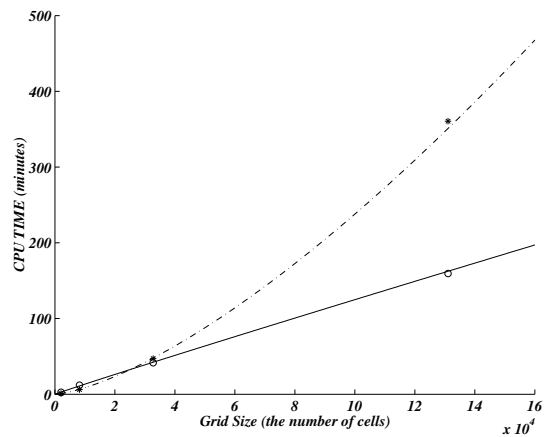


Fig. 10 CPU time versus grid size for $M_\infty = 0.1$. Solid: new approach $CPU = 1.7 \times 10^{-3} \times N^{0.971}$; dash: full coarsening $CPU = 1.53 \times 10^{-5} \times N^{1.44}$ where N is the number of cells.

standard method.

Conclusions and Future Work

The numerical results presented provide the proof of concept for the current elliptic/hyperbolic multi-grid method. Putting it differently: had the proposed strategy not worked for this problem, we might as well forget about making it work for more complex problems. The great advantage of the adopted test problem was that it was *a priori* known which strings of cells had to be semicoarsened. Thus, the investment in coding streamwise coarsening was minimal.

Obviously, we are just at the start of the development of a more robust method. The following points need to be attended in the future:

1. Development of a tightly-coupled restriction-prolongation strategy for the elliptic and hyperbolic residual components.
2. Formulation of the method on unstructured

grids, where one has to determine “on the go” which cells are strongly coupled in the flow direction. Here we can learn from previous experience by Mavriplis,⁵ who selects strings of cells in order to apply line solvers.

3. Overcoming the problem of reduced acoustic-residual smoothing near a sonic line or for large cell-aspect ratios (these problems do not arise with isotropic semicoarsening);
4. Extending the method to three dimensions. As mentioned above, one of the hyperbolic components can not be decoupled from the acoustic residual with conventional preconditioning; decoupling is only achieved if the preconditioning operator include differentiation entries. For certain discretizations (multi-D residual-distribution schemes) it may actually be possible to achieve this numerically without leaving the conservation form of the equations. Meanwhile, it remains to be seen how seriously multi-grid convergence would be effected if the decomposition were not complete.

Acknowledgement

Bram van Leer acknowledges valuable exploratory discussions with David Darmofal during a visit to the latter at Texas A&M in December 1997. This work was funded in part by a grant from the U.S. Civilian Research and Development Foundation.

References

- ¹D.L. Darmofal and K. Siu. A robust multigrid algorithm for the Euler equations with local preconditioning and semicoarsening. *Journal of Computational Physics*, 151:728–756, 1999.
- ²J. F. Lynn. *Multigrid Solution of the Euler Equations with Local Preconditioning*. PhD thesis, University of Michigan, 1995.
- ³J. F. Lynn and B. van Leer. Multi-stage schemes for the Euler and Navier-Stokes equations with optimal smoothing. AIAA Paper AIAA 93-3355-CP, 1993.
- ⁴J. F. Lynn, B. van Leer, and D. Lee. Multigrid solution of the Euler equations with local preconditioning. In *Fifteenth International Conference on Numerical Methods in Fluid Mechanics*. Springer, 1996.
- ⁵D. J. Mavriplis. On convergence acceleration techniques for unstructured meshes. AIAA Paper AIAA 98-2966, 1998.
- ⁶W. A. Mulder. A new approach to convection problems. *Journal of Computational Physics*, 83:303–323, 1989.
- ⁷P. L. Roe and E. Turkel. The quest for diagonalization of differential systems. In *Barriers and Challenges in Computational Fluid Dynamics*, pages 251–369. Kluwer, 1998.
- ⁸B. van Leer. Upwind-difference methods for aerodynamic problems governed by the Euler equations. In *Large-Scale Computations in Fluid Mechanics, Lectures in Applied Mathematics*, volume 22, Part 2, pages 327–336, 1985.
- ⁹B. van Leer. Euler solvers for transonic applications. In *NASA Transonic Symposium: Theory, Application, and Experiment*, 1989.

¹⁰B. van Leer and D. Darmofal. Steady Euler solutions in $O(N)$ operations. In E. Dick and K. Riemsdagh, editors, *Proceedings of the Sixth European Multigrid Conference, Gent, Belgium, September 27-30, 1999*. Springer, 2000.

¹¹B. van Leer, W. T. Lee, and P. L. Roe. Characteristic time-stepping or local preconditioning of the Euler equations. AIAA Paper 91-1552-CP, 1991.