20th AIAA Computational Fluid Dynamics Conference
27 - 30 June 2011, Honolulu, Hawaii

AIAA 2011-3230

A Comparison of Higher-Order Methods on a Set of Canonical Aerodynamics Applications

Julie Andren *
Aerospace Computational Design Laboratory, Massachusetts Institute of Technology
Haiyang Gao †
CFD Center, Iowa State University
Masayuki Yano ‡ David L. Darmofal §
Aerospace Computational Design Laboratory, Massachusetts Institute of Technology
Carl Ollivier-Gooch ¶
Advanced Numerical Simulation Laboratory, University of British Columbia
and Z.J. Wang ‖
CFD Center, Iowa State University

Higher-order discretizations have the potential to reduce the computational cost required to achieve a desired error level. In this study, we consider higher-order discretizations of the conservation equations suitable for unstructured, triangular grids. In particular, the methods studied include continuous (SUPG/GLS) and classical discontinuous Galerkin (DG) finite element methods, the correction procedure via reconstruction (CPR) formulations of the DG and spectral volume methods, and cell and vertex-centered finite volume (FV) algorithms. This paper presents subsonic and supersonic, inviscid results for a canonical set of aerodynamic applications. Error convergence and computational performance of these discretizations are compared, and preliminary results indicate that the methods perform relatively similarly. When singularities are present in the flow solutions and uniformly refined meshes are used, all methods fail to achieve optimal convergence rates, and the performance benefits of the higher-order discretizations are reduced; adaptive meshing improves the efficiency of the higher-order method and recovers optimal convergence rates.

I. Introduction

The potential for improved computational efficiency is of great interest to the CFD community, and the use of numerical methods based on higher-order discretizations is one potential solution. Like traditional CFD techniques, there are numerous ways to formulate higher-order methods, each with advantages and disadvantages. This paper highlights the preliminary results of a study comparing higher-order discretizations for unstructured grids. The following discretizations are considered in the study: the Galerkin Least Squares (GLS) method, the classical discontinuous Galerkin (DG) finite element method, a DG formulation using the correction procedure via reconstruction (CPR-DG), a CPR formulation of the spectral volume (SV) method, cell-centered finite volume method, and vertex-centered finite volume method with a median dual. Solutions are approximated with polynomials of degree $p = 1, 2$, and 3 on identical families of meshes for a number

*Masters candidate, AIAA Student Member, 77 Massachusetts Ave. 37-442, Cambridge, MA, 02139, jandren@mit.edu
†Doctoral candidate, AIAA Student Member, 2272 Howe Hall Ames, IA 50011-2151
‡Doctoral candidate, AIAA Student Member, 77 Massachusetts Ave. 37-442, Cambridge, MA, 02139
§Professor, AIAA Associate Fellow, 77 Massachusetts Ave. 37-207, Cambridge, MA, 02139
¶Professor, Senior Member AIAA, 6250 Applied Science Lane, Vancouver, BC V6T 1Z4, Canada
‖Wilson Professor of Engineering, AIAA Associate Fellow, 2249 Howe Hall, Ames, IA 50011

Copyright © 2011 by J. Andren. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission.

of fundamental aerodynamic flows. For problems with singularities, both adapted and non-adapted meshes are used to quantify the effect of solution irregularities on the error.

Various finite element discretizations can be derived from the weak form of the conservation laws by choosing different solution and test function spaces. The GLS method[1] belongs to a family of stabilized continuous Galerkin methods, where solution and test basis functions consist of continuous polynomials with local support. As in other finite element methods, the solution order is increased by simply using higher-degree polynomials for the basis functions. To overcome the instability of the Galerkin formulation for convection-dominated problems, the weak form is augmented by a stabilization term. Details on the stabilization matrix for the GLS method used in this study can be found in the reference.[2]

Discontinuous Galerkin methods use polynomials with element-wise compact support to construct the solution and test spaces. This results in solutions that are continuous within each element but discontinuous across element interfaces. These discontinuities increase the number of degrees of freedom for DG schemes relative to its continuous counterpart for the same solution order on a given mesh. However, the discontinuities provide a natural means of stabilizing the convection operator by introducing upwinding through a Riemann solver, e.g. Roe's approximate Riemann solver.[3] Since coupling between the elements are produced only through inter-elemental fluxes, the method results is an element-wise nearest neighbor stencil.

Wang and Gao have extended the one-dimensional flux reconstruction method to a two-dimensional correction procedure via reconstruction (CPR).[4,5] CPR transforms the weighted residual formulation of the conservation equation into a differential form. This new structure eliminates the need to calculate integrals via numerical quadrature, potentially increasing computational efficiency. It produces a general system that must be uniquely defined by a set of lifting coefficients. The lifting coefficients are a constant set of numbers for linear triangular and tetrahedral elements, thus they can be analytically calculated and directly loaded into the solver with minimal memory and computational overhead. By choosing the certain lifting coefficients, the CPR formulations can recover a number of discretizations. The DG and SV discretizations based on the CPR formulation are used in this study.

Both the cell-centered and vertex-centered finite volume methods identically apply the conservation law to each control volume in the domain, but differ in how the control volumes are constructed. The cell-centered finite volume method uses each cell as a control volume, while the vertex-centered finite volume method uses a median dual. Higher-order solutions are generated by computing a least-squares fitting on neighboring control volumes and introducing a flux limiter at the interfaces.[6,7]

Section II of the paper describes the procedure used to generate the adapted meshes for this study, the importance of understanding the stability of a discretization and the ability of a given space to approximate the solution when comparing different higher-order methods, and the cost metrics of interest in the analysis. Section III presents results for the following flows:

A. Subsonic, inviscid flow over a Gaussian bump

B. Non-lifting, subsonic, inviscid flow over a NACA 0012 airfoil

C. Lifting, subsonic, inviscid flow over a NACA 0012 airfoil

D. Non-lifting, supersonic inviscid flow over a diamond airfoil

Although these flows are simpler than problems typically solved in the industry, they are useful for comparing different discretizations. Future studies will extend this work to additional viscous and supersonic problems.

II. Methodology

A. Adaptation

All adapted meshes in Section III are generated using an output-based error estimation and adaptation strategy. The strategy is based on estimating the error in an output quantity of interest, such as lift or drag, and refining the regions contributing to large errors. For brevity, we only provide an overview of these techniques; a thorough review of the framework for aerodynamic applications is provided by Fidkowski and Darmofal,[8] and the particular scheme used in this work is described by Yano et al.[9]

The adaptation strategy iterates toward the dof-optimal mesh—the mesh that minimizes the output error for a given degree of freedom. The error estimation is based on the dual-weighted residual (DWR) method of

Becker and Rannacher,[10,11] which provides elemental indicators that quantify the local contribution to the output error. Based on the error estimate, the size of the elements are controlled to equidistribute the local errors using a fixed-fraction marking strategy. The shape of elements are selected based on the behavior of the local Mach number.[12,13] Using the size and shape information, high-order anisotropic simplex meshes are generated using the metric-conforming linear mesh generator BAMG[14] and an elasticity-based mesh curving technique.[15] Meshes generated using the strategy provide: 1) the element size distribution suited for computing the specified output, 2) strong element size grading toward singularities, 3) efficient resolution of arbitrarily-oriented directional features such as shocks, wakes, and boundary layers, and 4) high-order geometry information necessary for high-order methods.

B. Discretization Approximability and Stability

Different discretizations have varying ability to handle strongly graded meshes, in particular those that arise from adaptive mesh refinements. The quality of the solution, and thus the error, is a function of the approximability—the ability of a given space to represent the solution—and the stability of the discretization. Ideally, a discretization should be stable in the sense that the discrete solution is close to the best solution attainable for a given space. Mesh adaptation aims to improve the approximability of the space such that the true solution can be well approximated using finite degrees of freedom. The error can be reduced by improving the approximability through mesh refinement *assuming* the method remains stable; however, if the stability of the method is compromised on strongly graded meshes, the method cannot benefit from the improved approximability. A previous comparison of the stability of higher-order methods coupled with mesh adaptation is included in the work of Venkatakrishnan.[16]

While all methods considered in this work are high-order accurate, they exhibit different stability properties. This, in turn, results in the methods having similar convergence rates, but different absolute error levels, which may be of greater interest to a practitioner. The results of this study are inevitably affected by the fact that all adaptive meshes are provided based on error estimates from the classical DG discretization. Further study is necessary to characterize the impact of this bias.

C. Measuring Computational Cost

Computational cost is measured in terms of degrees of freedom and the number of non-zeros in the Jacobian matrix. Since all problems in this study are two-dimensional, the element size, h, scales with $1/\sqrt{dof}$, where dof are the degrees of freedom defined by the solution order and number of elements. As a result, the slope of error results plotted against $1/\sqrt{dof}$ illustrates the convergence rate with uniform mesh refinement.

Since all the schemes considered in this work use implicit time integration, solving the linear system at each iteration is a significant fraction of the computational cost, especially for large systems. The number of non-zero (NNZ) entries in the Jacobian matrix is indicative of the cost of solving the linear system. Assuming the each node is surrounded by six elements, the NNZ per element per state squared for each discretization can be approximated by the values in Table 1.

Table 1. Number of non-zero entries in the Jacobian matrix per element per state squared

p	DG	GLS	CPR-DG	CPR-SV	FV cell	FV vertex
1	36	3.5	27	27	19	19
2	144	23.0	90	90	19	39
3	400	76.5	220	220	31	39

Since the solvers used for this study are not production level codes, wall-clock time is not included in this analysis.

III. Results

A. Subsonic, inviscid flow over a Gaussian bump ($M_\infty = 0.5$)

The study begins by analyzing subsonic, inviscid flow over a Gaussian bump. The absence of viscosity, subsonic conditions, and infinitely differentiable geometry provide a simple problem that allows for easy

assessment of how these discretizations solve smooth flows.

The Gaussian bump geometry is defined by

$$y = He^{-a^2x^2}, \tag{1}$$

where $H = 0.0625$ and $a = 5$. The problem is solved on a family of uniformly refined, structured meshes; the coarsest mesh is shown in Figure 1. Nodes are equally spaced in the y-direction and logarithmically graded towards the bump peak in the x-direction. The inflow (fixed stagnation pressure and stagnation temperature) and outflow (fixed static pressure) conditions are consistent with a freestream Mach number of $M_\infty = 0.5$. The output of interest is the L^2 entropy error normalized by the freestream entropy error and the square root of the domain area:

$$s_{err} \equiv \sqrt{\frac{\int_\Omega \left(\frac{p/\rho^\gamma - p_\infty/\rho_\infty^\gamma}{p_\infty/\rho_\infty^\gamma}\right)^2 dV}{\int_\Omega dV}}, \tag{2}$$

where ρ and ρ_∞ are the density and freestream density, p and p_∞ are the static pressure and freestream static pressure, γ is the ratio of specific heats, and Ω is the computational domain.

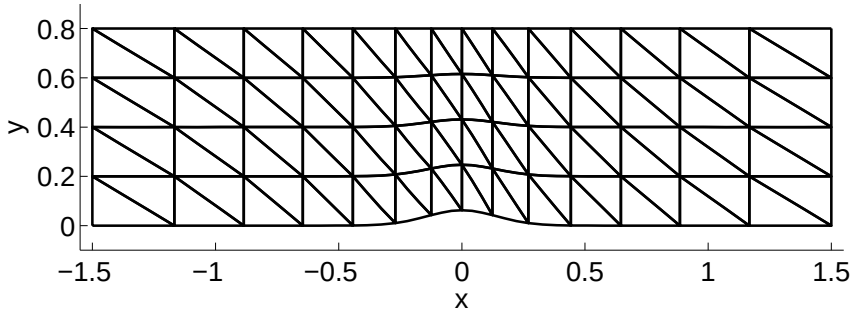

Figure 1. Gaussian bump coarse mesh (112 elements)

Figure 2a plots the variation of the L^2 entropy error with respect $1/\sqrt{dof}$, and Table 2 provides the convergence rates for solution orders $p = 1, 2$, and 3, for each discretization. The rates are computed using results from the two finest meshes; convergence rates fall between the values $p + 1/2$ and $p + 1$.

The horizontal dotted line at 10^{-4} in Figure 2a highlights two important differences between the solutions obtained by lower- and higher-order methods. Firstly, the variation among methods is greater for lower-order methods. The degrees of freedom necessary to obtain an entropy error level of 10^{-4} vary by an order of magnitude for $p = 1$, while the spread is much smaller for $p = 2$ (a factor of 1.5) and $p = 3$ (a factor of 1.625). This suggests that solutions obtained from different higher-order discretizations are more similar than those using lower-order approximations for smooth flows.

Secondly, the dotted line shows that for this smooth problem, fewer degrees of freedom are required for higher-order methods at low error levels. For example, regardless of the discretization, approximately 10^3 degrees of freedom are necessary to achieve an L^2 entropy error of 10^{-4} when using a solution order of $p = 3$, while at least three times as many degrees of freedom are required to achieve the same order of accuracy for $p = 1$. On the other hand, for error levels greater than 10^{-3}, no benefit is obtained for any method by increasing the solution order.

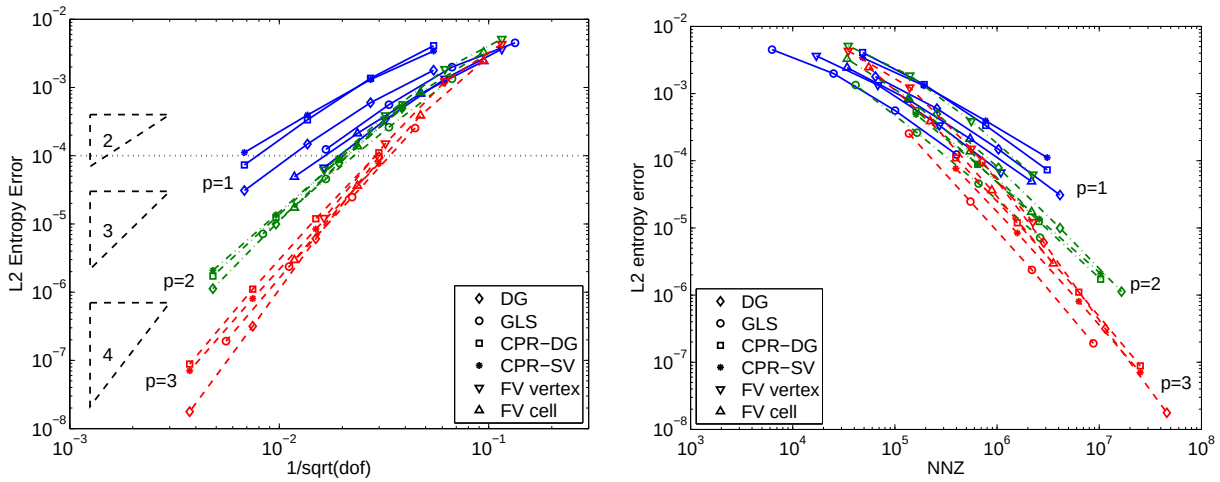

(a) vs. $1/\sqrt{dof}$ (b) vs. NNZ

Figure 2. Gaussian bump L^2 entropy error

Table 2. Gaussian bump L^2 entropy error convergence rates

p	DG	GLS	CPR-DG	CPR-SV	FV cell	FV vertex
1	2.264	2.176	2.202	1.821	2.110	2.466
2	3.146	2.673	2.853	2.702	2.993	2.980
3	4.169	3.634	3.643	3.515	3.620	4.267

The L^2 entropy error is plotted against the NNZ in Figure 2b. A slight advantage of using finite volume methods is seen for $p = 1$. For higher order, the GLS method has reduced NNZ compared to the other methods. The collocation-based formulation of the CPR-DG method results in a cheaper Jacobian than the more traditional DG method. However, because the classical DG method produces slightly lower error than the CPR-DG method, the L^2 entropy error relative to the cost of the Jacobian are comparable, regardless of the solution order.

B. Non-lifting, subsonic, inviscid flow over a NACA 0012 airfoil ($M_\infty = 0.5$)

Next we consider non-lifting, subsonic, inviscid flow over a NACA 0012 airfoil with a freestream Mach number of $M_\infty = 0.5$. The airfoil contour is closed by extending the upper and lower surfaces at the trailing edge until they meet; the airfoil is then scaled to have unit chord. These conditions allow us to understand how solution irregularities affect error convergence; the dominant singularity in this problem is geometry-induced and located at the trailing edge.

The far field boundary is located 500 chords away from the leading edge in order to reduce the effect of the boundary conditions on the solution. Restricting the problem to a finite domain and imposing freestream condition along this boundary is not an exact representation of inviscid flow in an infinite domain, introducing modelling error. Because the effect of the far field conditions is proportional to the inverse of the squared distance from the airfoil to the far field boundary, increasing the domain size greatly reduces this influence. Results show that the effect of having the far field boundary at 500 chords is insignificant, and the computed drag accurately represents the drag error for the error levels considered. The numerical flux at the far field boundary is computed using a Riemann solver, which combines the freestream conditions with the interior state.

A family of five meshes was constructed using a meshing technique that produces a set of meshes with similar refinement distributions as uniform refinement. Element distribution is prescribed by a function that defines the element size everywhere in the domain based on the total number of elements. These meshes will be referred to as non-adapted meshes. Since the degrees of freedom of finite volume methods depend only on the number of elements, and not the approximation order, both finite volume schemes were run on

two finer meshes in order to compare finite volume results with the other higher-order discretizations that require more degrees of freedom on a given mesh.

To further investigate the ability of higher-order methods to capture this type of flow, adapted meshes were also generated. A different mesh is generated for each combination of solution order and number of elements. Figure 3 shows the non-adapted, $p = 1$ adapted, and $p = 2$ adapted meshes near the airfoil (left) and at the trailing edge (right); all three meshes contain approximately 4800 elements. By design, the non-adapted mesh refines the leading and trailing edges more than other regions. However, mesh refinement at the trailing edge is more aggressive in the adapted meshes in order to better control the error near the singularity. Results show that appropriately refining the mesh reduces the cost necessary to achieve a given error level.

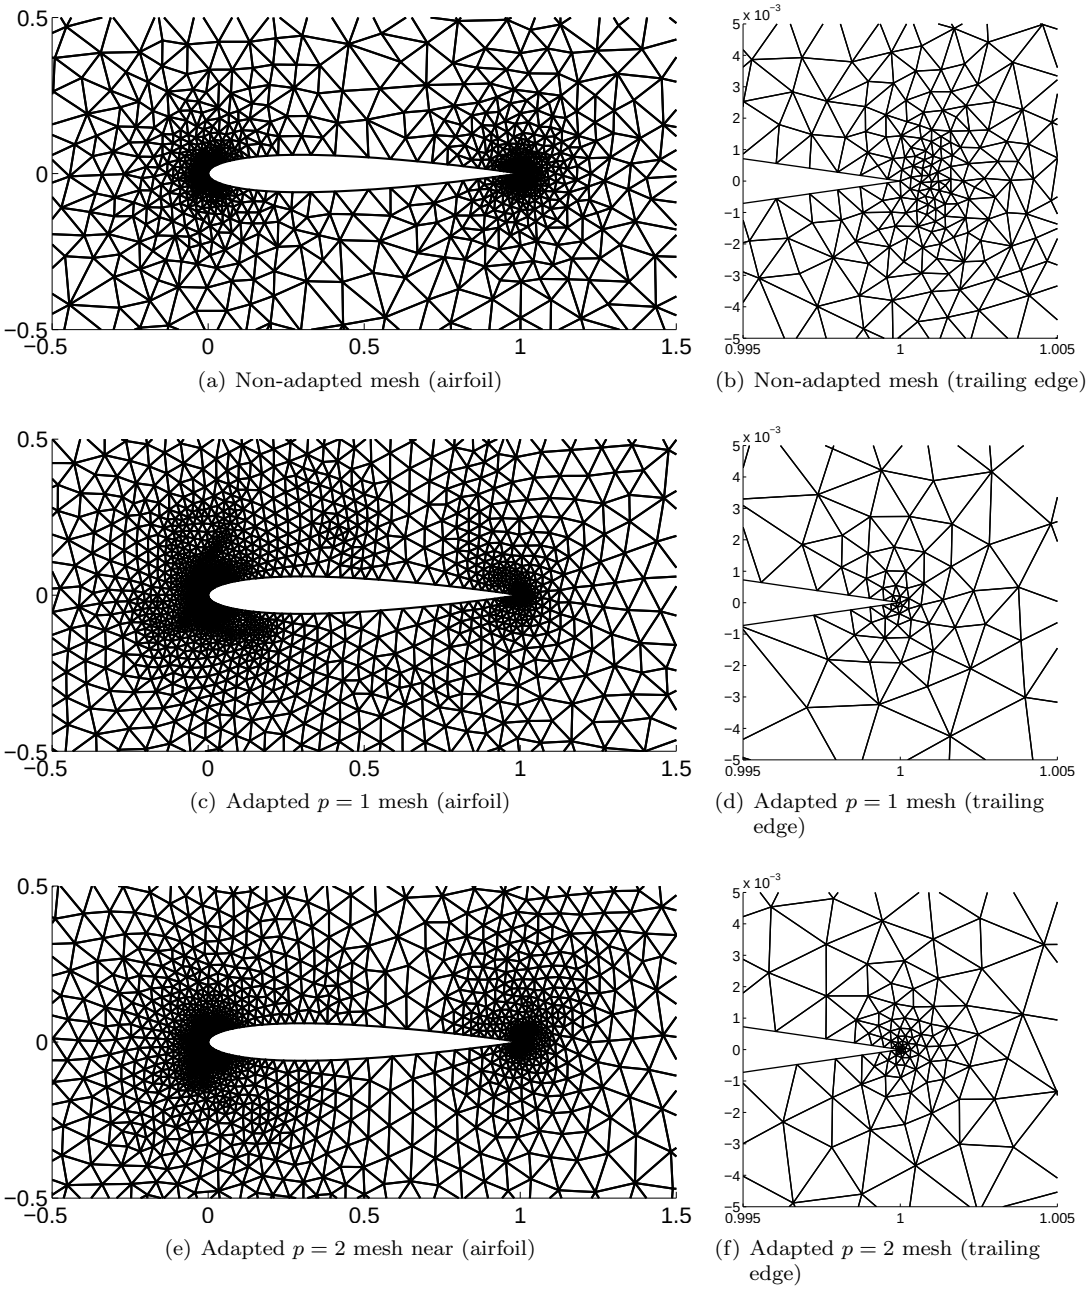

(a) Non-adapted mesh (airfoil)

(b) Non-adapted mesh (trailing edge)

(c) Adapted $p = 1$ mesh (airfoil)

(d) Adapted $p = 1$ mesh (trailing edge)

(e) Adapted $p = 2$ mesh near (airfoil)

(f) Adapted $p = 2$ mesh (trailing edge)

Figure 3. NACA 0012 meshes for non-lifting conditions ($\approx$ 4800 elements)

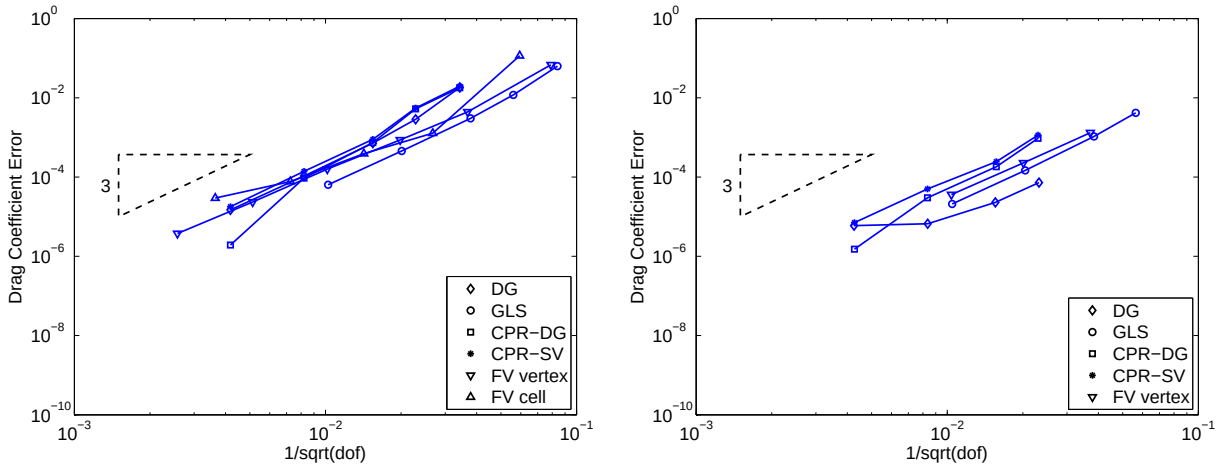

(a) p=1, non-adapted

(b) p=1, adapted

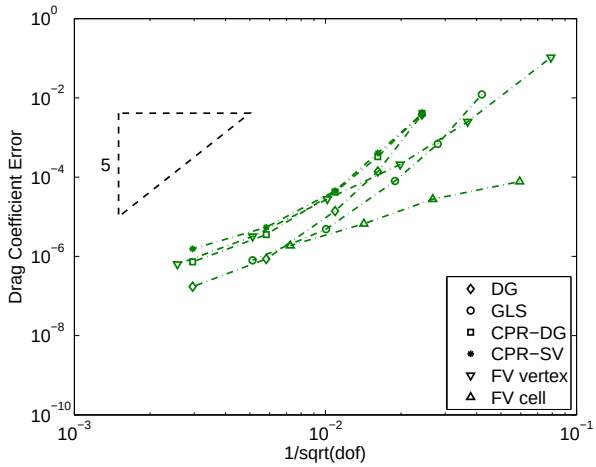

(c) p=2, non-adapted

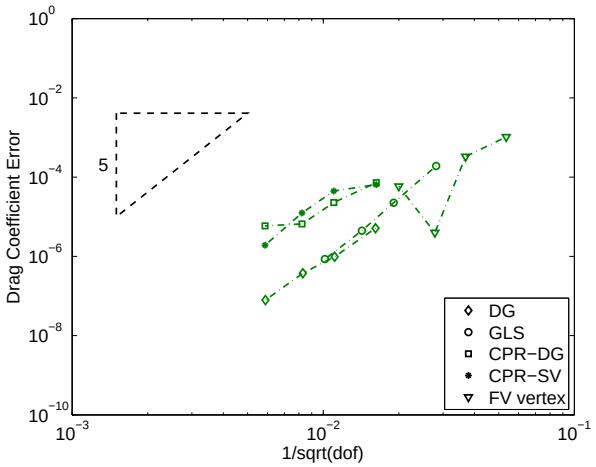

(d) p=2, adapted

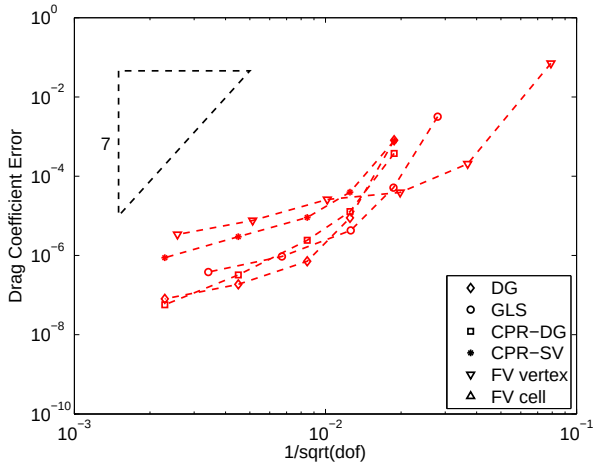

(e) p=3, non-adapted

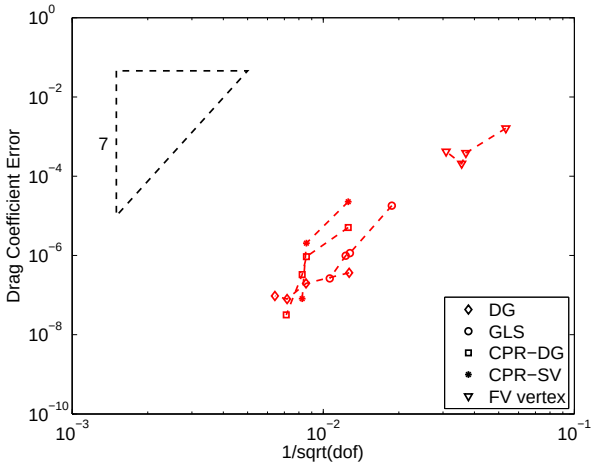

(f) p=3, adapted

Figure 4. Non-lifting, inviscid NACA 0012 drag coefficient error vs. element size

The output of interest is the drag coefficient of the airfoil; errors with respect to $1/\sqrt{dof}$ are plotted in Figure 4 for both non-adapted (left) and adapted (right) meshes. If the problem were smooth, as for the Gaussian bump, the output error convergence rates would be set by the solution order, specifically h^{2p+1}. However, since the trailing edge is inappropriately refined, the solution irregularity limits the asymptotic

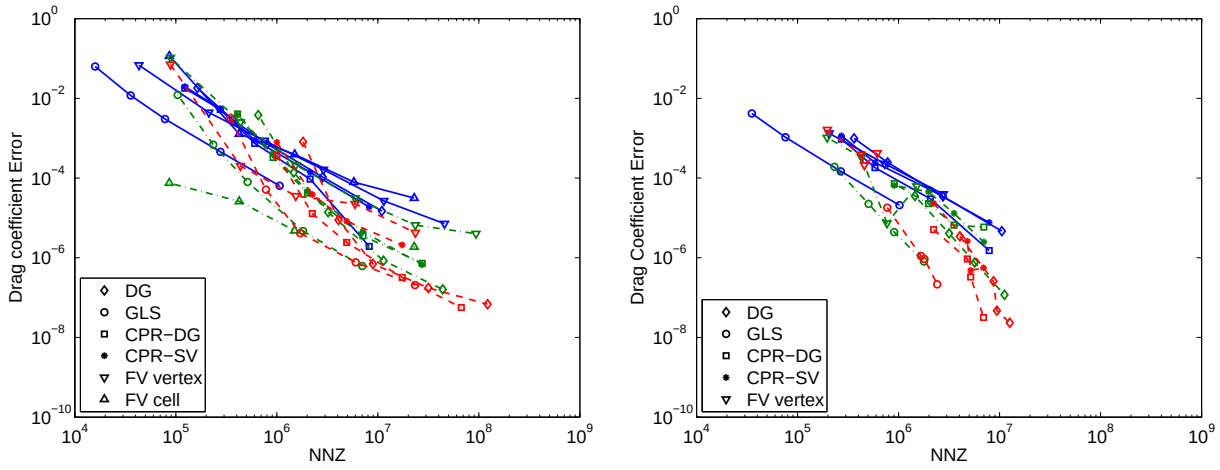

(a) Non-adapted (b) Adapted

Figure 5. Non-lifting, inviscid NACA 0012 drag coefficient error vs. NNZ ($p = 1, 2, 3$)

convergence rate, and there is no benefit of using higher-order methods; for $p = 2$ and $p = 3$, no discretization achieves the optimal convergence rate on non-adapted meshes.

Using the adapted meshes, nearly all discretizations achieve convergence rates of approximately h^{2p+1}. Unlike the results for the smooth Gaussian bump, the range of error levels for a given degree of freedom is approximately one to two orders of magnitude. The non-adaptive and adaptive results produce similar error values, but the degrees of freedom of the adapted meshes are lower than that of the non-adapted meshes for a given error level. Although there is little consistency as to which method achieves the lowest error for a given mesh and solution order, GLS is often a more competitive discretization in terms of accuracy per degree of freedom.

Let us compare the error convergence results for DG and CPR-DG in the context of approximability and stability. These two schemes have the same approximation spaces on a given mesh, thus the approximability of the discretizations are identical. Comparing the results obtained on the uniformly refined and adaptive meshes in Figure 4, we note that the difference in the output error between DG and CPR-DG is considerably larger for the adapted meshes. This result suggests that the CPR-DG method is less stable than the classical DG scheme on strongly graded meshes. In fact, the difference is largest for the $p = 1$ discretization, where CPR-DG uses the fewest number of the collocation points and its difference to the quadrature-based DG method is largest. For this particular problem, the quadrature-based formulations (DG and GLS) appear to be more stable on the highly graded meshes, producing smaller errors than the collocation-based schemes (CPR-DG and FV).

Figure 5 gives the non-adaptive and adaptive drag coefficient errors against the number of non-zeros in the Jacobian. Regardless of the solution order or mesh construction, GLS is clearly superior in terms of error per non-zero entry in the Jacobian.

C. Lifting, subsonic, inviscid flow over a NACA 0012 airfoil ($M_\infty = 0.5, \alpha = 2°$)

This case looks at the ability to solve lifting, subsonic, inviscid flow over a NACA 0012 airfoil. The freestream flow has a Mach number of $M_\infty = 0.5$ and an angle of attack of $\alpha = 2°$. The effect of the boundary conditions on the solution is more significant under lifting conditions, and the true drag value must be approximated in order to compute the error; the drag obtained from the mesh with the highest degrees of freedom (a combination of the number of elements and solution order) is used as the reference drag value separately for each discretization.

The problem was solved on both non-adapted and adapted meshes. The non-adapted meshes are the same non-adapted meshes for the non-lifting case, and new adapted meshes were generated for each solution order and number of elements under lifting conditions. The non-adaptive (left) and adaptive (right) results are shown in Figure 6. For the same reasons as under non-lifting conditions, little difference is seen between the adaptive and non-adaptive results when a lower solution approximation, $p = 1$, is used. Adapted meshes

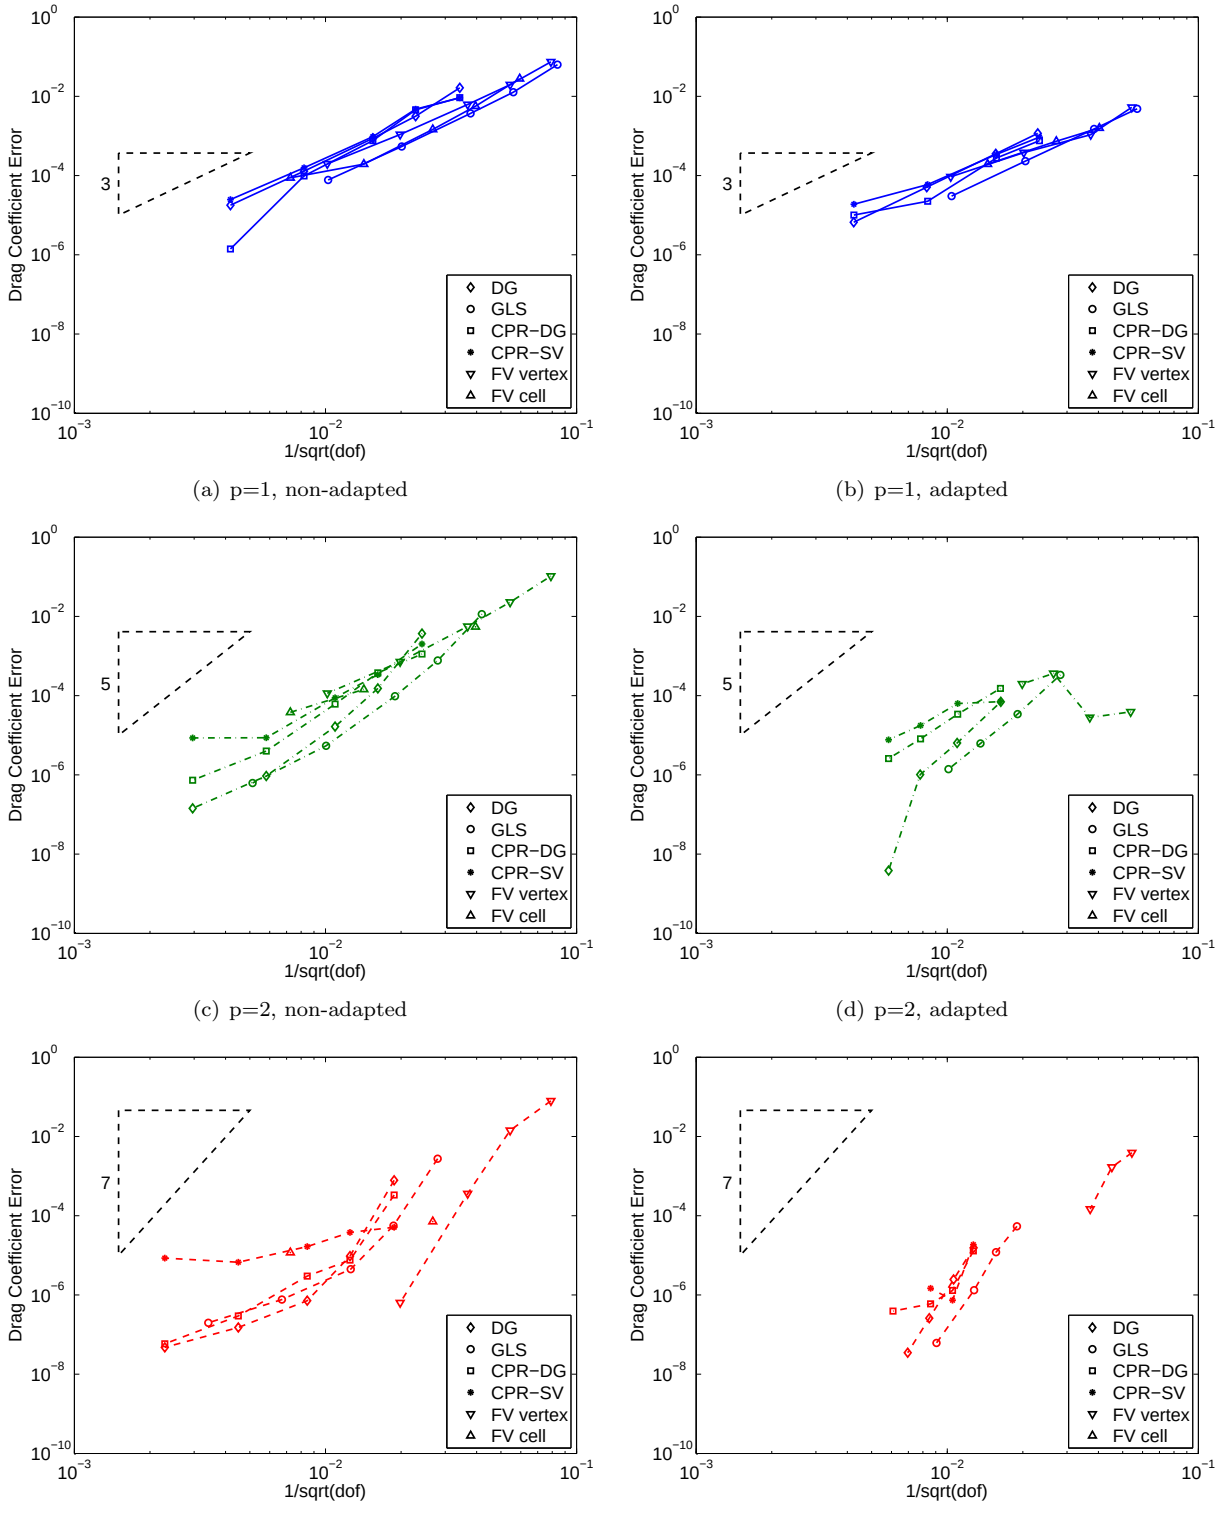

(a) p=1, non-adapted

(b) p=1, adapted

(c) p=2, non-adapted

(d) p=2, adapted

(e) p=3, non-adapted

(f) p=3, adapted

Figure 6. Lifting, inviscid NACA 0012 drag coefficient error vs. element size

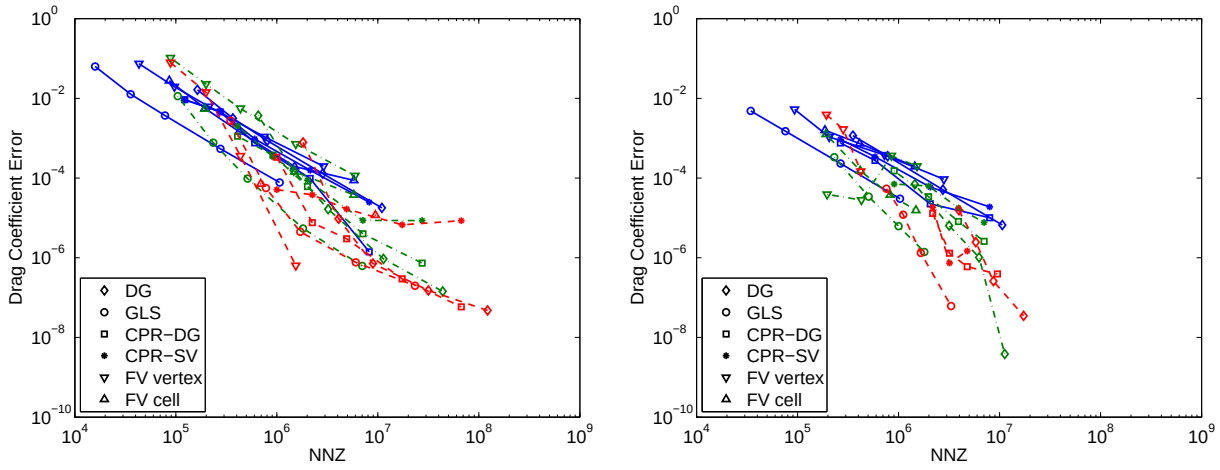

(a) Non-adapted (b) Adapted

Figure 7. Lifting, inviscid NACA 0012 drag coefficient error error vs. NNZ ($p = 1, 2, 3$)

improve convergence rates and the accuracy per degree of freedom. Although discretizations that require fewer degrees of freedom on a given mesh (GLS and finite volume methods) often result in higher error levels, their efficiency is comparable with the other methods, as they have fewer degrees of freedom. The large spread in error seen in the previous non-lifting case is again seen in these lifting results.

Figure 7 gives the non-adaptive and adaptive drag coefficient errors against the number of non-zeros in the Jacobian. Higher-order finite volume methods have less expensive Jacobian matrices for the non-adapted meshes, while GLS is significantly cheaper than most discretizations, regardless of the mesh construction and solution orders.

D. Non-lifting, supersonic, inviscid flow over a diamond airfoil ($M_\infty = 2.0$)

Finally, results are presented for non-lifting, supersonic, inviscid flow over a diamond airfoil with 5 percent thickness and a freestream Mach number of $M_\infty = 2$.

Using shock-expansion theory, the exact pressure drag on the airfoil can be computed and used to calculate the drag error. Figures 8 and 9 give the drag error with respect to element size and NNZ respectively. It is clear from these results that no method benefits from using a polynomial order greater than one. Although asymptotic convergence is not yet reached for most supersonic diamond airfoil results, in general, Figure 8 shows how the presence of shocks limits the overall convergence rates to approximately 1 with respect to h. Because the mesh is inappropriately refined near the shock, the flow features (shocks) are difficult to capture. If the problem were solved on a mesh where anisotropic elements were aligned with the shock and highly graded away from the shock, the flow features could be captured with increased efficiency. Adaptive results for the supersonic diamond airfoil were unable to make it into this study; these results will be included in future work.

IV. Conclusion

This paper compares the performance of higher-order methods for unstructured grids on a set of fundamental aerodynamic applications. The particular methods in the study include: Galerkin Least-Squares, classical discontinuous Galerkin finite element, CPR formulations of the discontinuous Galerkin finite element and spectral volume, cell-centered finite volume, and vertex-centered finite volume methods. These results focus on the accuracy and efficiency of each method when solving subsonic and supersonic, inviscid flows.

Flow over a Gaussian bump shows that all methods exhibit comparable levels of accuracy and obtain optimal convergence rates with mesh refinement for smooth flows. Subsonic flow over a NACA 0012 airfoil introduces a geometry-induced singularity at the trailing edge. When uniformly refined meshes are used for

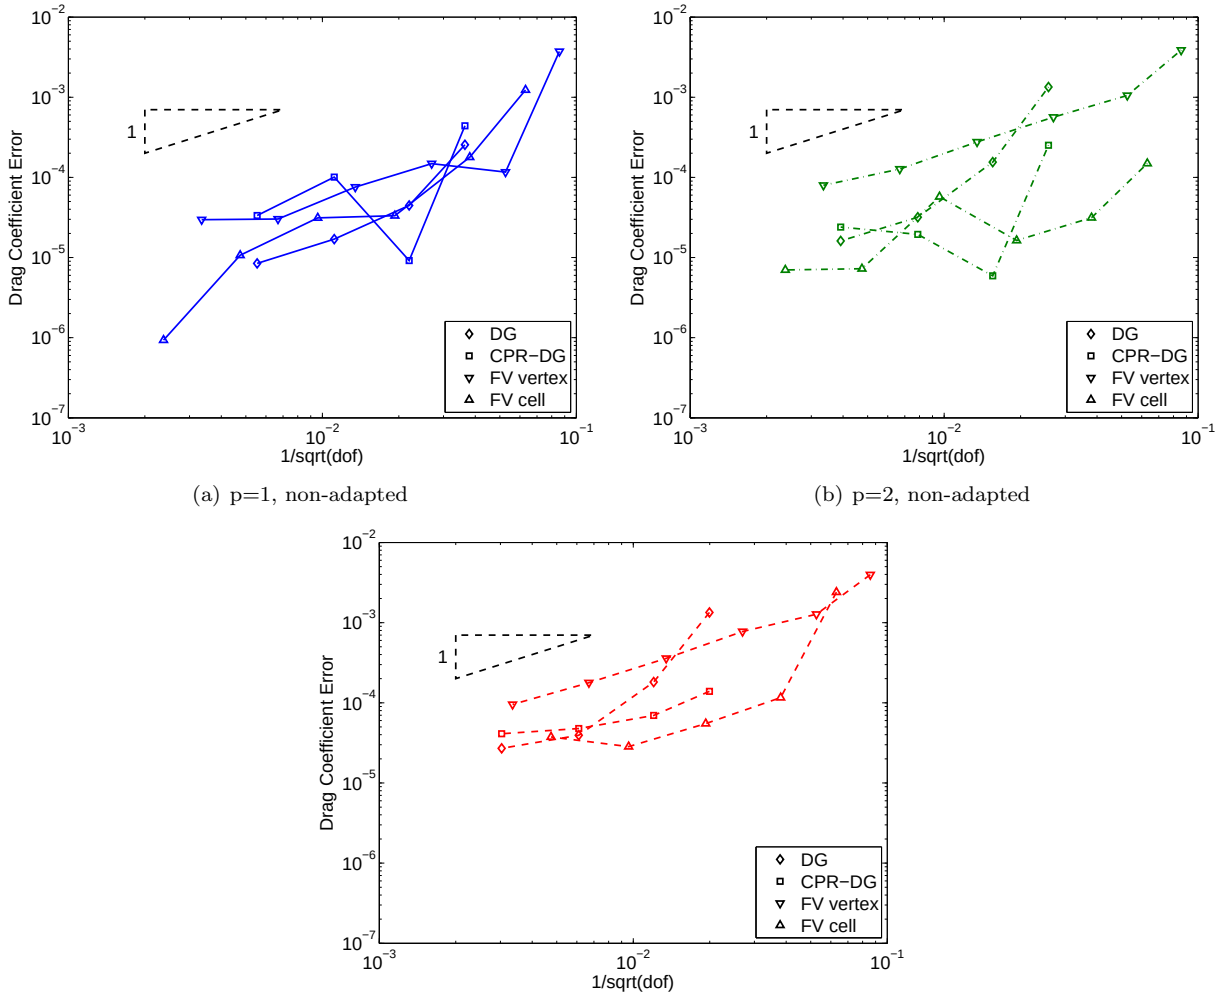

(a) p=1, non-adapted

(b) p=2, non-adapted

(c) p=3, non-adapted

Figure 8. Non-lifting, supersonic, inviscid diamond airfoil drag coefficient error vs. element size

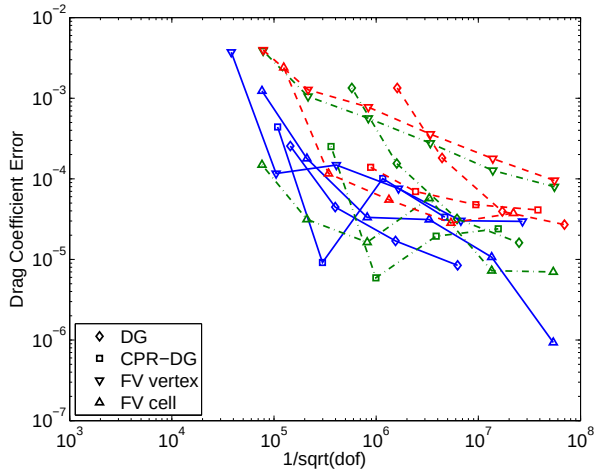

Figure 9. Non-lifting, supersonic, inviscid diamond airfoil drag coefficient error error vs. NNZ ($p = 1, 2, 3$)

a problem with solution irregularities, the error is dominant at the singularities, and convergence rates are limited, especially for higher-order approximations. Adapted meshes were generated to reduce the error by increasing mesh resolution near the singularity. In general, all higher-order methods were able to recover optimal convergence rates using the adapted grids. It is likely that the large range in error between methods that results from solving on adapted meshes are due to the strong grading near the trailing edge, which reveals different stability properties of the discretizations.

There are a number of areas that will be addressed as this study continues. Firstly, a more complete discussion of supersonic flow will be covered; supersonic flow over a diamond airfoil reiterated the idea that irregularities in the flow limit the convergence rates when using uniformly refined meshes. Secondly, we will test how well these methods discretize the Navier-Stokes equations by solving viscous flow over a flat plate and a NACA 0012 airfoil. Future studies will also include a more detailed analysis of the role stability and approximability on the error produced by different methods on strongly graded meshes. Because the adapted meshes in this study are generated based on solutions and error estimates from the classical DG method, it is important to understand how the different discretizations handle certain mesh characteristics that are requested by one particular adaptive scheme.

References

[1]Hughes, T. J. R., Franca, L. P., and Hulbert, G. M., "A new finite element formulation for computational fluid dynamics: VIII. The Galerkin/least-squares method for advective-diffusive equations," *Computer Methods in Applied Mechanics and Engineering*, Vol. 73, 1989, pp. 173–189.

[2]Yano, M., *Massively Parallel Solver for the High-Order Galerkin Least-Squares Method*, Masters thesis, Mass. Inst. of Tech., CDO, May 2009.

[3]Roe, P. L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372.

[4]Wang, Z. and Gao, H., "A unifying lifting collocation penalty formulation including the discontinuous, spectral volume/difference methods for conservation laws on mixed grids," *Journal of Computational Physics*, Vol. 228, 2009, pp. 8161–8185.

[5]Gao, H. and Wang, Z., "A High-order Lifting Collocation Penalty Formulation for the Navier-Stokes Equations on 2D Mixed Grids," AIAA 2009-3784, 2009.

[6]Nejat, A. and Ollivier-Gooch, C., "A High-Order Accurate Unstructured Finite Volume Newton-Krylov Algorithm for Inviscid Compressible Flows," *Journal of Computational Physics*, Vol. 227, No. 4, 2008, pp. 2592–2609.

[7]Ollivier-Gooch, C., Nejat, A., and Michalak, C., "On Obtaining and Verifying High-Order Finite-Volume Solutions to the Euler Equations on Unstructured Meshes," Vol. 47, No. 9, 2009, pp. 2105–2120.

[8]Fidkowski, K. and Darmofal, D., "Output error estimation and adaptation in computational fluid dynamics: Overview and recent results," AIAA 2009-1303, 2009.

[9]Yano, M., Modisette, J. M., and Darmofal, D., "The importance of mesh adaptation for higher-order discretizations of aerodynamic flows," AIAA 2011–0000, June 2011.

[10]Becker, R. and Rannacher, R., "A feed-back approach to error control in finite element methods: Basic analysis and examples," *East-West Journal of Numerical Mathematics*, Vol. 4, 1996, pp. 237–264.

[11]Becker, R. and Rannacher, R., "An optimal control approach to a posteriori error estimation in finite element methods," *Acta Numerica*, edited by A. Iserles, Cambridge University Press, 2001.

[12]Venditti, D. A. and Darmofal, D. L., "Anisotropic grid adaptation for functional outputs: Application to two-dimensional viscous flows," *Journal of Computational Physics*, Vol. 187, No. 1, 2003, pp. 22–46.

[13]Fidkowski, K. J. and Darmofal, D. L., "A triangular cut-cell adaptive method for higher-order discretizations of the compressible Navier-Stokes equations," *Journal of Computational Physics*, Vol. 225, 2007, pp. 1653–1672.

[14]Hecht, F., "BAMG: Bidimensional Anisotropic Mesh Generator," 1998, `http://www-rocq1.inria.fr/gamma/cdrom/www/bamg/eng.htm`.

[15]Oliver, T. A., *A Higher-Order, Adaptive, Discontinuous Galerkin Finite Element Method for the Reynolds-averaged Navier-Stokes Equations*, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2008.

[16]Venkatakrishnan, V., Allmaras, S. R., Kamenetskii, D. S., and Johnson, F. T., "Higher Order Schemes for the Compressible Navier-Stokes Equations," AIAA 2003-3987, 2003.