A High-Order Accurate Unstructured GMRES Solver for Poisson's Equation

Amir Nejat* and Carl Ollivier-Gooch†

Department of Mechanical Engineering, The University of British Columbia, BC V6T 1Z4, Canada

Email: nejat@mech.ubc.ca, cfog@mech.ubc.ca

ABSTRACT

In this paper, we present a fast higher-order accurate unstructured solver for Poisson's equation as a model problem for diffusive phenomena. The Krylov GMRES technique in matrix free form has been used to achieve the efficient convergence rate. By computing the solution over a series of coarse and fine meshes with different orders of accuracy, we show the possibility of reducing computational cost required for a given level of solution accuracy using higher-order methods. Numerical experiments also demonstrate dramatic improvement in convergence CPU-time with increasing size of the Krylov-subspace.

1. INTRODUCTION

Recent structured mesh flow solver research shows that for practical levels of accuracy using a higher-order accurate method can be more efficient both in terms of solution time and memory usage than a second-order scheme [15]. With higher-order accurate methods, the cost of flux computation, integration schemes and the other associated numerical calculation increase per control volume, but as we can use a coarser mesh, computation time and memory are saved overall and accuracy can be increased as well. To extend this favorable behavior, we want to apply high-order accurate methods on unstructured meshes to take advantage of well-documented flexibility in mesh generation and adaptation for unstructured meshes. High-order accurate method development for unstructured mesh finite-volume computations has historically focused on advective fluxes and are reasonably well established [2, 3, 4, 8]. For diffusive fluxes a new approach has been implemented successfully in Ref [9] by computing the solution gradients required for flux calculation from the solution reconstruction.

In spite of considerable progress in high-order unstructured algorithms, application of these methods for physically complicated flows is still a challenge due to slow rate of convergence. Degradation in convergence rate could increase the solution time noticeably and eliminate the efficiency benefits of higher-order unstructured discretization. Therefore convergence acceleration becomes the key issue for enabling practical use of higher-order discretizations of the fluid flow equations over unstructured topology.

Regardless of discretization technique, finding steady solutions for fluid flow problems requires the solution of a large system of non-linear algebraic equations. An important iterative technique for solving such systems of equations is the Newton family of methods. They have been used in CFD since the late 80's [1, 14] and are considered an attractive approach for solution convergence of steady flows due to their property of quadratic convergence, which is quite desirable. However, the cost of each iteration in terms of CPU-time and memory usage for pure Newton method is relatively high. Approximate or inexact Newton method can have satisfactory convergence behavior and lower cost per iteration [10]. In this type of the method the linear system of equations is solved approximately at each iteration. The formation and linearization of the Jacobian and the choice of preconditioner in iterative methods all affect the convergence rate. As a result there are a vast variety of approximate Newton or Quasi-Newton methods. The most common method especially for unstructured meshes is Newton-Krylov which uses the Krylov family of iterative solvers (mostly GMRES[13]).

For structured meshes, Pueyo and Zingg [11] presented an efficient matrix-free Newton-GMRES solver using multi-grid for steady aerodynamic flows and conducted some parametric studies. In a follow-up study [12] they investigated the efficiency of different strategies for Jacobian approximation and preconditioning for a variety of compressible

* Graduate Research Assistant.
† Assistant Professor.

inviscid, laminar and turbulent flows using GMRES. Barth and Linton [5] successfully applied both full-matrix and matrix-free GMRES iteration for computing compressible fluid flow over unstructured meshes. Blanco and Zingg [6] compared the use of two variations on quasi- and full-Newton method for solution of the Euler equations on unstructured grids. Standard (full-matrix) and matrix-free version of GMRES technique both were used in full Newton scheme and finally they proposed an hybrid strategy which combines the strengths of both full- and quasi-Newton implementations.

In this paper, we focus on efficient convergence for the Poisson equation for high-order accurate unstructured mesh discretizations using the matrix-free GMRES technique. Our expectation is that experience with this model problem will be useful in improving convergence for higher-order accurate unstructured mesh viscous flow solutions. Also we discuss the benefits of higher-order accurate solution both in terms of accuracy and efficiency. In the next section the theory and numerical procedure are described briefly. The following sections introduce test cases and present their results, focusing on the performance of our approach. Finally, we summarize the benefits and limitations of our approach and provide a prospective for future work.

2. THEORY

Poisson's equation, which is a model equation for diffusion problems, can be written for each control volume as eq.(1) :

$$\frac{1}{A}\left[\oint \nabla T.\hat{n}\,ds - \iint S(x,y)\,dA\right] \circ R(T) = 0 \quad (1)$$

In our discretization scheme, a higher-order accurate least-square reconstruction procedure [8] has been used in the interior of the domain. Eq. (1) includes diffusive fluxes which requires the gradient of the reconstructed solution. We compute these fluxes at control volume boundaries by averaging the gradients computed in each control volume. Having computed the fluxes, we integrate along each control volume boundary using Gauss quadrature integration technique with the proper number of points to get the required order of accuracy for the flux integral. Imposing the boundary conditions also has been done to high-order accuracy [9].

Because the Poisson equation is linear and intrinsically steady, solution of the problem requires, in principle, forming and solving a single system of linear equations. Since the flux integral is linear in T, discretization of eq.(1) over an unstructured domain results in a system of linear equation in the form of :

$$\left[\frac{\partial R}{\partial T}\right] T = \left[\frac{1}{A}\iint S(x,y)\,dA\right] \quad (2)$$

where $\partial R/\partial T$ is the Jacobian matrix. Making an initial guess at the solution T_0 and defining $DT = T - T_0$, we get:

$$\left[\frac{\partial R}{\partial T}\right] DT = -R(T_0) \quad (3)$$

We turn to iterative methods to solve this linear system. The linear system is asymmetric both in fill and in values making GMRES [13] the obvious first choice in iterative solvers. GMRES(k) minimises the norm of the residual vector over a Krylov subspace of size k. Increasing the number of search directions (increasing k) gives a more accurate update for each GMRES iteration. However, memory usage increases linearly and computation time quadratically with k. Also maintaining the orthogonality of the vectors in the Krylov subspace becomes more difficult with increasing k. So, there is a practical limit for increasing subspace size; we will discuss this issue in the result section.

The linear systems arising from third- and fourth-order accurate discretizations of the Laplacian have relatively large numbers of non-zero entries, 15-20 per row for 2D fourth-order discretizations. Because of the size of these matrices and the difficulty in computing their entries analytically, we have elected to use an entirely matrix-free implementation of GMRES. Since the GMRES algorithm only needs matrix-vector products and these products can be computed by matrix-free approach, matrix-free GMRES is a very attractive technique for dealing with the complicated Jacobian matrices resulting from higher-order discretization. Furthermore, this approach saves memory usage and in some cases even computation time. To obtain the matrix-vector product we take the directional derivative with respect to a nomalized vector V using numerical forward difference approximation [5].

$$\left[\frac{\partial R}{\partial T}\right].V \approx \frac{R(T + eV) - R(T)}{e} \quad (4)$$

In eq.(4) e is a small scalar value, typically chosen equal to the square root of machine precision. Thus we can calculate the required matrix-vector products simply by evaluating the residual based on a perturbed solution.

3. TEST CASES

To investigate the convergence characteristics and performance of the GMRES solver with a higher-order unstructured discretization, a series of numerical experiments has been conducted. The geometry is an annular segment and we use four different unstructured meshes with 62, 246, 987 and 3846 control volumes (Figure 1). The meshes for these test cases were generated using a guaranteed-quality meshing scheme [6]; all triangles had angles between 30º and 120º. To obtain a higher-order accurate solution, we must consider the curvature of curved domain boundaries instead of modelling them as linear segments. The grid generator provides the solver with necessary curvature information for implementation of boundary conditions and related integration over boundary cells. For all meshes, numerical solutions of Poisson equation have been computed using second, third and fourth order accurate reconstruction.

4. RESULTS

To assess solution accuracy, we plot the error in the computed solution with respect to the exact solution versus mesh size. Figure 2 shows the L2 norm of error in solution for all meshes. Second and fourth order results both demonstrate the expected behaviour. For the "third-order" solution the actual order of accuracy that we get is about second-order. This result is not unexpected; Ollivier-Gooch and Van Altena showed that a second order term appears in the truncation error for the discrete Laplacian, despite using a third-order accurate quadratic reconstruction [9].

Figure 3 shows the convergence history of the solution versus the number of GMRES outer iterations for the finest mesh (3846CV) and Krylov subspace size of 60. As expected, more iterations are needed for convergence when we use higher-order methods reflecting higher condition number for higher order of accuracy. This is added to the extra computational cost which reconstruction and integration techniques take for higher-order methods. Nevertheless, high-order methods are still more efficient for a given accuracy, as shown in Figure 4.

To give insight into relative efficiency in getting accurate solutions, Figure 4 plots the L2 norm of error versus the CPU-time for all meshes and orders of accuracy. This Figure indicates clearly the advantage of higher-order accurate solution with respect to the second-order cases for the Poisson equation. Using fourth-order method we can get better accuracy with considerable savings in time using a coarser mesh. For example, we compare 246CV-fourth-order with 987CV-third-order and 3846CV-second-order cases. The third-order case has nearly the same error that the second-order one has, but it has converged about three times faster. The fourth-order case gives us more accurate solution with more than one order of magnitude faster convergence than second-order. Likewise for a given CPU time, the fourth-order scheme is always most accurate. The fourth-order trend line has the maximum slope indicating that the advantage of the fourth-order scheme will grow with the mesh refinement.

Finally we want to study the effect of size of the Krylov subspace, k, on the efficiency of the GMRES solver. As we use a larger subspace size, we expect a more accurate update in each outer iteration. So the number of GMRES iterations should decrease, but the expense of each iteration will increase with increase in k. So there is trade-off issue based on the number of residual evaluations for increasing the Krylov subspace. In Figure 5 convergence history has been compared for different subspace sizes. Computational cost is compared on the basis of the number of residual evaluations required. Figure 5 shows that an increase in k decreases the total number of residual evaluations. Increasing k from 20 to 60 reduced the number of residual evaluations by half. However, we are still below the aforementioned trade-off point. As subspace size increases, eventually the number of residual evaluations should increase by increasing k. In addition, using large subspaces needs more memory, which finally would be a limiting factor as well.

Figure 6 compares the run-time for different order of accuracy versus the size of the Krylov subspace. Using larger k decreases the CPU-time, but this effect diminishes with the size of the subspace. The effect of increasing k is more noticeable for higher-order method especially for the fourth-order case, which seems to benefit from a more accurate update per GMRES iteration .

5. CONCLUDING REMARKS

A matrix-free GMRES algorithm has been successfully used for higher-order computation of solution of Poisson's equation. Accuracy-Run time diagram showed the advantage of high-order computation both in terms of accuracy and efficiency. The effect of Krylov subspace size was investigated: using a larger subspace size can reduce the computation cost especially for higher-order computations.

As Poisson's equation is simple and linear and we have used simple geometry, the resultant Jacobian

matrix should be well-conditioned; the convergence rate of our GMRES solver confirms this. However as GMRES convergence is sensitive to condition number of the Jacobian matrix, and in applied problems (especially for nonlinear equations) we are dealing with ill-conditioned Jacobians, using a preconditioner for GMRES becomes necessary. As we extend this research to inviscid and viscous fluid flows, we will concentrate on developing effective preconditioners for GMRES with high-order accurate unstructured mesh discretizations.

ACKNOWLEDGEMENTS

This research was supported in part by the Canadian Natural Science and Engineering Research Council under Grant OPG-0194467and in part by a University Graduate Fellowship from the University of British Columbia.

REFERENCES

1. Bailey, H. E. and Beam R. M., Newton's Method Applied to Finite Difference Approximations for Steady State Compressible Navier-Stokes Equations, Journal of Computational Physics, Vol. 108, pp. 108-127, 1991.

2. Barth, T. J. and Fredrickson, P. O., Higher-Order Solution of the Euler Equations on Unstructured Grids Using Quadratic Reconstruction, AIAA conference paper 90-0013, 1990.

3. Barth, T. J., Aspects of Unstructured Grids and Finite-Volume Solvers for the Euler and Navier-Stokes Equations, in Unstructured Grid Methods for Advection-Dominated Flows, AGARD-R787, 1992.

4. Barth, T. J., Recent Development in High-Order K-Exact Reconstruction on Unstructured Meshes, AIAA conference paper 93-0668, 1994.

5. Barth, T. J. and Linton S.W., An Unstructured Mesh Newton Solver for Compressible Fluid Flow and Its Parallel Implementation, AIAA conference paper 95-0221, 1995.

6. Blanco, M., and Zingg D. W., A Fast Solver for the Euler Equations on Unstructured Grids Using a Newton-GMRES Method, AIAA conference paper 97-0331, 1997.

7. Boivin, C., and Ollivier-Gooch C., Guranteed-quality Triangular Mesh Generation for Domains with Curved Boundaries, International Journal for Numerical Methods in Engineering, Vol. 55, pp.1185-1213, 2002

8. Ollivier-Gooch, C., Quasi-ENO Schemes for Unstructured Meshes Based on Unlimited Data-Dependent Least Squares Reconstruction., Journal of Computational Physics, Vol. 133, pp. 6-17, 1997.

9. Ollivier-Gooch, C., and Van Altena, M., A Higher-Order Accurate Unstructured Mesh Finite-Volume Scheme for the Advection-Diffusion Equation, Journal of Computational Physics, Vol. 181, pp. 729-752, 2002.

10. Orkwis, P. D., Comparison of Newton's and Quasi-Newton's Method Solvers for Navier-Stokes Equations, AIAA Journal, Vol.31, pp. 832-836, 1993.

11. Pueyo, A., and Zingg, D. W., An Efficient Newton-GMRES Solver for Aerodynamic Computations, AIAA conference paper 97-1955, 1997.

12. Pueyo, A., and Zingg, D. W., Improvement to a Newton-Krylov Solver for Aerodynamic Flows, AIAA conference paper 98-0619, 1998.

13. Saad, Y. and Schultz, M. H., A Generalized Minimal Residual Algorithm for Solving Non-Symmetric Linear Systems, SIAM J. Sci., Stat. Comp. Vol. 7, pp. 856-869, 1986.

14. Venkatakrishnan, V., Newton Solution of Inviscid and Viscous Problems, AIAA conference paper 88-0413, 1988.

15. Zingg, D. W., De Rango, S., Nemec, M. and Pulliam T. H., Comparison of Several Spatial Discretizations for the Navier-Stokes Equations, Journal of Computational Physics, Vol.160, pp.683-704, 2000.

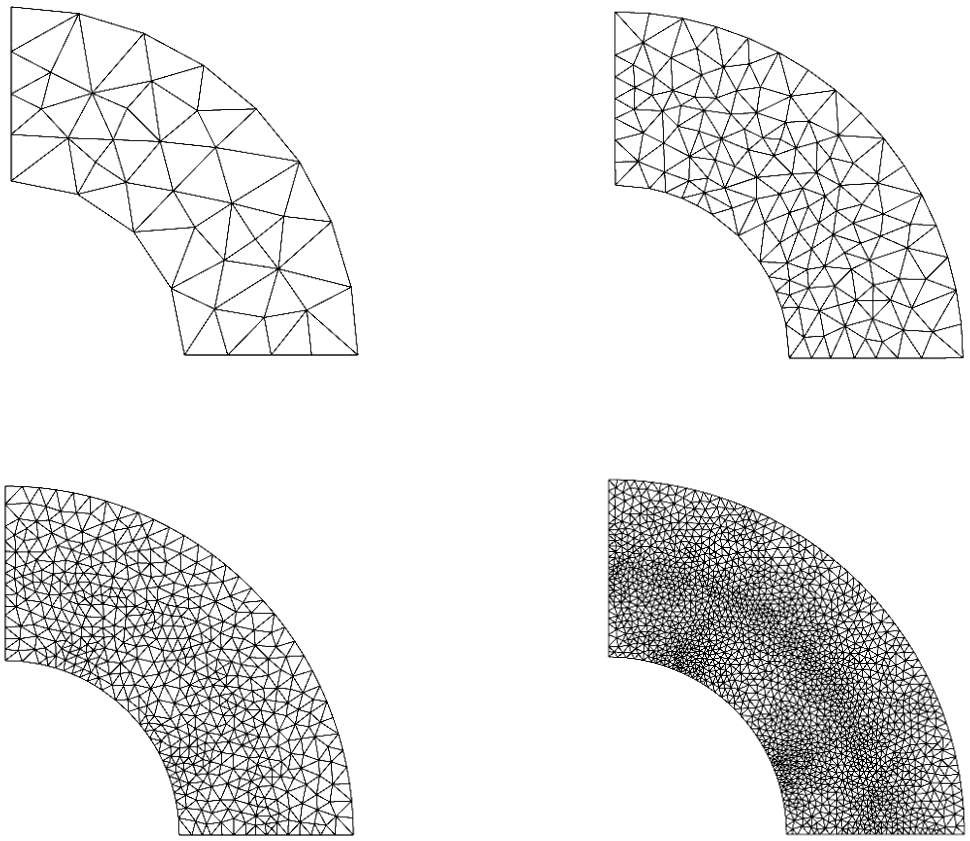

Figure 1. Different meshes for annulus geometry with 62, 246, 987, 3846 control volumes.

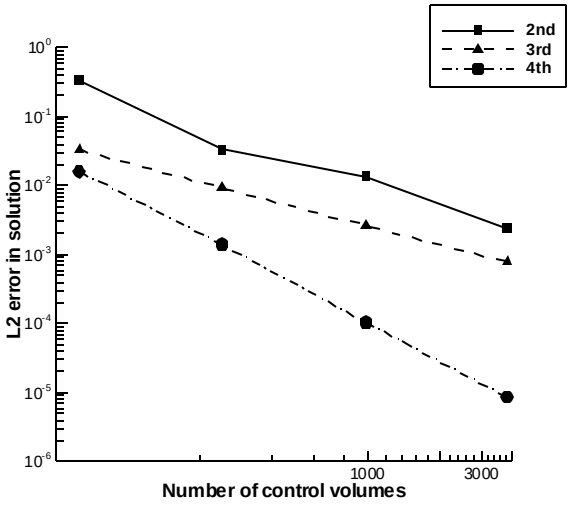

Figure 2. Accuracy assessment versus number of control volumes.

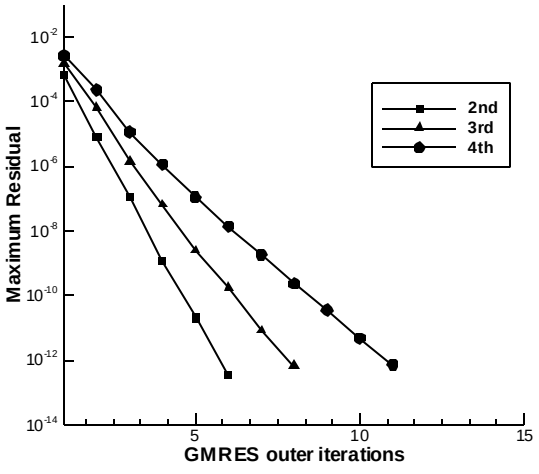

Figure 3. Convergence history for the finest mesh.

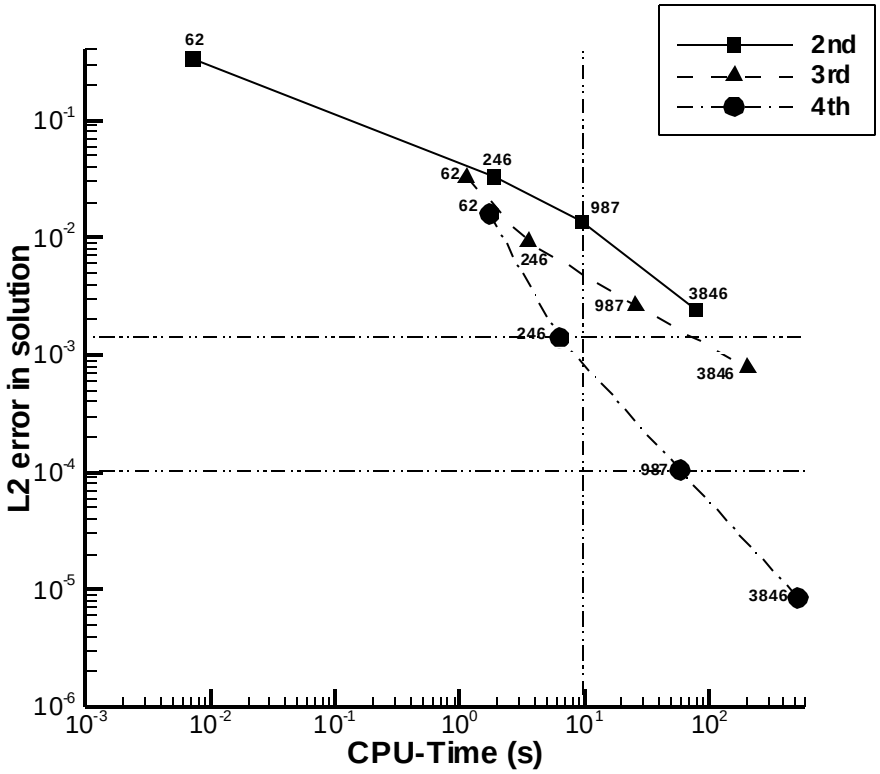

Figure 4. Accuracy-Run time diagram for different meshes.

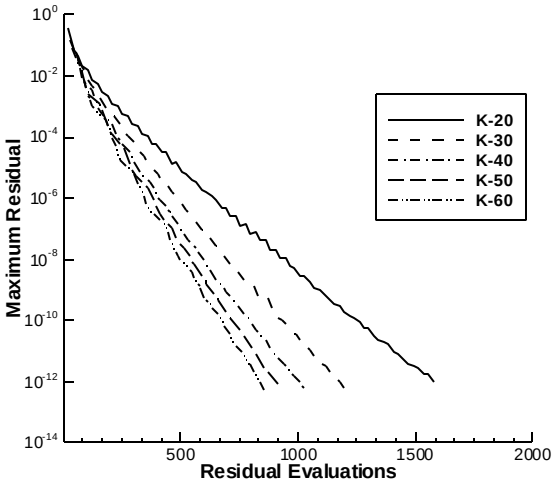

Figure 5. Comparison of convergence history based on size of the Krylov-subspace for the finest mesh and fourth-order of accuracy.

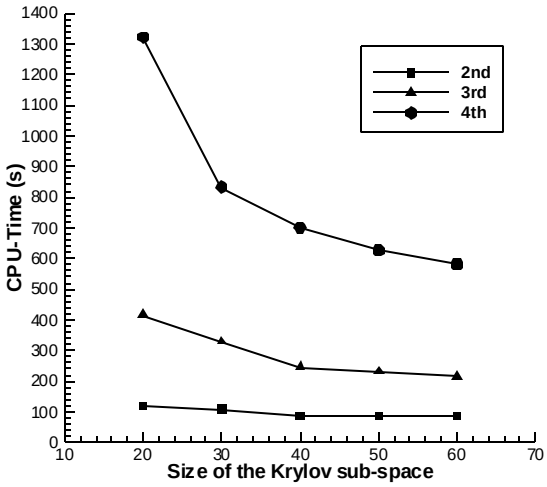

Figure 6. Run time versus Krylov subspace size for the finest mesh.