Analysis of "p"-Multigrid for Continuous and Discontinuous Finite Element Discretizations

Brian T. Helenbrook*
Dept. Mech. and Aero. Eng.
Clarkson University
Potsdam, NY 13699-5725

Dimitri Mavriplis†
National Institute of Aerospace
144 Research Drive
Hampton, VA 23666

Harold L. Atkins‡
NASA Langley Research Center
MS 128
Hampton, VA 23681-2199

Analysis and numerical experiments examining the behavior and performance of p-multigrid (p = polynomial degree) for solving hp-finite element (FEM) discretizations are presented. We begin by demonstrating the mesh and order independent properties of p-multigrid when used to solve a C^0 continuous FEM discretization of the Laplace equation. We then apply p-multigrid to both continuous and discontinuous FEM discretizations of the convection equation. Although 1D Fourier analysis predicts that mesh independent results should be possible for both discretizations, in 2D the results are sensitive to both the mesh resolution and the degree of polynomial approximation. Examining the solutions, we find that for both discretizations, the slowest converging mode is long-wavelength along the streamwise direction and short wavelength normal to this direction. Because of the isotropic coarsening of p-multigrid, this mode is not damped on coarse levels.

Introduction

p-multigrid is an iterative algorithm in which systems of equations arising from compact, high-order finite element discretizations, such as spectral/hp or discontinuous Galerkin (DG) formulations, are solved by recursively iterating on solution approximations of different polynomial order. For example, to solve equations derived using a polynomial approximation order of 4, the solution can be iterated on at an approximation order of $p = 4$, 2, and 1. When a low order is reached, i.e. $p = 1$ or 0, a conventional multi-grid (grid coarsening) algorithm can be applied to solve for the low-order components of the solution. The p component of this algorithm was proposed by Rønquist and Patera[1] and analyzed by Maday and Munoz[2] for a one-dimensional, Galerkin spectral element discretization of the Laplace equation. Helenbrook[3] combined p-multigrid with standard low-order multigrid and applied it to an unstructured streamwise-upwind-Petrov-Galerkin (SUPG) discretization of the incompressible Navier-Stokes equations.

This work provides a detailed examination of the performance of p-multigrid. We begin by presenting results for a continuous hp-FEM discretization of the Laplace equation in two dimensions. Using this result as the baseline, we then examine SUPG and DG discretizations of the convection equation. Results are obtained using a 5-stage Runge-Kutta method for a relaxation scheme.[4] For DG, we also examine element Jacobi relaxation.[5] Fourier analysis is used to explain the results of the calculations.

Laplace Equation

We begin by examining the performance of p-multigrid when applied to a FEM discretization of the Laplace equation. For second-order differential equations, it is natural to use a C^0 continuous basis, therefore we only obtain results for a continuous Galerkin formulation. The results of Rønquist and Patera[1] have shown that in 1D, p-insensitive convergence rates can be obtained. To verify that this is the case in

*Assistant Professor, Department Mechanical and Aerospace Engineering
†Research Fellow, National Institute of Aerospace
‡Senior Research Scientist, Computational Modeling and Simulation Branch

two dimensions, numerical results are obtained using a mesh of triangles and the hierarchical triangular basis developed by Dubiner.[6]

The problem is solved on the unit square, $(x, y) \in [0, 1] \times [0, 1]$. A homogenous Dirichlet boundary condition is enforced on the left side, periodicity is enforced at the top and bottom, and a homogeneous Neumann boundary condition is enforced at the right side. The exact solution is simply $u(x, y) = 0$. Initial conditions are given by

$$u = (\sin(2\pi x) + \sin(100\pi x))\,(\sin(2\pi y) + \sin(100\pi y)) \tag{1}$$

These initial conditions have both high and low wavenumber components to ensure that the multigrid algorithm can damp the full range of error modes. The mesh is a structured Cartesian mesh with each square element subdivided to create two triangles. Resolutions of 2×2 to 64×64 are studied. Basis sets having a maximum polynomial degree, $p = 1$, 2, and 4 are studied.

The multigrid algorithm begins at the highest approximation order and halves the approximation order each level. Residuals are transfered to the lower-order spaces by simply eliminating those components resulting from the integration of the governing equation with the higher-order polynomials. Since the low-order space is a subset of the high-order space, this can always be done independent of the form of the polynomial basis functions. Once a linear approximation level has been reached, standard geometric multigrid[7] is used. A V-cycle is used to drive the iterations. This cycle includes both the different p levels and the geometric meshes. For example, starting from a 4×4 mesh with $p = 4$, the sequence would be: $(p = 4 : 4 \times 4)$, $(p = 2 : 4 \times 4)$, $(p = 1 : 4 \times 4)$, $(p = 1 : 2 \times 2)$, $(p = 1 : 1 \times 1)$ and then back again.

At each coarsening stage of the V-cycle, a Runge-Kutta (RK) scheme[4] is used as an iterative smoother. This smoother is an evolution in "pseudo-time" of the following equation

$$M\frac{\partial \vec{u}}{\partial \tau} + K\vec{u} = 0 \tag{2}$$

where M is the mass matrix and K is the stiffness matrix from the higher-order discretization. $\vec{u}$ is the vector of coefficients describing the solution and τ is the pseudo-time. To avoid global inversion of the mass matrix, mass-lumping is used; the mass matrix is approximated such that $M\vec{u}$ is exact for any $\vec{u}$ in the space formed using polynomials of degree $p - 1$. This allows the matrix to be inverted using only local operations. The mass-lumping operation procedure will be described in a future paper. The time step for the RK scheme is proportional to h^2/p^4 where h is the element length. This is approximately the inverse of the maximum eigenvalue of equation 2.

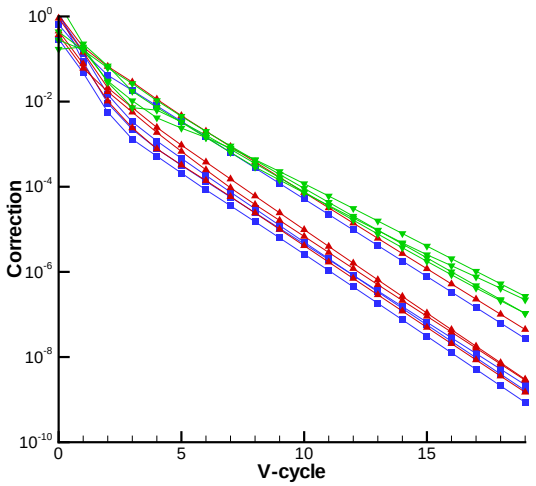

Fig. 1 Convergence results. (blue square: p=1, 8x8, 16x16, 32x32, 64x64), (red up-triangle: p=2, 4x4, 8x8, 16x16, 32x32), (green down-triangle: p=4, 2x2, 4x4, 8x8, 16x16)

For each $p = 1$, 2, and 4, four different fine mesh resolutions are studied. These are chosen such that, independent of p, the lowest resolution case has 64 unknowns and the highest resolution case has 4096 unknowns. For the lowest resolution case, 4 multigrid levels are used and for the highest resolution case 8 multigrid levels are used (independent of p). For all cases, the coarsest approximation used in the multigrid V-cycle is a mesh of two triangles with linear polynomials.

Figure 1 shows the magnitude of correction made to the solution after each V-cycle iteration. The p=1 (blue square) curves correspond to standard geometric multigrid since p is equal to one on each mesh. From this case, we see that, as expected, the results are insensitive to the mesh resolution. The first column in Table 1 shows the convergence rates calculated using an exponential fit in the 5-20 V-cycle region. For this relaxation scheme a convergence rate of around 0.4 is obtained. For unknown reasons, there is a slight decrease in convergence rate for the finest mesh. This may be because the initial conditions have higher wavelengths than the mesh resolution and thus change from mesh to mesh.

Examining the p=2 (red up-triangle) curves on the figure, we see that there is almost no change between the p=1 and p=2 convergence results. Similarly, the $p = 4$ (green down-triangle) curves are also grid independent, but generally converge at a slightly slower rate than the $p = 2$ and $p = 1$ cases. This can be seen more precisely from the rates given in the table. These results demonstrate that p-multigrid is an effective method for solving equations resulting from higher-order discretizations of the Laplace equation.

Results can be obtained in nearly the same number of cycles at p=4 as $p = 1$, and the rates are also independent of the mesh resolution.

Table 1 Convergence rates for the Laplacian

Levels	$p = 1$	$p = 2$	$p = 4$
4	0.38	0.39	0.46
5	0.39	0.38	0.49
6	0.38	0.38	0.47
7	0.42	0.43	0.46

Convection

We now examine the performance of p-multigrid for convective problems. The governing equation is given by

$$\cos(\alpha)\frac{\partial u}{\partial x} + \sin(\alpha)\frac{\partial u}{\partial y} = 0 \tag{3}$$

The problem studied is on the same domain and has the same initial and boundary conditions as given for the Laplace equation. The only difference is that at the right hand boundary, a convective outflow boundary condition is enforced.

Two different discretizations are studied. The first is a SUPG[8] stabilized continuous discretization. Stabilization is necessary to ensure that there is coupling of the high wave numbers (odd/even modes) and thus a unique solution to the problem. The exact form of the equations are

$$\int_\Omega \left(w + \left[\cos(\alpha)\frac{\partial w}{\partial x} + \sin(\alpha)\frac{\partial w}{\partial y} \right] \tau \right) \times \left(\cos(\alpha)\frac{\partial u}{\partial x} + \sin(\alpha)\frac{\partial u}{\partial y} \right) d\Omega = 0 \tag{4}$$

where Ω is the domain and τ is a stabilization parameter which is taken as $h/(p^2)$ (assuming a wave speed of one). The dependence of the stabilization parameter on p is discussed in.[3,9] The above equation is enforced for all functions w in our approximation space.

The second discretization is a discontinuous Galerkin method. In this case, a Galerkin formulation is used on each element with a basis which is discontinuous between adjacent elements. Inter-element coupling is provided by an approximate Riemann flux evaluated at element boundaries. For the current test case, this reduces to evaluating the flux using data from the upwind side of the element boundary. As for SUPG, the upwinding of edge fluxes ensures that there is coupling of the high wave numbers and thus a unique solution to the problem. The main distinction between the two approaches is that, because continuity is not enforced for DG, the basis functions span only a single element while for a continuous formulation the basis functions can span two or more elements.

The multigrid iteration and the relaxation scheme are the same as for the Laplacian equation. In this case, the RK time step is proportional to h/p^2. For DG, the only additional change is that, because the mass matrix is totally decoupled between elements, it is inverted without using mass lumping.

For DG, we also use "element Jacobi" (eJ) as a relaxation scheme.[5] In element Jacobi, on each element we form the element stiffness matrix and invert it. Inter-element coupling terms are neglected such that the matrices are local to each element. This allows the process to be performed element by element. There is no easy way to do this for a continuous formulation.

1D Fourier Analysis

We begin by performing a 1D Fourier analysis of the multigrid iteration scheme. The solution (or error) is assumed to be of the form

$$u_k(\xi, \tau) = \hat{u}(\xi, \tau)e^{(ik\theta)},$$

where k denotes the element index, ξ denotes the local element coordinate, and $0 < \theta < \pi$. We are basically assuming that there is a periodic variation of the basis coefficients across elements. Given this periodic form, we then analyze the eigenvalues of a two-level multigrid iteration: This consists of relaxing the solution on the fine grid then solving the equations for the coarse grid correction exactly and adding the correction to the fine grid solution. The maximum eigenvalue of this process determines the convergence rate. The limit of $\theta = 0$ corresponds to a very long wavelength mode relative to the mesh. If in this limit the eigenvalues approach a value less than one, we have grid independent results.

Table 2 shows the maximum eigenvalues for the various discretizations. In all cases, we have eigenvalues

Table 2 Maximum eigenvalue for 2-level multigrid iteration

Levels	SUPG-RK	DG-RK	DG-eJ
$p = 2$ to $p = 1$	0.43	0.25	0.5
$p = 4$ to $p = 2$	0.94	0.83	0.5

that are below one in the limit of θ going to zero so we have grid independent results. For both SUPG and DG when using the RK scheme, the results are dependent on the polynomial degree. This is somewhat expected since the RK smoother is basically a time advancement; With a higher-order scheme, there is less damping on the high-wavenumber modes. For element Jacobi, the results are totally independent of p and the mesh. In this case, the fine grid iteration does a good job eliminating the intra-element error modes for both $p = 4$ and 2 while the coarse meshes eliminate the long wavelength errors. Actually with element Jacobi, one fine iteration eliminates all error modes interior to the element, and at the next level one can switch to $p = 1$ independent of the initial p and damp all the remaining modes.

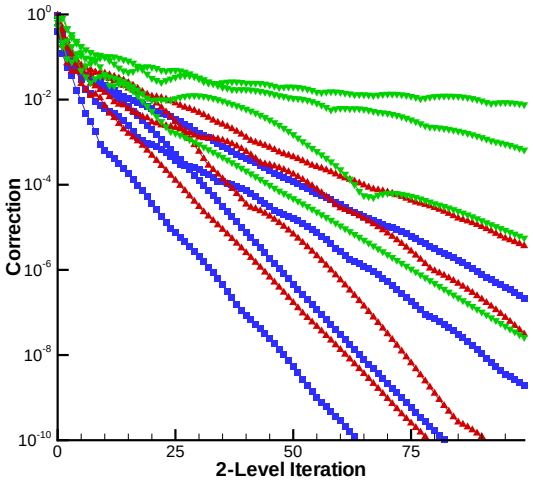

Fig. 2 Convergence results for SUPG-RK convection. (blue square: p=1, 8x8, 16x16, 32x32, 64x64), (red up-triangle: p=2, 4x4, 8x8, 16x16, 32x32), (green down-triangle: p=4, 2x2, 4x4, 8x8, 16x16)

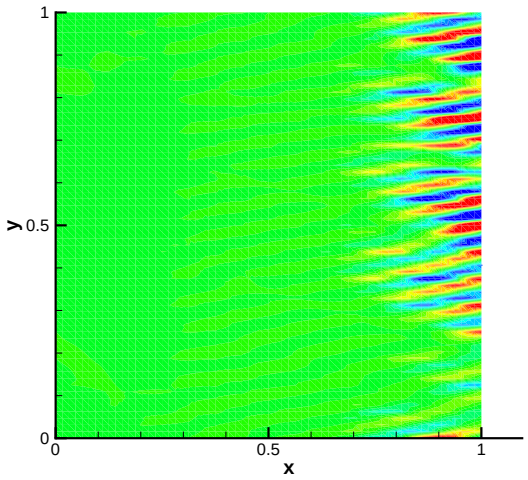

Fig. 3 Slowest converging mode for SUPG-RK with $p = 4$ and a 16x16 mesh

2D Numerical Results

We now examine the two-dimensional system. The resolutions and polynomial degrees studied are the same as for the Laplace system. To facilitate comparison to analysis, a two-level multigrid iteration is used. This makes it easier to evaluate the performance of the finest levels of the multigrid system.

Figure 2 shows the convergence results for the SUPG discretization. As before, the p=1 (blue) curves correspond to standard geometric multigrid. From these curves, we see that, unlike the Laplacian results, there is a large sensitivity to the mesh resolution even for $p = 1$ (standard multigrid). Table 3 shows the convergence rates calculated using an exponential fit in the 25-100 cycle region. For this relaxation scheme the convergence rate varies form 0.7 to 0.9 with increasing mesh resolution.

Table 3 Convergence rates for continuous SUPG-RK convection

Levels	$p = 1$	$p = 2$	$p = 4$
4	0.74	0.77	0.86
5	0.76	0.77	0.89
6	0.85	0.86	0.95
7	0.88	0.90	0.98

With increasing p the convergence results get worse even at the same resolution. For the highest resolution, at $p = 4$, the convergence degrades to 0.98. To understand why the convergence is much worse than that predicted by the 1D Fourier analysis, we examine a contour plot of the solution for $p = 4$ and a 16x16 mesh after 75 V-cycles. This is shown in figure 3. From this figure we see that the slowest converging mode is high-wavenumber across the streamwise direction and long wavelength in the streamwise direction (10 degrees). Since this is a high-wavenumber mode, it is not represented well on coarser meshes. In addition, there is little damping of this mode on the finest mesh because there is no coupling normal to the streamwise direction for equation 3. The damping must come from the propagation of information from the left boundary along the streamline. Since this takes longer on fine meshes, the convergence rates degrade with increasing mesh resolution.

Figure 4 shows the convergence results for the DG discretization. In this case, we have only studied $p = 4$ and $p = 2$. For $p = 1$, the coarsening should be to $p = 0$ since this is possible for a discontinuous basis, however Dubiner's triangular basis does not allow $p = 0$ so we have not studied this case. The trends are almost exactly the same as for the SUPG discretization with the convergence rate being sensitive to both the grid resolution and the polynomial degree. We have also found that the slowest converging mode for this problem is similar to that shown in figure 3. The convergence rates are also comparable.

The final cases we study are DG using element Jacobi. The results are calculated for the same cases shown in figure 4. In the element Jacobi scheme, we use a relaxation factor of 0.5 which gives good damping on the highest wavenumber modes. Figure 5 shows the results for element Jacobi. Although grid sensitive, the results are insenstive to p. In fact, for the same number of unknowns, the higher-order basis converges faster! Examining the slowest converging mode we again find the same mode as in the previous cases; a mode which is only damped on the fine mesh and must move across the domain to be eliminated.

The reason the higher-order basis converges faster for the same number of unknowns is that the error

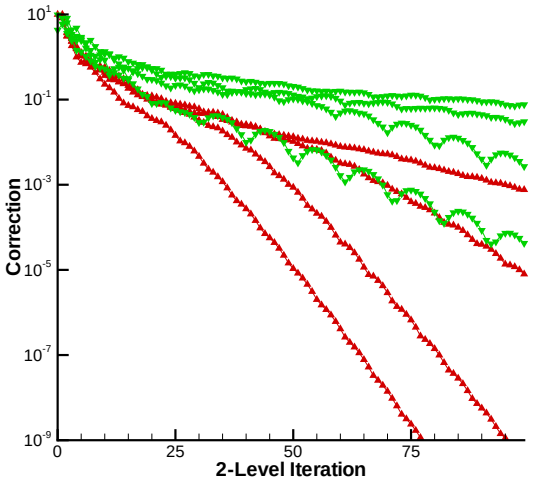

Fig. 4 Convergence results for DG-RK convection. (red up-triangle: p=2, 4x4, 8x8, 16x16, 32x32), (green down-triangle: p=4, 2x2, 4x4, 8x8, 16x16)

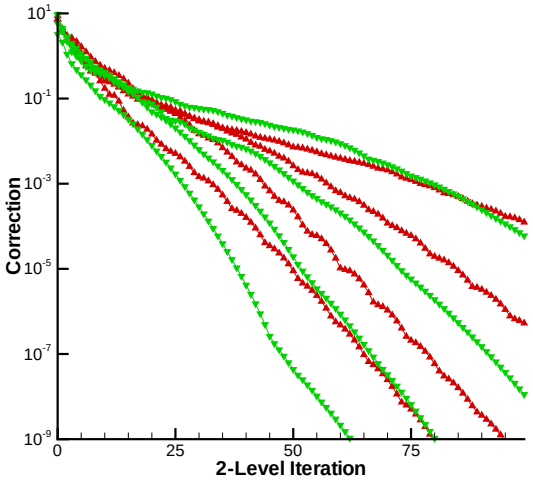

Fig. 5 Convergence results for DG-eJ. (red up-triangle: p=2, 4x4, 8x8, 16x16, 32x32), (green down-triangle: p=4, 2x2, 4x4, 8x8, 16x16)

modes move across the domain at a rate which is independent of p (approximately 1/2 an element per iteration). Since the mesh is actually coarser at higher p when we keep the same number of unknowns, the error modes are swept out faster. For the mass matrix approach, there is a stability limitation which depends on p, thus the modes are removed slower at higher p.

2D Fourier Analysis

To determine whether our explanation of the poor 2D performance in 2D is accurate, a 2D Fourier analysis is performed for square elements by assuming solution has the form

$$u_{j,k}(\xi,\eta,\tau) = \hat{u}(\xi,\eta,\tau)e^{i(j\theta_x + k\theta_y))}$$

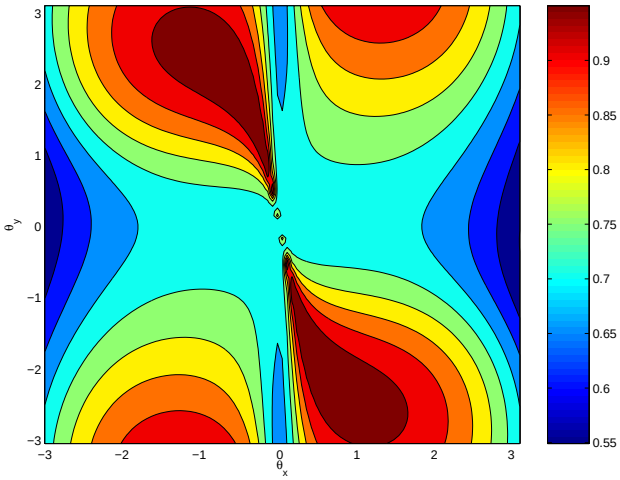

Fig. 6 2D Fourier Analysis of 2-level p-multigrid for convection. (DG-RK)

where j and k denote the element index, ξ and η denote local element coordinates, and $-\pi < \theta_x, \theta_y < \pi$. The analysis is for a DG discretization with the RK relaxation scheme and a two-level multigrid iteration going from $p = 2$ to $p = 1$. Figure 6 shows the magnitude of the maximum eigenvalue as a function of θ_x and θ_y.

Along $\theta_y = 0$, the analysis qualitatively agrees with the 1D results. The maximum eigenvalue along this line is 0.7. The most difficult modes to eliminate are along the line $(\theta_x, \theta_y) \cdot (-ay, ax) = 0$ passing through zero. (10 degrees counter-clockwise from the y-axis.) These modes are long wavelength in the streamwise direction and short wavelength in the cross-stream direction. The eigenvalues approach one so we no longer have grid independent results. This is consistent with our previous observations. For this problem however, the inlet and outlet boundary conditions have an important role in the convergence and should be included in the analysis.[10] For the doubly-periodic case, there is no damping on cross-stream modes because the equations have no coupling in this direction. Thus, this is the worst case scenario, but it does confirm the reasons for the poor performance in 2D.

Concluding Remarks

The p-multigrid solution algorithm provides an efficient means of solving the equations produced by high-order discretizations of diffusive systems as verified by the 2D tests of the Laplace equation. For convective systems however, the severely anisotropic nature of the equations cause difficulties for the isotropic p-multigrid iteration. Modes which are short wavelength normal to the streamlines but long wavelength along the streamlines do not get damped well. This was verified both by numerical experiments and 2D Fourier analysis. The only way to eliminate this problem is to either use a relaxation scheme which has strong damping along the

entire streamline or use anisotropic multgrid coarsening.

As a relaxation scheme, element Jacobi was shown to have convergence properties which were totally p independent in 1D for convection. However, this does not eliminate the problem discussed above for convection in 2D. A strong possibility for p and mesh independent solutions to the wave equation is a sweeping element-Jacobi iteration since element Jacobi can be applied element by element,. This will provide strong damping along the streamlines and thus eliminate the long-wavelength streamwise modes. Unfortunately element Jacobi is not easily extendable to continuous formulations.

References

[1]Rönquist, E. M. and Patera, A. T., "Spectral Element Multigrid. I. Formulation and Numerical Results," *J. Sci. Comput.*, Vol. 2, No. 4, 1987, pp. 389–406.

[2]Maday, Y. and Munoz, R., "Spectral Element Multigrid Part 2: Theoretical Justification," Tech. Rep. 88-73, ICASE, 1988.

[3]Helenbrook, B. T., "A Two-Fluid Spectral Element Method," *Comp. Meth. Appl. Mech. Eng.*, Vol. 191, 2001, pp. 273–294.

[4]Martinelli, L., *Calculations of Viscous Flow with a Multigrid Method*, Ph.D. thesis, Princeton University, Princeton, NJ, October 1987.

[5]Atkins, H. L. and Shu, C.-W., "Analysis of Preconditioning and Relaxation Operators for the Discontinuous Galerkin Method Applied to Diffusion," *AIAA*, Vol. 01-2554, 2001.

[6]Dubiner, M., "Spectral methods on triangles and other domains," *J. Sci. Comput.*, Vol. 6, No. 4, 1991, pp. 345–390.

[7]Mavriplis, D. J., "Multigrid Techniques for Unstructured Meshes," Tech. Rep. 95-27, ICASE, April 1995.

[8]Hughes, T. J. R. and Mallet, M., "A New finite element formulation for computational fluid dynamics: III. The generalized streamline operator for multidimensional advective diffusive systems," *Comp. Meth. App. Mech. Eng.*, Vol. 58, 1986, pp. 305–328.

[9]Gervasio, P. and Saleri, F., "Stabilized Spectral Element Approximation for the Navier-Stokes Equations," *Num. Meth. Part. Diff. Eq.*, Vol. 14, 1998, pp. 115–141.

[10]Diskin, B. and Thomas, J. L., "Half-Space Analysis of the Defect-Correction Method for Fromm Discretization of Convection," *J. Sci. Comput.*, Vol. 22, No. 2, 2001, pp. 633–655.