Application of "p"-multigrid to discontinuous Galerkin formulations of the Poisson equation

B. T. Helenbrook

Clarkson University

Mechanical and Aeronautical Engineering Department

Potsdam, NY 13699-5725

H. L. Atkins

Computational Modeling and Simulation Branch

NASA Langley Research Center

MS 128

Hampton, VA 23681-2199

Abstract

We investigate p-multigrid as a solution method for several different discontinuous Galerkin (DG) formulations of the Poisson equation. Different combinations of relaxation schemes and basis sets have been combined with the DG formulations to find the best performing combination. The damping factors of the schemes have been determined using Fourier analysis for both one and two-dimensional problems. One important finding is that when using DG formulations, the standard approach of forming the coarse p matrices separately for each level of multigrid is often unstable. To ensure stability the coarse p matrices must be constructed from the fine grid matrices using algebraic multigrid techniques. Of the relaxation schemes, we find that the combination of Jacobi relaxation with the spectral element basis is fairly effective. The results using this combination are p sensitive in both one and two dimensions, but reasonable convergence rates can still be achieved for moderate values of p and isotropic meshes. A competitive alternative is a block Gauss-Seidel relaxation. This actually out performs a more expensive line relaxation when the mesh is isotropic. When the mesh becomes highly anisotropic, the implicit line method and the Gauss-Seidel implicit line method are the only effective schemes. Adding the Gauss-Seidel terms to the implicit line method gives a significant improvement over the line relaxation method.

1 Introduction

p-multigrid is an iterative algorithm in which systems of equations arising from high-order finite element discretizations such as spectral/hp formulations, are solved by recursively iterating on solution approximations at different polynomial order, p. 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 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 [13] and analyzed by Maday and Munoz [12] for a Galerkin spectral element discretization of the Laplace equation. Helenbrook [8] combined p-multigrid with geometric multigrid and applied it to an unstructured streamwise-upwind-Petrov-Galerkin (SUPG) discretization of the incompressible Navier-Stokes equations. Recently there has also been work combining overlapping Schwarz relaxation methods with multigrid for spectral element discretizations [7].

All of the above work has been for continuous formulations. For discontinuous formulations, we have performed some analysis of p-multigrid for hyperbolic systems [9], but not for elliptic equations. In this paper, we analyze the efficiency of p-multigrid when applied to discontinuous Galerkin formulations of the Poisson equation. Because discontinuous formulations of the Poisson equation are relatively new, there has been little analysis of this combination. There are many factors that can affect the performance of p-multigrid for these formulations. In the following we examine various combinations of: DG formulation, polynomial

basis functions, and relaxation scheme. Fourier analysis of both one and two-dimensional problems is performed to assess the iterative efficiency. In 2D, we also examine the effect of mesh aspect ratio on the performance of the iteration. The first section of the paper summarizes the DG formulations for the poisson equation. Sections 3 and 4 give a description of the polynomial bases used and the relaxation schemes investigated. Section 5 describes the multigrid scheme and the analysis techniques used to determine its efficiency. The last two sections give results for one-dimensional and two-dimensional problems respectively.

2 Discontinuous Galerkin Formulations

The model problem we analyze is the Poisson equation in one and two dimensions. To simplify the analysis, we use a unit segment or unit square domain with periodic boundary conditions. When formulating a DG scheme, the Poisson equation is usually written in the following form

$$\sigma = -\nabla u \tag{1}$$

$$\nabla \cdot \sigma = f(x) \tag{2}$$

where u is the solution to the Poisson problem, x is the position vector, σ is a flux vector, and $f(x)$ is a given source function.

We then introduce a finite dimensional space space of functions to represent the solution. The domain is subdivided into either into uniform length (1D) or rectangular (2D) elements, and on each element, we use a polynomial basis to describe the solution, u, and the flux

vector σ. Following the notation in [1], we define the following spaces

$$V_h := \left\{ v \in L^2(\Omega) : v|_K \in P(K) \; \forall K \in \mathcal{T}_h \right\} \tag{3}$$

$$\Sigma_h := \left\{ \tau \in [L^2(\Omega)]^2 : \tau|_K \in \Sigma(K) \; \forall K \in \mathcal{T}_h \right\} \tag{4}$$

where $L^2(\Omega)$ is the space of square integrable functions on the domain Ω. $\mathcal{T}_h$ is the set of segments or quadrilateral elements K that triangulate the domain, and the subscript h refers to the element length associated with a particular mesh. In one dimension, $P(K) = \mathcal{P}_p(K)$ is the space of polynomial functions of degree at most p on segment K. In two dimension, $\mathcal{P}_p$ is formed from the tensor product of the one-dimensional space of polynomial functions. $\Sigma(K)$ is equal to $[P(K)]^2$.

All of the DG formulations we analyze are based on the weak form of equations 1 and 2. Multiplying these equations by a scalar test function, $v_h \in P(K)$ and a vector test function $\tau_h \in \Sigma(K)$ respectively and then integrating by parts gives the weak form which is used to find $u_h \in V_h$ and $\sigma_h \in \Sigma_h$

$$\int_K \sigma_h \cdot \tau \; dx = -\int_K u_h \nabla \cdot \tau dx + \int_{\partial K} \widehat{u}_K n_K \cdot \tau ds \;\; \forall \tau \in \Sigma(K) \tag{5}$$

$$-\int_K \sigma_h \cdot \nabla v dx = \int_K f v dx + \int_{\partial K} \widehat{\sigma}_K \cdot n_K \; v \; ds \;\; \forall v \in P(K) \tag{6}$$

n_K is the outward normal to the element boundary, ∂K, and $\widehat{u}_K$ and $\widehat{\sigma}_K$ are boundary flux functions. These functions are evaluated along element edges using information from both sides of the element and thus provide the inter-element coupling in a DG scheme.

The choice of the boundary fluxes distinguishes the various DG schemes [1]. We analyze p-multigrid with the schemes that are listed in table 1. The flux functions for these schemes

Scheme	$\widehat{u}_K$	$\widehat{\sigma}_K$
Bassi and Rebay [3]	$\{u_h\}$	$\{\sigma_h\}$
Brezzi et al. [5]	$\{u_h\}$	$\{\sigma_h\} - \alpha_r(\llbracket u_h \rrbracket)$
local DG (LDG) [6]	$\{u_h\} - \beta \cdot \llbracket u_h \rrbracket$	$\{\sigma_h\} + \beta \llbracket \sigma_h \rrbracket - \alpha_j \llbracket u_h \rrbracket$
interior penalty [11]	$\{u_h\}$	$\{\nabla_h u_h\} - \alpha_j \llbracket u_h \rrbracket$
Bassi et al. [4]	$\{u_h\}$	$\{\nabla_h u_h\} - \alpha_r(\llbracket u_h \rrbracket)$

Table 1: DG schemes analyzed and their numerical fluxes.

are also shown. The notation shown again follows that of [1]: Braces { } denote the average of a quantity along an edge. Double brackets $\llbracket \ \rrbracket$ denotes the jump in a quantity along an edge. For a scalar, q, the jump is a vector given by $q_1 n_1 + q_2 n_2$ where q_1 and q_2 are the values of the scalar evaluated from the elements adjacent to the edge and n_1 and n_2 are the opposing outward unit normals of these two elements. if q is a vector, the jump is given by $q_1 \cdot n_1 + q_2 \cdot n_2$

In the LDG scheme, the constant β can take values between -1/2 and 1/2. The constant α_j is given by ηh^{-1} where η is an $O(1)$ constant, and h is a characteristic mesh length normal to the edge. $\alpha_r(\llbracket u_h \rrbracket)$ is defined by a "lifting operator" [1]. Briefly, the following equation defines a vector function, $r_e \in \Sigma_h$, which is nonzero only on the elements on either side of the edge, e

$$\int_\Omega r_e \cdot \tau dx = - \int_e \llbracket u_h \rrbracket \cdot \{\tau\} ds \ \ \forall \tau \in \Sigma_h \tag{7}$$

$\alpha_r(\llbracket u_h \rrbracket)$ is then given by $\eta\{r_e\}$ where η is an $O(1)$ constant and the braces again denote the average of the discontinuous function r_e along the edge e. In one-dimension, $\alpha_r(\llbracket u_h \rrbracket)$ simplifies to the multiplication of $\llbracket u_h \rrbracket$ by a constant that depends on the order of the basis and the length of the elements. Except for the first scheme in the table, the above schemes are all consistent and stable. The first scheme is consistent but not stable. It is included so

5

Basis	Symbol	Property
Legendre	$P_n(\xi)$	$\int_{-1}^{1} P_m(\xi)P_n(\xi)d\xi = \delta_{m,n}$
$\int$ Legendre	$I_n(\xi) = \int_{0}^{\xi} P_{n-1}(\xi')d\xi'$	$\int_{-1}^{1} \dfrac{dI_m(\xi)}{d\xi}\dfrac{dI_n(\xi)}{d\xi}d\xi = \delta_{m,n}$
Monomial	$M_n(\xi) = \xi^n$	quadrature free integration [2]
Spectral Element	$G_n(\xi)$	nodal at Gauss-Lobatto points (Lagrangian)

Table 2: Basis sets used.

that we can understand how the stability of the scheme affects the multigrid iteration. Note that if β is zero, the LDG scheme and the Brezzi scheme are only different by the magnitude of α_r versus α_j. This is also true of the interior penalty scheme and the Bassi et al. scheme.

3   Basis Functions

Although the specific form of the basis functions used to represent the space $\mathcal{P}_p(K)$ does not affect the final solution for u_h, it can have a strong effect on the efficiency of the relaxation schemes. We investigate several different sets of basis functions. Unlike continuous formulations for which the form of the basis is constrained by continuity requirements, almost any basis can be easily used in a DG formulation. Table 2 shows the sets of one-dimensional basis functions we investigate and their special properties. The bases are defined on the domain $\xi \in [-1, 1]$. The two-dimensional bases are a tensor product of the one-dimensional bases.

The Legendre basis is orthogonal which gives a diagonal mass matrix. The mass matrix for a particular element K is defined as $M = \int_K \phi^T \phi dx$ where ϕ is the vector of basis

6

functions. The integrated Legendre basis is orthogonal with respect to the typical bilinear operator we would get in a continuous formulation of the Poisson equation. The monomial basis is included because it is simple and can be efficient if implemented properly [2]. The spectral element basis is typically used for continuous formulations. Operations at element boundaries are simpler with this basis because only a small subset of the bases are nonzero at the element boundaries. It has also been shown that this basis is well-suited for p-multigrid solutions of continuous formulations of the Poisson equation [13, 12, 9].

4   Relaxation Schemes

Before explaining the relaxation schemes, it is useful to introduce some matrix notation for the linear systems of equations generated by the DG formulations. Because DG formulations are dominated by operations on elements, we will label the solution coefficients by element such as u_j. This corresponds to the vector of coefficients used to describe the solution on element j. The solution on this element can then be written as $\phi^T u_j$ where ϕ is again the vector of basis functions. In one dimension, the vector of coefficients and basis functions is of length $p + 1$. In two dimensions, it is of length $(p+1)^2$. In 2D, we use a multidimensional indexing for the elements i.e. $u_{j,k}$ where j is the horizontal index and k is the vertical index.

The following matrix is an example of the general form of one-dimensional DG formulations

for a mesh with 4 elements and periodic boundary conditions.

$$\begin{bmatrix} M & 0 & 0 & 0 & D_M & D_R & 0 & D_L \\ 0 & M & 0 & 0 & D_L & D_M & D_R & 0 \\ 0 & 0 & M & 0 & 0 & D_L & D_M & D_R \\ 0 & 0 & 0 & M & D_R & 0 & D_L & D_M \\ S_M & S_R & 0 & S_L & U_M & U_R & 0 & U_L \\ S_L & S_M & S_R & 0 & U_L & U_M & U_R & 0 \\ 0 & S_L & S_M & S_R & 0 & U_L & U_M & U_R \\ S_R & 0 & S_L & S_M & U_R & 0 & U_L & U_M \end{bmatrix} \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ F_1 \\ F_2 \\ F_3 \\ F_4 \end{bmatrix} \qquad (8)$$

Each entry in the matrix is a block of dimension $(p+1) \times (p+1)$. The first four equations are the discrete equivalent of equation 5. The last four correspond to equation 6. The F_j terms are the vectors that result from the variational integration of the source term f on each element. Because all of the DG formulations choose the flux $\widehat{u}_K$ to be independent of σ, there is no inter-element coupling of σ, and thus the upper left quarter of the matrix is block diagonal. This allows σ to be found using local operations and eliminated from the problem.

Replacing σ_j by $M^{-1}(D_L u_{j-1} + D_M u_j + D_R u_{j+1})$ we arrive at a circulant penta-diagonal matrix with block entries

$$\begin{aligned} A_{LL} &= 0 & -S_L M^{-1} D_L \\ A_L &= U_L & -S_L M^{-1} D_M - S_M M^{-1} D_L \\ A_M &= U_M & -S_L M^{-1} D_R - S_M M^{-1} D_M - S_R M^{-1} D_L \\ A_R &= U_R & -S_M M^{-1} D_R - S_R M^{-1} D_M \\ A_{RR} &= 0 & -S_R D_R \end{aligned} \qquad (9)$$

All of the iterative schemes are applied to this form of the equations. The two-dimensional system is similar with a 9 element stencil. For the interior penalty method and the Bassi et al. scheme, the $\widehat{\sigma}$ fluxes do not involve σ. Therefore, S_L, S_M, and S_R are zero and the system is tri-diagonal. For the LDG method, if β is chosen uniformly as $+1/2$, the $\widehat{u}$ flux for the edge between element j and $j+1$ becomes one-sided involving only u_{j+1}. This implies that D_L is equal to 0. Similarly for $\widehat{\sigma}$, S_R becomes zero. We then arrive at a tri-diagonal

matrix again. For non-periodic problems however, β must be chosen to be compatible with the boundary information and it is in general not possible to produce a uniformly compact stencil with the LDG approach.

Given the above form for the discrete governing equations, we can now describe the iterative schemes. All of the schemes can be written in the form

$$R\Delta u + (Au - F) = 0 \tag{10}$$

R is the relaxation matrix, u is the vector of unknown coefficients for all elements, and A is the stiffness matrix which is composed of the block entries shown in equation 9. F is the entire vector of source terms consisting of the F_j vectors from each element. The first scheme we investigate is essentially a physical time advancement scheme. R is block diagonal with the element mass matrix scaled by a constant ω used as the diagonal blocks. This is exactly what we would arrive at if analyzed the heat equation instead of the poisson equation. The inverse of ω corresponds to the physical time-step and is taken as the maximum eigenvalue of $R^{-1}A$ calculated with $\omega = 1$. This relaxation scheme allows us to verify that the DG implementations are correct because it should give results consistent with the physical behavior of the heat equation. The results should also be independent of the form of the basis functions because the iterative matrix can be obtained from a bi-linear functional.

The next scheme is a Jacobi scheme with R composed of the diagonal elements of A. This is again scaled by a constant ω that is again taken as the magnitude of the maximum eigenvalue of $R^{-1}A$ calculated with $\omega = 1$. This is the least computationally intensive scheme. It is also the only scheme that gives results that are *dependent* on the form of the polynomial

9

basis. When using Legendre polynomials, this approach is very similar to the mass matrix approach because the mass matrix is diagonal. The only difference is that when using Jacobi, each mode effectively has its own time step. This is also true of the spectral-element basis, because the mass matrix of the spectral-element matrix is also diagonal when integrated using the $p+1$ point Gauss-Lobatto integration rule.

The last scheme is a block Jacobi scheme. In this case we take R to be the block diagonal matrices of A for each element , i.e. A_M. The results for this scheme should also be independent of the basis because the block diagonal term can be obtained from a variational form. The block Jacobi scheme is much more expensive than either of the first two schemes because it involves the inversion of a block matrix. Unlike the mass matrix scheme, a simple inversion that can be reused on various element geometries does not exist.

In practice, any of the above schemes can be transformed into a block Gauss-Seidel scheme by calculating the residual, $Au - f$, and updating the solution element by element. If the elements are ordered left to right then bottom to top, this corresponds to adding to R the block matrices of A that couple the element being updated to the elements to the left and down (essentially the lower triangle of A). To keep R circulant, these block matrices are added even when the position of the corresponding elements has wrapped around the domain due to periodicity. This maintains the periodicity of the problem, which makes the analysis easier.

5 p-multigrid

To implement the p-multigrid algorithm, restriction and prolongation operators are needed in addition to the relaxation scheme. Restriction consists of moving solution residuals from a space of high polynomial order to a lower order. We will typically choose the order of the polynomial spaces such that the coarse space has a degree, $p_c = p/2$. In some cases, however we will skip directly from order p to order 1. Prolongation is the reverse operation in which the solution correction from the low-order space is transferred to the higher-order space. For a basis ϕ_c which is contained in the space spanned by a higher-order basis, ϕ, the prolongation operator on an element is given by

$$I_{p,p_c} = M_K^{-1} \int_K \phi \phi_c^T dx \tag{11}$$

This is a matrix of dimension $p \times p_c$ which takes a correction represented using the basis ϕ_c and gives an equivalent representation using the basis ϕ. For hierarchical bases such as $P_n(\xi)$, $I_n(\xi)$, and $M_n(\xi)$ (bases in which the lower order basis functions are a subset of the higher-order functions) , the prolongation operator takes a very simply form. If the functions are organized from low to high order then the operator is simply an identity matrix of dimension $p_c + 1$ followed by $p - p_c$ rows of zeros. In all cases, the restriction operator is the transpose of prolongation.

The multigrid V-cycle algorithm can be written as a recursive subroutine as follows

cycle(p) {

    if (p = smallest p) {

$$u_{[p]} = A_{[p]}^{-1}(F_{[p]})$$

return

}

Relaxation:

$$u_{[p]} = u_{[p]} + R^{-1}(F_{[p]} - A_{[p]}u_{[p]}) \quad \nu = 0, ...n_d$$

Restriction:

$$p_c = p/2$$

$$F_{[p_c]} = I_{p,p_c}^T(F_{[p]} - A_{[p]}u_{[p]})$$

$$u_{[p_c]} = 0$$

cycle(p_c)

Prolongation:

$$u_{[p]} = u_{[p]} + I_{p,p_c}u_{[p_c]}$$

Relaxation:

$$u_{[p]} = u_{[p]} + R^{-1}(F_{[p]} - A_{[p]}u_{[p]}) \quad \nu = 0, ...n_u$$

return

}

The subscript $[p]$ indicates which polynomial space is being used. For the highest order space, $u_{[p]}$ is the solution to the discrete Poisson equation. For lower values of p, $u_{[p]}$ is a solution correction that will be prolongated to a higher-order space. For each level, n_d relaxation steps are performed on the first entry to the subroutine, and n_u relaxations are performed after the prolongation step. For the results presented here, we use $n_d = 1$ and $n_u = 0$. At the coarsest level, the matrix equations are directly inverted. Unless otherwise

noted, the results presented here are for a cycle with only two levels. The function shown uses $p_c = p/2$. When using block Jacobi relaxation, we also investigate the case of a direct jump to $p_c = 1$.

The stiffness matrices, $A_{[p]}$ at the coarse levels are determined using two different techniques. In the first case, the standard approach is used in which these matrices are evaluated in the same was as outlined above for the fine level. In the results, this will be unimaginatively referred to as the "standard" approach. In the second case, the coarse level stiffness matrices are evaluated using an "algebraic" approach. In this case, we use the restriction and prolongation operators

$$A_{[p_c]} = I_{p,p_c}^T A_{[p]} I_{p,p_c} \tag{12}$$

For some of the schemes, these two approaches yield the same results. Specifically, the interior penalty scheme can be formulated totally in terms of u and does not require static inversion of σ. Furthermore, the constant α_j does not change with p. Thus, compression by equation 12 simply reduces the order of the stiffness matrix yielding exactly the same results as if it were derived directly. The constant α_r in the Bassi et al. scheme changes with p resulting in a small difference in the two approaches. The first three schemes in Table 1 all require static inversion of σ. In this case, the two different approaches of deriving $A_{[p_c]}$ give very different results.

6   Analysis Techniques

For the analysis, we will assume that the source function f is zero everywhere because the source term has no affect on the convergence rate. To determine the convergence rates, we examine the eigenvalues of the multigrid iteration. For a two-level iteration with no source term, one multigrid cycle can be simplified to the following form

$$u^{[\nu+1]} = \left(I - R^{-1}A_{[p]}\right)^{n_u} \left(I - I_{p,p_c}A_{[p_c]}^{-1}I_{p,p_c}^{T}A_{[p]}\right)\left(I - R^{-1}A_{[p]}\right)^{n_d} u^{[\nu]} \tag{13}$$

where ν is the iteration counter and I is the identity matrix.

Because the matrices are circulant, the discrete Fourier transform can be used to determine the eigenvalues. In one dimension, we assume the solution on each element has the form

$$u_j = \tilde{u}e^{ij\theta} \tag{14}$$

where $\tilde{u}$ is a vector of dimension $(p+1)$ and θ can take values of $-\pi$ to π by increments of $2\pi/N$ where N is the number of elements. For all of the results presented N is chosen large enough such the results are essentially continuous functions of θ. Substitution of the form given by equation 14 reduces the dimension of the eigenvalue problem to $p+1$. This is then solved numerically at any θ. The maximum eigenvalue over the range $\pi < \theta < \pi$ is the damping factor of the iterative scheme.

The results for two dimensions are obtained in a similar way. In this case, we Fourier transform in both directions using $e^{i(j\theta_x+k\theta_y)}$. This reduces the problem to a $(p+1)^2$ eigenvalue problem that can be solved for each combination of θ_x and θ_y in the domain $[-\pi, \pi]^2$.

7 1D Results

In this section, we present one-dimensional results for the mass matrix iterative scheme, the Jacobi relaxation scheme, the block Jacobi iterative scheme, and lastly the improvement to these schemes when implemented as a Gauss-Seidel relaxation. We begin with the mass matrix scheme because this iteration is physically analogous to the unsteady heat equation. By comparing the two, we can validate the analysis techniques and also examine the accuracy of the various DG formulations in Fourier space. The remainder of the section focuses on the efficiency of p-multigrid.

7.1 Mass Matrix Relaxation

The eigenvalues of the matrix $R^{-1}(Au)$ calculated using the mass matrix preconditioner with $\omega = 1$ should correspond to the eigenvalues determined from the Fourier transform of the continuous problem

$$\frac{\partial u}{\partial t} - \frac{\partial^2 u}{\partial x^2} = 0 \tag{15}$$

Figure 1 shows the difference between the eigenvalues of the discrete schemes and the continuous scheme as a function of θ for $p = 4$. θ can be interpreted as the wavenumber of the eigenmodes nondimensionalized by Δx. The analytic eigenvalues in terms of θ are given by $-\theta^2/\Delta x^2$. In the figure, the error in the eigenvalues is non-dimensionalized by Δx^2. We also have unrolled the $p + 1$ eigenvalues in θ; At any value of θ we have $p + 1$ eigenvalues. The larger eigenvalues correspond physically to eigenfunctions with a wavenumber that is a shifted by an integer multiple of 2π from θ. We order the eigenvalues by magnitude and
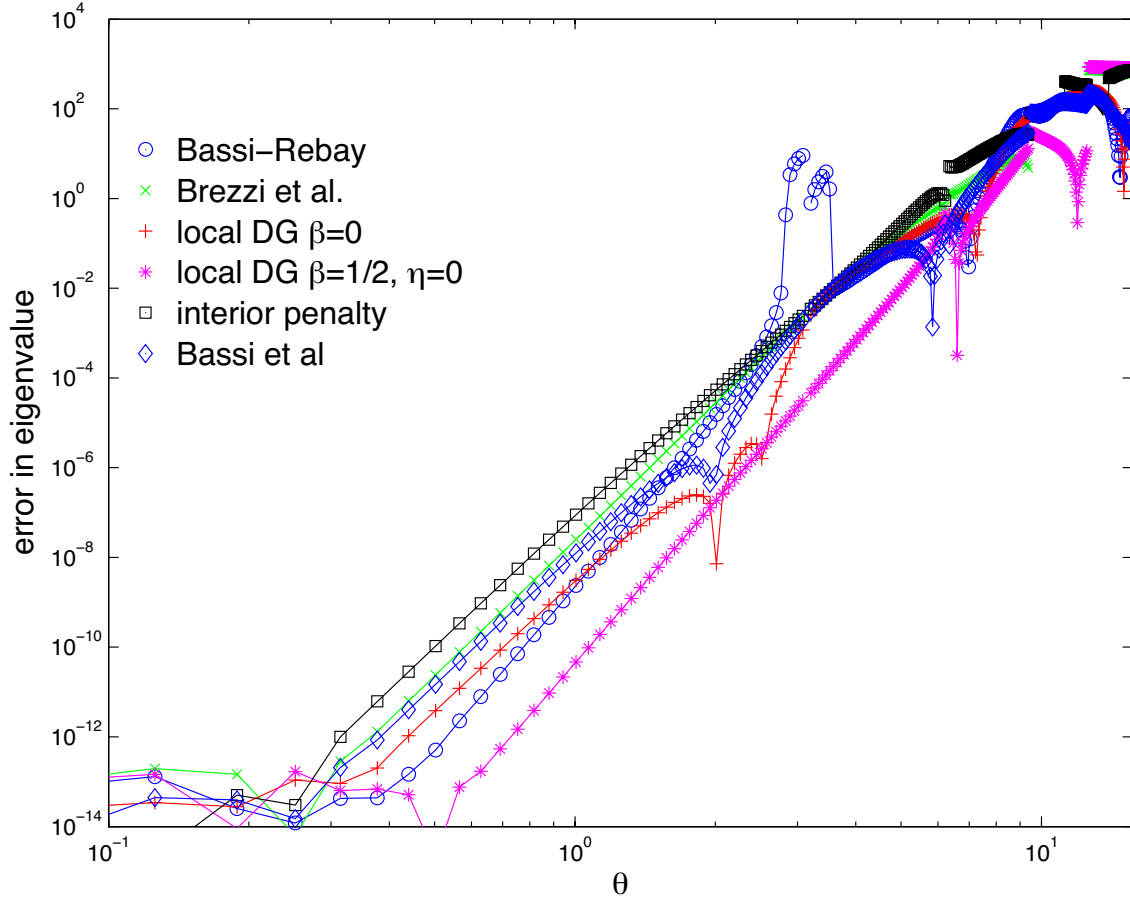
Figure 1: Accuracy of the eigenvalues for several DG formulations; $p = 4$.

then assign the eigenvalues to the positions θ, $|\theta - 2\pi|$, $2\pi + \theta$, $|\theta - 4\pi|$, $4\pi + \theta$... This allows us to plot them using the domain $\theta \in [0, (p+1)\pi]$. For these results, the constant η for the penalty terms of the DG schemes is chosen as 1.

The eigenvalues for a continuous formulation are expected to converge at a rate of θ^{2p+1} [10]. All of the discontinuous schemes converge at this rate, θ^9, except two, the Bassi and Rebay scheme and the one-sided LDG scheme ($\eta = 0, \beta = 1/2$). These two schemes converge with a rate of θ^{12}. These are also the only schemes that are consistent when p is equal to 0. Unfortunately, both of these schemes have other drawbacks. The Bassi and Rebay scheme

16

is not stable, and this manifests itself as the large increase in error around $\theta = \pi$ in the figure. At $\theta = \pi$ one of the eigenvalues is zero which indicates that there is an element scale odd-even decoupling problem. The uniformly one-sided LDG scheme can only be used in periodic problems or in problems where the flux is specified on one side of the domain and the temperature on the other. For any other configuration, β must change sign to be compatible with the boundary information. if β changes sign and no α_j term is included, this scheme also is not stable.

The one other exceptional result is the interior penalty scheme. With η equal to unity, the interior penalty method converges well for wavelengths greater than the element scale ($\theta < \pi$), however at higher wavenumbers the scheme loses stability and the eigenvalues actually change sign. This is impossible to see from the figure because it shows the absolute value of the error, and all the errors are large at high wavenumbers. By adjusting the penalty constant, η, we can make the interior penalty scheme and the Bassi et al. scheme identical. This is also true for the Brezzi et al. scheme and the LDG scheme with $\beta = 0$. At this value of p, the Bassi et al. scheme is equivalent to the interior penalty scheme with $\eta = 12.5$. To avoid this loss of stability, for most of the following cases we choose $\eta = 20$ when using the interior penalty scheme. We also parametrically investigate the effect of η on our results.

Before beginning our examination of p-multigrid, we establish some baseline properties of the mass matrix relaxation scheme. Figure 2 shows the damping of the mass matrix scheme as a function of θ for $p = 4$ using the Brezzi et al. scheme, whose properties are typical of most of the DG schemes considered. The plotting versus θ is done in the same way as for figure 1. The plot shows that the mass matrix relaxation scheme has good damping at
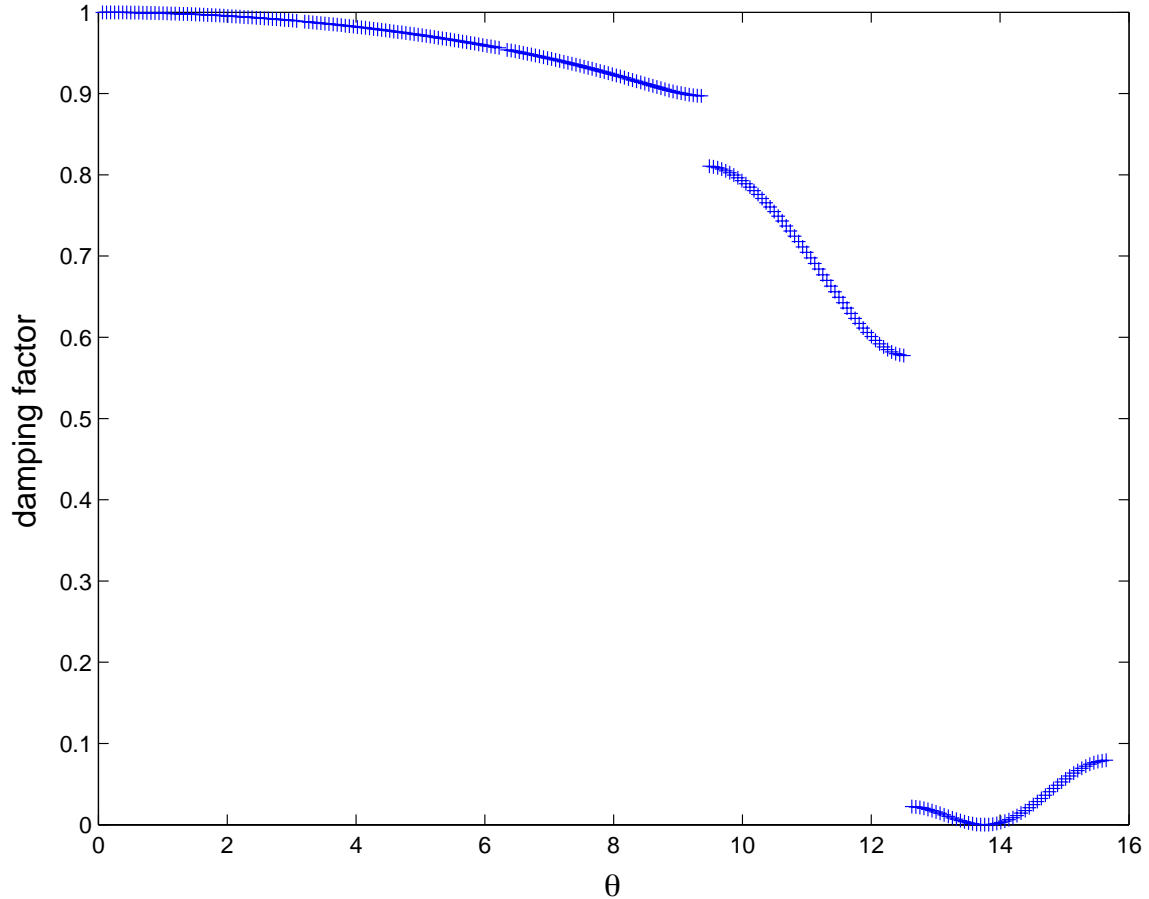
Figure 2: Damping factor for the mass matrix preconditioner applied to the Brezzi et al. scheme at $p = 4$.

high-wavenumbers but the damping rapidly deteriorates at smaller wavenumbers. All the other schemes are similar except for the Bassi and Rebay scheme. Because the Bassi and Rebay scheme has a zero eigenvalue at $\theta = \pi$, the damping factor is exactly 1 at this point.

Figure 3 shows the magnitude of the eigenvalue spectrum for the multigrid iteration when using the mass matrix iteration applied to the Brezzi et al. scheme with $p = 4$. Two sets of curves are shown. The curves marked with a "+" are obtained using the standard approach for obtaining the coarse grid stiffness matrix. The curves marked with a "□" are obtained using the algebraic approach. We again unroll the results in θ, but because of the

non-monotonic behavior of the eigenvalues, it is difficult to be sure that we have correctly shifted the eigenvalues to the correct θ domain. Examining the figure, we see that multigrid iteration does an excellent job of eliminating the error modes that are not removed by relaxation alone. The algebraic approach totally removes three of the five error components for all values θ. This is because the coarse p operator is perfectly consistent with the fine p operator and solving the coarse p operator completely eliminates the $p/2+1$ low-order error modes of the high-order system. For the standard approach, there is a difference between the coarse and fine space operators and for this reason only 2 of the three modes of the coarse space are totally eliminated. The overall damping factor of the standard approach is slightly better: 0.84 versus 0.87 for the algebraic case. We show later that this is exceptional; In most cases the algebraic approach works better. In either case, the maximum damping factor is independent of the number of elements in the grid so the convergence rates are grid independent .

Table 3 gives the damping factors for all of the schemes at polynomial degree of $p = 1, 2,$ 4, and 8. In each case, $p_c = p/2$. Most of the schemes are not consistent when $p = 0$, so we do not expect good results for the $p = 1$ case. However, it would be useful if we could coarsen to $p = 0$ because this system could then be solved using standard mesh-based multigrid techniques. In addition to examining the dependence on p, we have also varied the stabilization constant η from $1/4$ to 4 times its baseline value. For the Bassi and Rebay scheme and the LDG 1-sided scheme, the dashed entries imply that the scheme does not depend on η. Unless otherwise noted, the baseline value of η for all the schemes is one except for the interior penalty for which we use 20 for the reasons discussed above. In the
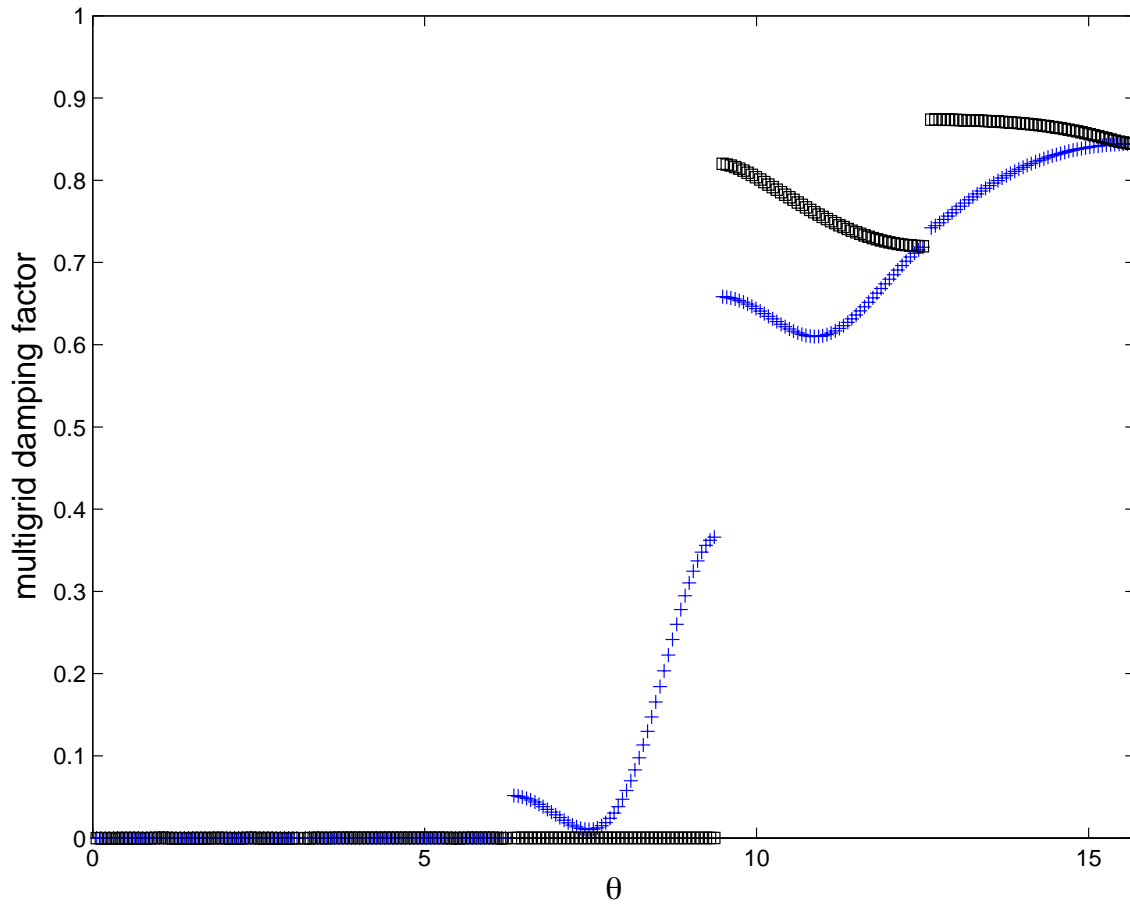
Figure 3: Damping factor for the mass matrix preconditioner applied to the Brezzi et al. scheme at $p = 4$ with multigrid. The $+$ is using the standard coarse grid operator. The $\square$ is the algebraically derived coarse grid operator.

A_c:		standard			algebraic		
η/η_0:		1/4	1	4	1/4	1	4
Bassi and Rebay	$p=1$	-	u	-	-	0.998	-
Brezzi et al.		u	u	u	0.77	0.72	0.89
LDG, $\beta=0$		u	0.69	0.80	0.85	0.69	0.80
LDG, $\beta=1/2$, $\eta=0$		-	u	-	-	0.75	-
interior penalty		0.8	0.95	0.988	0.8	0.95	0.988
Bassi et al.		u	u	u	u	0.54	0.88
Bassi and Rebay	$p=2$	-	0.998	-	-	0.998	-
Brezzi et al.		0.78	0.64	0.86	0.78	0.72	0.88
LDG, $\beta=0$		0.93	0.79	0.71	0.93	0.79	0.71
LDG, $\beta=1/2$, $\eta=0$		-	0.68	-	-	0.74	-
interior penalty		0.6	0.87	0.97	0.6	0.87	0.97
Bassi et al.		u	0.75	0.83	u	0.67	0.94
Bassi and Rebay	$p=4$	-	u	-	-	0.9992	-
Brezzi et al.		u	0.84	0.95	0.85	0.87	0.96
LDG, $\beta=0$		u	u	0.81	0.98	0.92	0.83
LDG, $\beta=1/2$, $\eta=0$		-	u	-	-	0.93	-
interior penalty		u	0.85	0.97	u	0.85	0.97
Bassi et al.		u	u	0.94	u	0.83	0.86
Bassi and Rebay	$p=8$	-	u	-	-	0.9997	-
Brezzi et al.		u	0.98	0.993	0.96	0.98	0.993
LDG, $\beta=0$		u	u	u	0.994	0.98	0.95
LDG, $\beta=1/2$, $\eta=0$		-	u	-	-	0.98	-
interior penalty		u	u	0.98	u	u	0.98
Bassi et al.		u	u	0.99	u	0.98	0.991

Table 3: Multigrid damping factors in 1D with the mass matrix relaxation scheme.

table, a "u" indicates that the iteration was unstable.

A scan of the entries reveals that there are many more "u"'s when using the standard multigrid approach to determine A_c. The only unstable entries when using the algebraic approach occur for the interior penalty scheme and the Bassi et al. scheme. Both of these schemes lose stability when η is too small. Thus, the problem is with the scheme itself not the multigrid iteration. For the Bassi et al. scheme, as long as η is kept larger than one, this problem is avoided. For the interior penalty scheme however, as we change p from 4

to 8 with η/η_0 fixed at one, the scheme changes from stable to unstable. Thus, the penalty parameter for the interior penalty scheme must be adjusted with p. All of the remaing u's in the table are caused by inconsistency between A_c evaluated with the standard approach and A. Although, the damping factor is not universally better when using the algebraic A_c, we conclude that the algebraic approach for evaluating A_c is correct because it guarantees that the coarse $'p'$ problem is consistent with the high $'p'$ problem. To reduce the possible combinations to be investigated, in the remaining results only the algebraic A_c is used. We will also drop the interior penalty scheme because it is of the same form as the Bassi et al. scheme but requires manual adjustment of η with p. The Bassi and Rebay scheme will be dropped as well because the damping factor for this scheme is much worse than the other schemes. In fact, we can get results arbitrarily close to unity for this scheme by increasing the resolution in θ. This is again because of the lack of stability of the scheme.

Looking at the dependence of the schemes on η for the algebraic A_c cases, we see that the Brezzi et al. scheme performs best at $\eta = 1$ or $\eta = 1/4$. The LDG scheme with $\beta = 0$ performs best at $\eta = 4$ for all p except $p = 1$. At any p, the Brezzi et al. scheme and the LDG $\beta = 0$ scheme only differ by the constant α_j and α_r. Because we are no longer going to use the standard method for evaluating A_c, the variation of this constant with p becomes irrelevant for multigrid. Examining the combined results of Brezzi et al and LDG noting that for $p =$1, 2, 4, and 8, the ratio of α_r to α_j is 0.5, 4.5, 12.5, and 40 respectively, we find that the minimum in damping factor occurs when $\eta\alpha_j\Delta x = 2.0$ for all values of p. Thus, we will also eliminate the Brezzi et al. scheme and only use the LDG $\beta = 0$ scheme with $\eta = 4$ as the baseline in the remaining results. The Bassi et al. scheme works best with $\eta = 1$,

which will be used as the baseline for this scheme in the remaining results.

Examining the remainder of the entries in the table, we see that there is a strong sensitivity to p in the damping factor. The $p = 2$ and $p = 1$ entries are comparable when using the algebraic coarsening technique, so we may be able to coarsen all the way to $p = 0$. At higher p the damping factor degrades. To show why this occurs, figure 4 shows the damping factor for the same conditions as shown in figure 2 except at $p = 8$. For the multigrid scheme to work well, the wavenumbers that will be truncated in moving to the coarser space must be damped well by the relaxation scheme. Figure 4 shows that only the highest order mode is suitably damped when using the mass matrix scheme. This implies that we must either move to the space $p_c = p - 1$ or use a more effective relaxation scheme.

7.2   Jacobi Relaxation

We next investigate the Jacobi relaxation scheme. The Jacobi relaxation results depend on the form of the polynomial basis used so in this case we investigate different bases. Table 4 shows the damping factors for various basis sets and DG formulations. These results show that the spectral-element basis is superior to all of the other bases when using Jacobi relaxation. The integrated Legendre basis is very poor. Although this basis is orthogonal with respect to the bi-linear product associated with the diffusive operator, this property is not valuable because of the boundary coupling. The monomial form also performs very badly. This is not unexpected because monomials lack any orthogonality property. The Legendre basis performs moderately well, but not quite as well as the spectral element basis. This may be because all of the modes of the Legendre basis are nonzero at the element
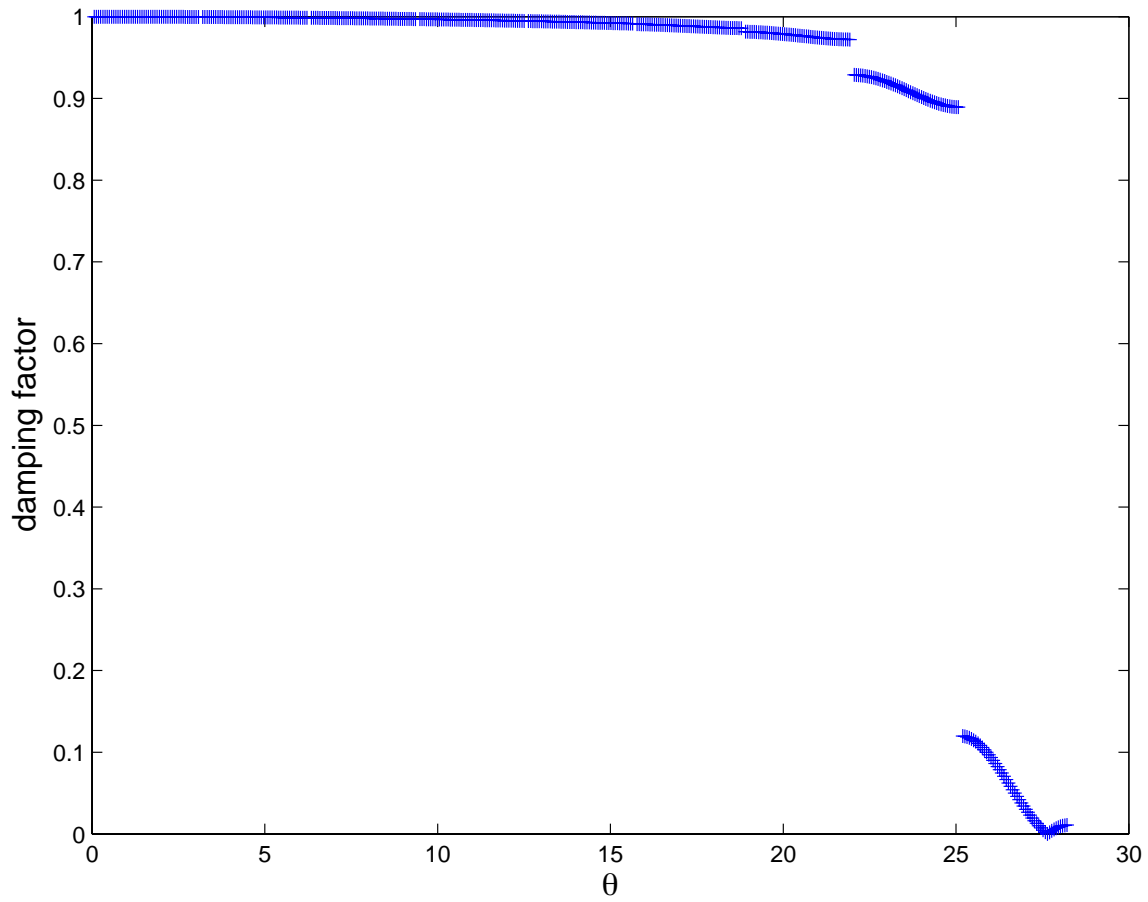
Figure 4: Damping factor for the mass matrix preconditioner applied to the Brezzi et al. scheme at $p = 8$.

Basis:		$P_n(\xi)$	$I_n(\xi)$	$M_n(\xi)$	$G_n(\xi)$
LDG, $\beta = 0$	$p = 1$	0.86	0.86	0.86	0.88
LDG, $\beta = 1/2$, $\eta = 0$		0.82	0.82	0.82	0.80
Bassi et al.		0.73	0.73	0.73	0.73
LDG, $\beta = 0$	$p = 2$	0.83	0.92	0.92	0.67
LDG, $\beta = 1/2$, $\eta = 0$		0.86	0.94	0.94	0.65
Bassi et al.		0.80	0.90	0.90	0.69
LDG, $\beta = 0$	$p = 4$	0.90	0.97	0.995	0.69
LDG, $\beta = 1/2$, $\eta = 0$		0.96	0.97	0.997	0.77
Bassi et al.		0.91	0.98	0.996	0.78
LDG, $\beta = 0$	$p = 8$	0.97	0.99	1.0	0.91
LDG, $\beta = 1/2$, $\eta = 0$		0.99	0.98	1.0	0.86
Bassi et al.		0.97	0.99	1.0	0.90

Table 4: Multigrid damping factors in 1D with the Jacobi relaxation scheme for different basis sets.

boundaries which leads to greater inter-element coupling.

Compared to the continuous case, the discontinuous p-multigrid iteration does not perform as well. For a continuous spectral element formulation, the combination of Jacobi preconditioning and p-multigrid gives p independent results in one dimension [13]. Table 4 shows that there is a sensitivity to p. However, the damping factors are still reasonably good for moderate values of p.

The table also shows there is only weak sensitivity to the DG formulation chosen even though the schemes are very different. This suggests that the damping factors may not vary significantly with the parameters of the DG schemes. To verify this, we recalculate the results with the spectral element basis with η four time larger. All of the results with η four times smaller are worse so these are not shown. The results for the LDG $\beta = 0$ scheme and the Bassi et al. scheme are shown in Table 5. Large improvements in the damping factor can be obtained by increasing η, however increasing η degrades the accuracy of the scheme.

η/η_0:	1	4
LDG, $\beta = 0$ $p = 1$	0.88	0.95
Bassi et al.	0.73	0.88
LDG, $\beta = 0$ $p = 2$	0.67	0.55
Bassi et al.	0.69	0.51
LDG, $\beta = 0$ $p = 4$	0.69	0.59
Bassi et al.	0.78	0.60
LDG, $\beta = 0$ $p = 8$	0.91	0.76
Bassi et al.	0.90	0.68

Table 5: Multigrid damping factors in 1D with the Jacobi relaxation scheme for different values of η

For the LDG scheme, with $\eta = 16$, the error at any θ is nearly a factor of 10 larger than the values shown in figure 1, which are calculated with $\eta = 1$.

We have also checked that the Jacobi relaxation scheme works well in a multilevel multigrid cycle. Because we are using the algebraic approach to determine the coarse space matrices, at any p these matrices will be different depending on the fine space they were started from. This means the behavior of the $p = 4$ to $p = 2$ transition will behave differently depending on the polynomial degree we started from. To ensure that this does not have any adverse effects, we investigate a V-cycle with 4 levels starting at $p = 8$ and ending at $p = 1$. The damping factor for LDG, $\beta = 0$ using the Jacobi scheme and $\eta = 4$ is 0.91. This is exactly the same as the two-level case. The result for the Bassi et al. case with $\eta = 1$ is also identical. This is because the highest order modes are the most difficult to damp and thus determine the damping factor. If we stop the cycle at $p = 0$ instead however the results become worse. For both the LDG scheme and the Bassi et al. schemes, the damping factor is 0.98. A similar behavior is seen if we start from $p = 4$. For the remainder of the results, we will only investigate a two-level iteration because these results provide an accurate prediction of the

η/η_0:		1/4	1	4
LDG, $\beta = 0$	$p = 2$	0.62	0.27	0.09
LDG, $\beta = 1/2$, $\eta = 0$		-	0.00	-
Bassi et al.		u	0.53	0.10
LDG, $\beta = 0$	$p = 4$	0.81	0.52	0.20
LDG, $\beta = 1/2$, $\eta = 0$		-	0.00	-
Bassi et al.		u	0.61	0.08
LDG, $\beta = 0$	$p = 8$	0.94	0.8	0.49
LDG, $\beta = 1/2$, $\eta = 0$		-	0.00	-
Bassi et al.		u	0.76	0.10

Table 6: Multigrid damping factors in 1D with the block Jacobi relaxation scheme.

damping factor of a full V-cycle iteration, and we will not investigate the $p = 1$ to $p = 0$ transition.

7.3   Block Jacobi Relaxation

The last relaxation scheme we examine is block Jacobi relaxation. The block Jacobi precon-ditioner gives results that are independent of the basis however there is a strong sensitivity to the DG parameters so we again investigate various values of η/η_0. The results are shown in table 6. In the table, the results for one-sided LDG are identically zero for all p. In this case, block Jacobi corresponds to a static inversion of the higher-order modes and then a direct solve of the low-order equations. This results in a direct inversion of the equations. For the other two schemes, we get very good damping factors, especially when using the larger values of η.

Because block Jacobi is such a strong relaxation scheme, we also investigate the case in which we move directly to $p_c = 1$. Table 7 gives these results. Compared to the damping factors given in table 6, there is little change if we move directly to $p = 1$. This may recover some

η/η_0:		1	4
LDG, $\beta = 0$	$p = 4$	0.58	0.25
LDG, $\beta = 1/2$, $\eta = 0$		0.00	-
Bassi et al.		0.68	0.11
LDG, $\beta = 0$	$p = 8$	0.83	0.55
LDG, $\beta = 1/2$, $\eta = 0$		0.00	-
Bassi et al.		0.80	0.13

Table 7: Multigrid damping factors in 1D with the block Jacobi relaxation scheme and $p_c = 1$.

of the additional cost of block Jacobi compared to the simpler relaxation schemes. We have also tried moving directly to $p = 0$, but this again causes a large degradation in performance.

7.4   Gauss-Seidel

As mentioned previously, any of the previous relaxation schemes can used with a Gauss-Seidel approach of updating the solution. We find that there is little improvement to either the mass matrix preconditioner or the Jacobi preconditioner results when the Gauss-Seidel terms are added. There is obviously no improvment to the block Jacobi approach for the LDG one-sided scheme since this is a direct solve anyway. For the two other DG formulations, the block Jacobi preconditioner with Gauss-Seidel gives the improved results shown in Table 8. Since the additional computational cost of evaluating the Gauss-Seidel terms is much less than the cost of inverting the diagonal blocks, it is definitely beneficial to add the Gauss-Seidel terms to the iteration.

LDG, $\beta = 0$ $p = 2$	0.15
Bassi et al.	0.32
LDG, $\beta = 0$ $p = 4$	0.34
Bassi et al.	0.41
LDG, $\beta = 0$ $p = 8$	0.68
Bassi et al.	0.63

Table 8: Multigrid damping factors in 1D with the block Jacobi Gauss-Seidel relaxation scheme.

8   2D Results

We next investigate the performance of p-multigrid in multiple dimensions. We again We begin with Jacobi preconditioning using the spectral element basis, and then look at block Jacobi relaxation. In the last part of the section, we examine Gauss-Seidel iteration and line relaxation schemes. The results are obtained for both isotropic and high aspect ratio grids. The reason the line relaxation schemes are introduced is that they have been shown to be effective in finite volume solvers for high-aspect ratio grids.

8.1   Jacobi

Table 9 shows the damping factors for various values of η. Compared to the 1D results shown in table 5, these results are all worse. This behavior also occurs when using p-multigrid to solve continuous spectral element formulations [9]. The explanation for this behavior that is usually given is that a high-order finite-element simulation corresponds to a discretization that is on a high aspect ratio mesh; the spacing of Gauss Legendre points near $\xi = -1$ or 1 goes like $1/p^2$ as compared to $1/p$ if the spacing is uniform. This causes the same aspect ratio difficulties that occur for geometric multigrid iterations on high aspect ratio meshes; at

η/η_0:		1/4	1	4
LDG, $\beta = 0$	$p = 2$	0.94	0.90	0.95
LDG, $\beta = 1/2$, $\eta = 0$		-	0.93	-
Bassi et al.		u	0.86	0.94
LDG, $\beta = 0$	$p = 4$	0.96	0.90	0.90
LDG, $\beta = 1/2$, $\eta = 0$		-	0.93	-
Bassi et al.		u	0.92	0.94
LDG, $\beta = 0$	$p = 8$	0.987	0.96	0.92
LDG, $\beta = 1/2$, $\eta = 0$		-	0.96	-
Bassi et al.		u	0.96	0.93

Table 9: Multigrid damping factors in 2D with the Jacobi relaxation scheme.

the high-wavenumbers there is a directional dependence in the eigenvalues which results in poor damping of some of the high wavenumber modes. Although, the performance decreases with p, at moderate values of p the damping rates are still moderately good ≈ 0.9.

8.2   Block Jacobi

Table 6 shows the damping factors for the block Jacobi relaxation scheme. Again we see a significant reduction in performance in two-dimensions. In one-dimension, the modes that are undamped by block Jacobi are well represented in the coarse space. In two dimensions, there are modes that are high wavenumber in one direction and low wavenumber in the other that are not damped well by block Jacobi but not represented well in the coarse space either. These modes cause the performance to degrade. Because of the greater computational cost of block Jacobi, it may be more efficient to use the Jacobi iteration in two dimensions.

η/η_0:		1/4	1	4
LDG, $\beta = 0$	$p = 2$	0.70	0.59	0.79
LDG, $\beta = 1/2$, $\eta = 0$		-	0.63	-
Bassi et al.		u	0.61	0.77
LDG, $\beta = 0$	$p = 4$	0.86	0.68	0.75
LDG, $\beta = 1/2$, $\eta = 0$		-	0.73	-
Bassi et al.		u	0.73	0.85
LDG, $\beta = 0$	$p = 8$	0.95	0.88	0.84
LDG, $\beta = 1/2$, $\eta = 0$		-	0.89	-
Bassi et al.		u	0.86	0.94

Table 10: Multigrid damping factors in 2D with the block Jacobi relaxation scheme.

LDG, $\beta = 0$	$p = 2$	0.44
LDG, $\beta = 1/2$, $\eta = 0$		0.48
Bassi et al.		0.42
LDG, $\beta = 0$	$p = 4$	0.51
LDG, $\beta = 1/2$, $\eta = 0$		0.58
Bassi et al.		0.52
LDG, $\beta = 0$	$p = 8$	0.74
LDG, $\beta = 1/2$, $\eta = 0$		0.73
Bassi et al.		0.69

Table 11: Multigrid damping factors in 2D with the Block Gauss-Seidel relaxation scheme.

8.3   Gauss-Seidel

As is the case in one-dimension, the Gauss-Seidel terms do little to improve the damping factors of the Jacobi iteration, but the block Jacobi iteration, shown in table 11, again improves significantly compared to results without the Gauss-Seidel terms (table **??**). Thus in multiple dimensions, it is again beneficial to include the Gauss-Seidel terms in the block relaxation scheme.

$\Delta x/\Delta y$:	1/10	1	10
Jacobi $p = 2$	0.997	0.90	0.997
Block Jacobi	0.978	0.58	0.988
Block Gauss-Seidel	0.96	0.44	0.96
Line Solve	0.28	0.53	0.989
Line Gauss Seidel	0.16	0.36	0.96
Jacobi $p = 4$	0.997	0.90	0.997
Block Jacobi	0.988	0.69	0.988
Block Gauss-Seidel	0.95	0.51	0.95
Line Solve	0.51	0.61	0.987
Line Gauss Seidel	0.32	0.43	0.95
Jacobi $p = 8$	0.9987	0.96	0.9987
Block Jacobi	0.991	0.85	0.991
Block Gauss-Seidel	0.96	0.73	0.96
Line Solve	0.79	0.81	0.992
Line Gauss Seidel	0.66	0.66	0.96

Table 12: Multigrid damping factors in 2D: effect of aspect ratio for various iterative schemes and polynomial degree. The DG formulation is the LDG $\beta = 0$ scheme with $\eta = 4$.

8.4   Aspect Ratio Effects

The last issue we examine is the effect of mesh aspect ratio on the damping rates. The effect of aspect ratio on all the schemes is nearly the same so we only show results for the LDG, $\beta = 0$ scheme. Table 12 shows the damping factors for various relaxation schemes as a function of aspect ratio. We have introduced two new relaxation schemes that are commonly used for high aspect ratio problems. The first is a block line solver which is a block tri- or penta-diagonal preconditioner that must be inverted along lines. It is penta-diagonal for the LDG $\beta = 0$ and tri-diagonal for the LDG one-sided scheme and the Bassi et al. scheme. The lines are oriented in the x-direction. The second scheme is similar except that the line updates are performed sequentially from bottom to top and each line uses the most recent information in its update i.e. it is a line solve in x and Gauss-Seidel in y.

Examining the results, we see that the line solvers are the only effective relaxation scheme when the mesh has a high aspect ratio. To get good performance, the lines must be aligned with the compressed direction of the mesh. The line solvers actually improve in performance as the mesh aspect ratio increases. Considering that the cost of a line-solve is much larger than evaluating Gauss-Seidel terms it is highly beneficial to use a the line Gauss-Seidel relaxation version. When the mesh is isotropic, the regular line solve actually performs worse than block Gauss-Seidel which is much less expensive. For the $p = 8$ and $p = 4$ cases, we have checked that these results are reproducible using a V-cycle to $p = 1$. As was found in the 1D case, the damping factors for the V-cycle are either identical or very close to those shown in Table 12.

9 Conclusions

p-multigrid can be a very effective way to solve discontinuous Galerkin formulations of the Poisson equation. A key finding is that the coarse space matrix operators must be derived by applying using the restriction and prolongation operators to fine space matrices. The standard approach of re-evaluating these matrices for each space often results in an unstable iteration.

Of the relaxation schemes evaluated, the Jacobi relaxation scheme with a spectral element basis gives reasonable results on isotropic meshes. An alternative is the block Gauss-Seidel scheme. This scheme gives damping factors on the order of 0.5 for $p = 4$ and 0.7 for $p = 8$ in 2 dimensions. On high-aspect ratio meshes, the most effective scheme is the line Gauss-

Seidel iteration. This scheme gives aspect ratio independent results and damping factors on the order of 0.4 for $p = 4$ and 0.6 for $p = 8$.

A problem that remains is that the multigrid efficiency degrades if the polynomial space is coarsened beyond $p = 1$. For moderate p discretizations, after coarsening to $p = 1$ the remaining system may still be fairly large and expensive to solve directly. Standard geometric multigrid techniques could be easily applied to the $p = 0$ discontinuous system, but there is no obvious generalization of these techniques to the $p = 1$ discontinuous system.

References

[1] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. Unified analysis of discontinuous galerkin methods for elliptic problems. *SIAM J. Numer. Anal.*, 39(5):1749–1779, 2002.

[2] H. L. Atkins and C.-W. Shu. Quadrature-free implementation of discontinuous Galerkin method for hyperbolic equations. *AIAA Journal*, 36(5):775–782, May 1998.

[3] F. Bassi and S. Rebay. A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier–Stokes equations. *J. Comp. Phys.*, 131:267–279, 1997.

[4] F. Bassi, S. Rebay, G. Mariotti, S Pedinotti, and M. Savini. A high-order accurate discontinuous finite element method for inviscid and viscous turbomachinery flows. In R. Decuypere and G. Dibelius, editors, *Proceedings of the 2nd European Conference*

on Turbomachinery, Fluid Dynamics, and Thermodynamics, pages 99–108, Antwerpen, Belgium, 1997. Technologisch Instituut.

[5] F. Brezzi, M. Manzini, D. Marini, P. Pietra, and A. Russo. Discontinuous finite elements for diffusion problems. *Ist. Lomb. Acc. Sc. Lett.*, 16, 1999.

[6] B. Cockburn and C.-W. Shu. The local discontinuous Galerkin method for time-dependent convection-diffusion systems. *SIAM J. Numer. Anal.*, 35(6):2440–2463, 1998.

[7] P. F. Fischer and J. W. Lottes. Hybrid/multigrid schwarz algorithms for the spectral element method. *J. Sci. Comput.*, to appear, 2004.

[8] B. T. Helenbrook. A two-fluid spectral element method. *Comp. Meth. Appl. Mech. Eng.*, 191:273–294, 2001.

[9] B. T. Helenbrook, H. L. Atkins, and D. J. Mavriplis. Analysis of "p"-multigrid for continuous and discontinuous finite element discretizations. In *The 16th AIAA Computational Fluid Dynamics Conference*, Orlando, Florida, June 2003. AIAA-2003-3989.

[10] T. J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis.* Prentice Hall, Englewood Cliffs, New Jersey, 1987.

[11] Jr. J. Douglas and T. Dupont. Interior penalty procedures for elliptic and parabolic galerkin methods. In *Computing Methods in Applied Sciences*, volume 58 of *Lecture Notes in Physics.* Springer-Verlag, Berlin, Germany, 1976.

[12] Y. Maday and R. Munoz. Spectral element multigrid part 2: Theoretical justification. Technical Report 88-73, ICASE, 1988.

[13] E. M. Rönquist and A. T. Patera. Spectral element multigrid. I. formulation and numerical results. *J. Sci. Comput.*, 2(4):389–406, 1987.