

Output-Based Mesh Adaptation for High-Speed Flows

James G. Coder¹, Benjamin L. S. Couchman², Marshall C. Galbraith²,
Steven R. Allmaras², and Nick Wyman³
Corresponding author: jcoder@utk.edu

¹ University of Tennessee, Knoxville, TN, USA

² Massachusetts Institute of Technology, Cambridge, MA, USA

³ Cadence Design Systems, Fort Worth, TX, USA

Abstract: High-speed Computational Fluid Dynamics calculations often rely on structured meshes in order to facilitate aligning the mesh with shocks either manually or via some semi-automated process. The shock alignment of the mesh is often considered critical in order to obtain sufficiently accurate outputs such as integrated surface heat transfer. However, creating shock-aligned structured meshes with complex geometries and/or multiple interacting shocks is challenging, and may not even be feasible for some configurations. Unstructured meshes offer greater flexibility focusing mesh resolution to capture complex geometry and flow features. This work demonstrates that integrated pressure drag computed using automated output-based adapted unstructured meshes can achieve higher degrees of accuracy compared shock fitted structured meshes for a canonical high-speed test case.

Keywords: Computational Fluid Dynamics, High-Speed Flows, Mesh Adaptation.

1 Introduction

Properly resolving critical flow features for high-speed aerodynamic flow calculations is critical for accurate predictions of outputs of interests, such as lift, drag, or surface heat transfer, just to name a few. This raises an underlying question then of what aspects of the flow are the most critical to resolve. Traditionally, much emphasis is placed on resolving the shockwave, especially the leading shockwaves of the geometry. There is strong justification for this approach, as these shockwaves fully set the flow conditions for everything downstream on the geometry. Poor resolution of the shockwave and/or misalignment of the mesh can introduce spurious, non-physical entropy gradients that contaminate the flow and can lead to numerical phenomena such as carbuncles. Consequently, there is a strong preference for generating grids that are well-aligned with the shockwaves. For structured meshes and simple geometries that lack strong interactions of the shockwaves, this may be accomplished with straightforward heuristics, but for complex geometries and unstructured grids, the process can become very labor intensive with significant user intervention across multiple iterations.

Nevertheless, such feature-based adaptation is not always guaranteed to improve the measured results for output quantities of interest. One of the classical examples of this outcome is the “fish-tail” shock that airfoils exhibit at high transonic speeds. The shockwave in this scenario is very strong, but because it emanates from the trailing edge of the airfoil, it is completely decoupled from the drag functional. Therefore, increasing resolution around such a shock does not improve the predicted drag accuracy. Conversely, it can be found that the stagnation streamline has a leading-order impact on the drag functional, but the flow field typically does not trigger any feature-based grid refinement such as one would apply to a shockwave.

Output-based mesh adaptation, however, offers a robust alternative to feature-based adaptation. Output-based methods use the adjoint of a target functional, such as lift, drag, or heat transfer, and connects it to the error of the discretization. In this way, the grid can be refined in regions of the flow with the strongest influence on the output quantity of interest, and coarsened in regions with the weakest influence to keep the

total degrees of freedom in the solution near a target level. In this paper, output-based mesh adaptation is applied to two high-speed test cases, one inviscid and one viscous, using a finite-element-based flow solver and multiple mesh generation tools. The results for the inviscid case are compared with finite-difference solutions obtained using a more industrial approach with structured, shock-fitted grids. It is demonstrated from these results that the output-based adapted unstructured meshes perform as well or better than the shock-fitted structured grids, and reveal influential flow features otherwise missed by shock fitting alone.

2 Test Case

The primary test case used in this study is the inviscid cylinder case described by Bova et al. [1]. They considered flow of ideal nitrogen at $M_\infty = 20$. As the present work is concerned with the inviscid behaviors at these flow conditions, the exact composition of the gas does not affect the equations so long as $\gamma = 1.4$. For consistency across all methods considered, the domain is taken to be the windward half of the cylinder with the outer boundary set to be a constant 2 radii from the surface. With the high free-stream Mach number, this boundary location is sufficient to not intersect with the shock structure and cause undue influence on the solution.

A viscous test case was also considered based on the Gnoffo's cylinder [2]. For this case, the free-stream Mach number is $M_\infty = 17.6$ and the unit Reynolds number based on radius is approximately $Re = 3.76 \times 10^5$. However, difficulties were encountered in converging the OVERFLOW solver, so only partial results are shown in this paper.

3 Flow Solvers

3.1 Solution Adaptive Numerical Simulator (SANS)

The MIT SANS code base includes both stabilized continuous (CG) and discontinuous (DG) Galerkin finite-element methods. Results are presented for inviscid flows using VMSSD (Variational Multiscale Method with Discontinuous Subscales)[3, 4] discretization. While SANS supports higher-order discretizations, only the 2nd-order accurate discretization with linear P1 solution and linear Q1 geometric elements are considered here.

A Homotopy Continuation (HTC) damped Newton's method is used to solve the nonlinear FEM system of equations where the complete residual is linearized using operator overloaded automatic differentiation[5]. The HTC method is used to improve the robustness of Newton's method by interpolating the between a simple L^2 projection residual and the discrete Euler residual. The final iteration of the Homotopy algorithm is a simple Newton solve. At each HTC iteration, the linear system is solved using direct LU via the PETSc[6, 7, 8]. The adjoint system is solved using the same linear solver as the primal linear solver. The linear and non-linear system residuals are driven to approximately machine zero.

3.2 OVERFLOW

The OVERFLOW code [9] is used here to provide representative, industrial-type flow solutions as a basis of comparison for the solution-adaptive, finite-element solutions. OVERFLOW is a structured, finite-difference solver of the compressible Navier-Stokes equations. The viscous terms can be suppressed so that the Euler equations can be solved instead. For this work, the flow equations are discretized using a 3rd-order accurate MUSCL reconstruction of the primitive variables to the half-node locations. The HLLE++ scheme [10] is then used to find the upwinded approximate Riemann flux. This flux is specially tailored for high-Mach number flows, as it features additional dissipation in the vicinity of shocks to damp the formation of carbuncles. The cut-in for the added dissipation is controlled by the DELTA parameter. The recommended value of $DELTA = 5$ from Ref. [10] is used here, as it was sufficient to damp the instability while still permitting machine-zero convergence.

4 Size Field Optimization and Mesh Generation

4.1 Metric Optimization via Error Sampling and Synthesis

The Metric Optimization via Error Sampling and Synthesis (MOESS) [11, 12, 13, 14] mesh adaptation framework (implemented in the SANS software package) is used to generate optimal adapted meshes. These meshes are optimal in the sense that an output error indicator is minimized for a given cost. The Dual Weighted Residual (DWR) method [15] provides an error estimate for an output functional of interest, which is then localized to allow the adaptation process to target areas with high output error. Since directly optimizing the discrete mesh is intractable, the minimization is instead performed on a continuous metric field: a smooth field composed of symmetric positive definite matrices $\mathcal{M}(x)$. This approach is possible due to the duality between a discrete mesh, \mathcal{T}_h , and a continuous mesh represented by a Riemannian metric field $\mathcal{M} = \mathcal{M}(x)|_{x \in \Omega}$, as demonstrated by Loissele [16, 17, 18].

The MOESS mesh adaptation problem in the continuous form is stated as

$$\mathcal{M}^* = \arg \min_{\mathcal{M} \in \mathcal{M}(\Omega)} \mathcal{E}(\mathcal{M}), \quad \text{s.t.} \quad \mathcal{C}(\mathcal{M}) \leq N, \quad (1)$$

where \mathcal{E} and \mathcal{C} are surrogate models for the error and cost respectively. Here, \mathcal{C} is defined as the number of Degrees of Freedom (DOF) in the mesh and is constrained by some target upper bound N . The error model is obtained by summing local error contributions:

$$\mathcal{E}(\mathcal{M}) = \int_{\Omega} \eta(\mathcal{M}(x), x) dx. \quad (2)$$

where the local error model, η , is assumed to be only a function of the local metric.

By definition, a metric conforming mesh has edge lengths e that are approximately unit under the metric, i.e., $e^T \mathcal{M}(x) e \approx 1$. To generate such a mesh, the metric field is passed to a metric-conforming mesh generator. This work uses `Pointwise` for isotropic mesh generation and both `avro` [19] and `refine` software packages are used for metric-conforming mesh generation.

To perform a mesh refinement study using the MOESS adaptation algorithm, a sequence of increasing target DOFs is first selected. Then, a set number of MOESS adaptation iterations are executed for each target DOF. Here and in Ref. [20], 50 adaptation iterations are used in order to ensure that the mesh adaptation has found an optimal mesh for any given target DOF. Even after multiple adaptations at the target DOF count, there is always some variation in the mesh. Hence, results presented are shown as the average over the last 5 adaptation iterations along with the actual values over the same iterations. Finally, the last mesh for a given target DOF is also used as the initial mesh for the next target DOF.

4.2 Pointwise

4.2.1 Isotropic Size Fields and Re-meshing

A size field is a discrete scalar field used to define the expected isotropic element size at any location over the domain [21]. A solver-neutral approach to mesh adaptation utilizes unstructured point clouds to augment the size field to improve the mesh avoiding unnecessary communication of solver data [22].

Mesh adaptation occurs by requesting a new mesh conformal to the updated size field and independent of the previous mesh. The re-meshing process importantly maintains the association with the supporting geometry. Mesh adaptation without conforming to the geometry, or even methods which conform to a discrete approximation of the geometry, add systemic errors into the solution. Re-meshing is “user-independent” as the `Pointwise` mesh generation software creates a mesh conforming to the updated size field from the original mesh topology, with all the geometry associativity, meshing parameters, and CFD solver attributes.

4.2.2 Boundary Layer Meshing

An advancing layer scheme is used in `Pointwise` to generate high aspect ratio elements orthogonal to wall boundaries. The layer generation scheme advances a layer point-by-point, which allows maximum flexibility to alter local topology to circumvent, for example, sharp concave regions where fixed topology would fail or

produce poor quality mesh. Because `Pointwise` re-meshes conformally to the updated size field, the surface mesh is made conformal prior to the boundary layer advancement resulting in a mesh capturing the desired size field transverse spacing in the viscous region.

4.3 avro

`avro`¹ is a dimension-independent anisotropic mesh adaptation package originally developed at MIT and has been demonstrated on up to 4D mesh generation problems [23]. It is based on local cavity operators inspired by the work of Coupez [24] and Loseille [25] and uses a combination of edge split, edge collapse, edge swaps, facet $((d - 1)$ -simplex) swaps and vertex smoothing to conform to a prescribed metric. While `avro` does not use a re-meshing strategy, the mesh modifications are performed to always maintain a valid mesh. Hence, the output mesh is always guaranteed to be valid upon termination of the mesh modification process. `avro` associates each mesh vertex with an EGADS [26] geometry entity in order to check the topological validity of each mesh operator and for projecting vertices to the geometry, and hence maintaining the original geometric fidelity defined by the EGADS kernel.

4.4 refine

The `refine` open source anisotropic grid adaptation mechanics package was developed at NASA². The current version uses the combination of edge split and collapse operations proposed by Michal and Krakos [27]. Vertex relocation is performed to improve adjacent element shape. A convex combination of these ideal vertex locations is chosen to yield a new vertex location update that improves the element shape measure in the anisotropic metric [28]. The vertices of elements with poor shape measures are also relocated with the nonsmooth optimization of Frietag and Ollivier-Gooch [29]. Geometry is accessed through the EGADS application program interface, and parallel execution is facilitated by EGADSLite [30].

4.5 Structured Mesh Generation and Shock Fitting

Single-block structured meshes for use with the `OVERFLOW` solver were generated parametrically based on the analytical definition of the cylinder surface. Points along the surface were distributed uniformly along the surface between $\theta = [\pi/2, 3\pi/2]$. Wall-normal points were distributed using an algebraic mapping function,

$$r_j = \frac{r_1 r_2}{r_2 + \xi^p (r_1 - r_2)} \quad (3)$$

along lines of constant θ , with

$$\xi = \frac{j - 1}{j_{max} - 1} \quad (4)$$

The inner and outer radii are taken to be $r_1 = 1$ and $r_2 = 3$, respectively. The exponent p controls the clustering near the wall, with larger values providing tighter spacing. A value of $p = 1$ was found to be sufficient for the inviscid case, while $p = 3$ produced grids suitable for the viscous case based on typical best practices. Images of the coarsest (97×49) and finest (257×129) grids for the inviscid case are shown below in Fig. 1.

The shock fitting algorithm employed here is heuristic in nature, and begins with a converged (or partially converged) flow solution. A search is performed along each wall-normal line from the outer boundary inward to find the shock. For this application, the shock is identified as the radial station where the static pressure first rises at least 5% above that of the free stream. The azimuthal array of shock locations are stored and then smoothed using Gauss-Seidel sweeps. Next, the grid is redistributed for each azimuthal station so that a specified number of points are located between the shock and the surface using a bicubic distribution. Outward from the shock, the remaining points are distributed using a constant stretching ratio while maintaining the same boundary location. The total number of points remains unchanged through the fitting process.

¹`avro` is available via gitlab.com/philipclaude/avro under GNU LGPL, Version 2.1.

²`refine` is available via github.com/NASA/refine under the Apache License, Version 2.0.

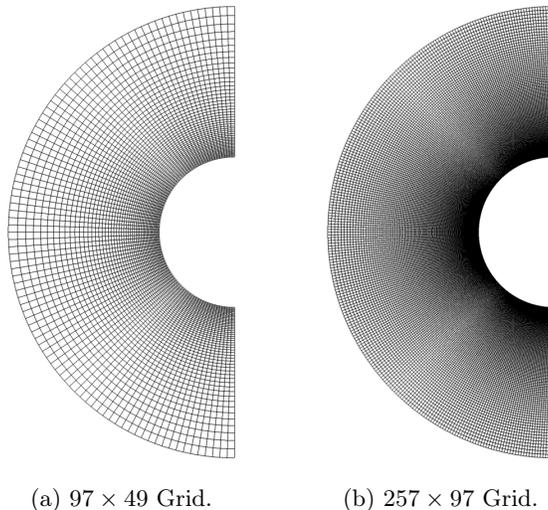


Figure 1: Algebraically generated structured meshes.

All shock-fit **OVERFLOW** solutions presented here began on a uniform grid and were subject to the fitting procedure three times, each after fully converging the solution on the prior grid. It was not found to be necessary to interpolate the solution onto the new grid before restarting.

5 Results

5.1 Inviscid, “Bova cylinder”

A full set of grid refinement studies were performed for the $M_\infty = 20$ “Bova cylinder” [1] case using **OVERFLOW** with the algebraic and shock-fit grids, and **SANS** with adapted grids obtained using **Pointwise**, **avro**, and **refine**. Density contours from the final converged Euler solutions with the finest shock-fit and adapted meshes are shown in Fig. 2. The initial and final meshes for each mesh generator are shown in Figs. 3 and 4 respectively. The **OVERFLOW** initial algebraic mesh is much finer since the shock fitting procedure employed here does not alter the grid dimensions, whereas the output-based adaptation successively refines the meshes based on a target limit to the number of DOFs.

The final meshes strongly reflect the underlying physics used in adaption as well as constraints of the mesh-generation tool. For instance, the structured, shock-fit mesh is feature-based, and it responds strongly to the shock all the way to the outflow boundaries at the top and bottom of the cylinder and produces a single grid line aligned with the feature. As a consequence, the **OVERFLOW** solution produced a very crisp shockwave as was seen in Fig. 2. The **avro** and **refine** meshes respond to the shock feature all the way to the boundary, and owing to their anisotropic capabilities, are able to produce a much more crisp shock than **Pointwise**, which is constrained to be fully isotropic and have more limited gradation between elements than what is permissible by **avro** and **refine**. For all three output-based adapted meshes, there is a prominent reduction in element density in triangular regions between the surface and the shock at the top and bottom of the cylinder. This is a direct result of the adaptation functional being a surface-integrated quantity, as this region of the grid is outside of the domain of dependence for the surface. It is downstream of the sonic line (which itself is refined), so the solution is hyperbolic in nature in this region. An additional feature of note in the output-based adaption is the response to the stagnation streamline. All three of these solutions placed an emphasis on this region, but this is unsurprising as it both experiences the strongest jump conditions across the shock wave and corresponds to the region of the surface whose normal is completely aligned with the drag direction.

The integrated drag coefficient as it varies across mesh refinement level for all four adaption strategies is shown in Fig. 5. The two sets of **OVERFLOW** solutions (non-fitted and shock-fit) exhibit monotonic convergence of the drag. The shock-fitting procedure does appear to reduce the functional error in that the slope is reduced

on these axes, implying a reduced error coefficient. The `avro` and `refine` results show a marked improvement in solution quality with adaptation. As the DOF count became more similar to that used with `OVERFLOW`, the curves rapidly flatten towards grid independence. The `Pointwise` results, on the other hand, deviate substantially until the grid is much finer than was used with the other methods. The slower convergence is attributed to `Pointwise` being constrained to isotropic meshes for this case, while the structured, `avro`, and `refine` grids can all support arbitrarily strong levels of anisotropy in the mesh.

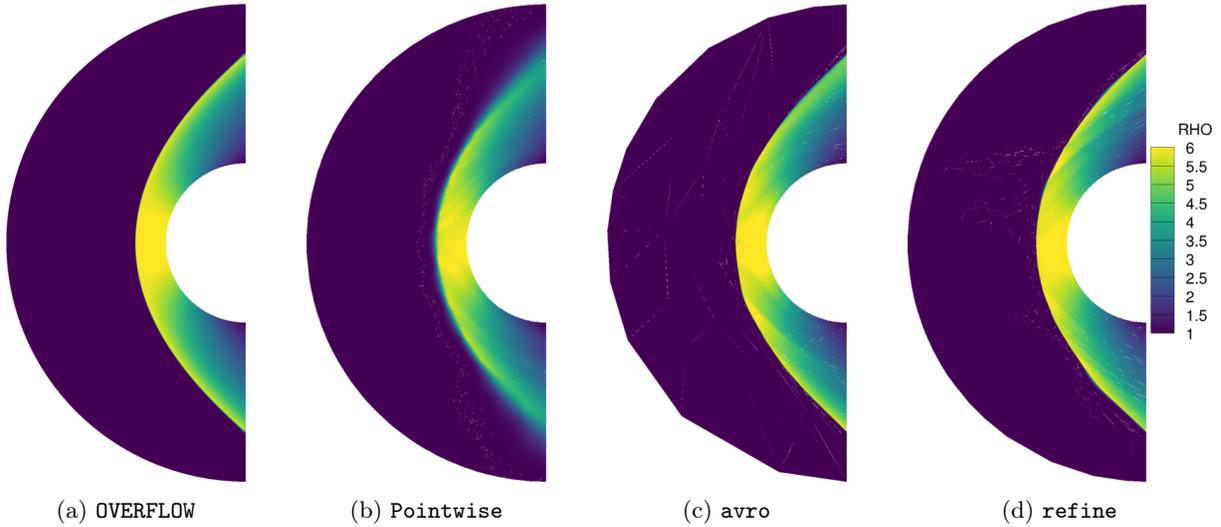


Figure 2: Density contours on final adapted meshes for the Bova cylinder with $M_\infty = 20$, Inviscid

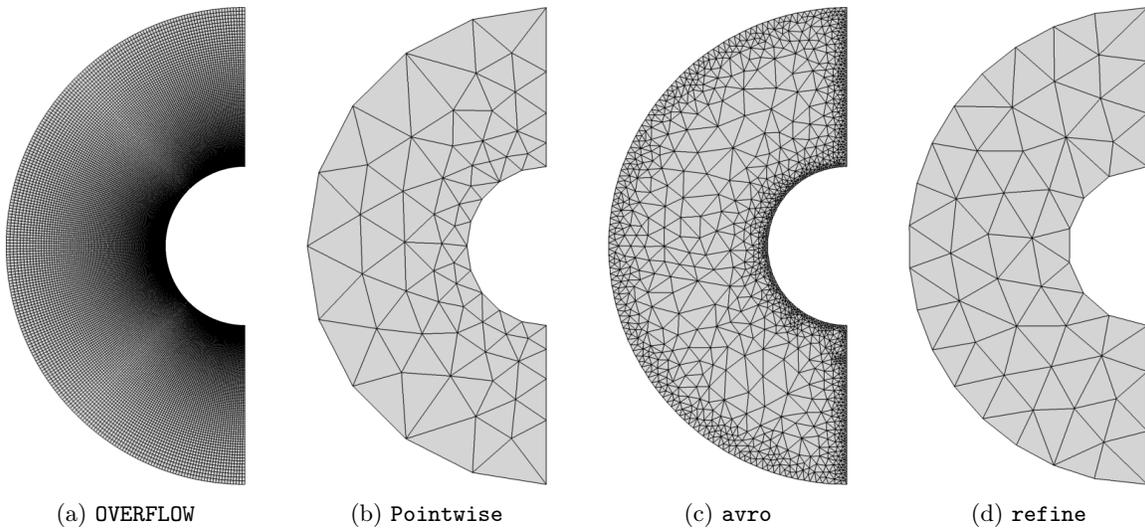


Figure 3: Initial meshes for the inviscid Bova cylinder

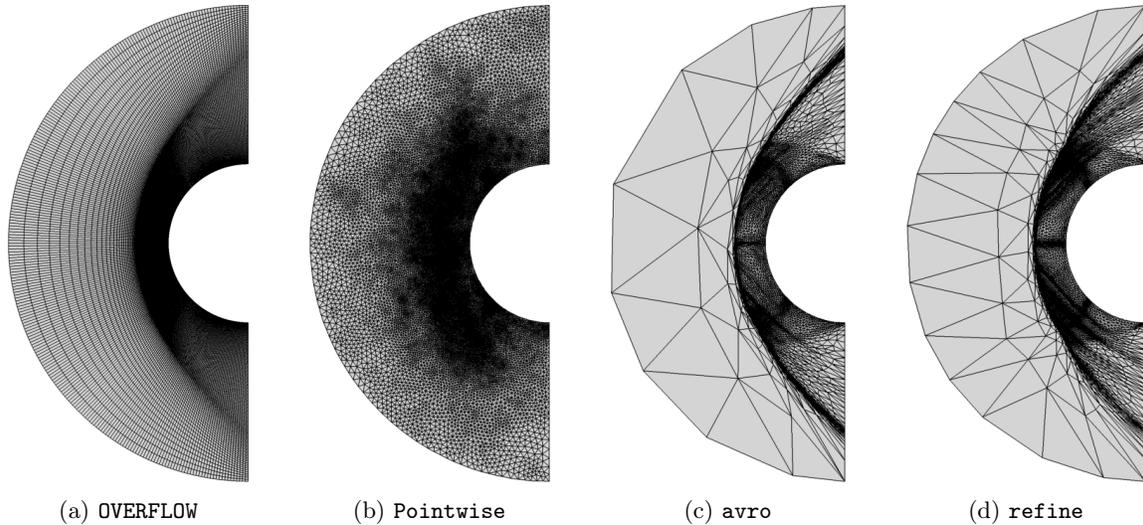


Figure 4: Final adapted meshes with comparable DOF counts for the inviscid Bova cylinder

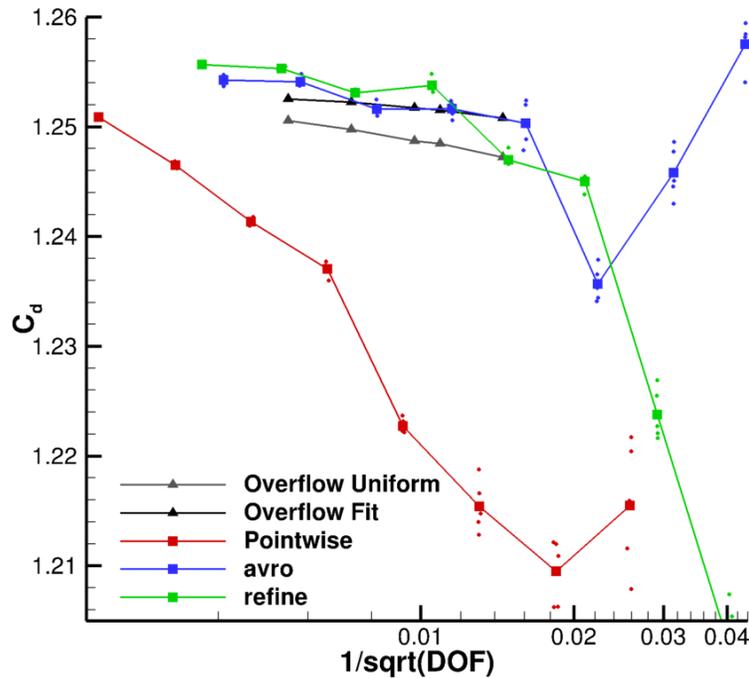


Figure 5: Grid convergence with mesh adaptation of the inviscid Bova cylinder integrated drag coefficient

5.2 Viscous, “Gnoffo cylinder”

The viscous “Gnoffo cylinder” case provided an increase in problem complexity by requiring the resolution of a boundary layer in addition to the shock structure. The final adapted meshes obtained using **Pointwise** and **avro** are shown in Fig. 6 with a close-up of the cylinder surface in Fig. 7 to illustrate the clustering in the boundary layer. The **avro** results for this case show qualitatively very similar behavior away from surface as it did for the inviscid case. The shock structure is captured very similarly, with the same triangular regions outside of the domain of dependence being much coarser. The zoomed-in image shows an arc of mesh clustering around $\phi 3\pi/4$, which corresponds to the sonic line behind the bow shock. The **Pointwise** adapted grid uses *a priori* boundary layer insertion to capture the boundary layer, but does not capture a

distinct refinement region around the shockwave as it did for the inviscid case. This is attributed to the use of a predefined boundary-layer insertion method that consumed much of the target DOF count. The grid refinement results for drag coefficient with `Pointwise` and `avro` are plotted in Fig. 8. The behavior is again qualitatively similar to what was observed with the inviscid case, with the `Pointwise` meshes requiring much higher DOF counts for the same accuracy.

Simulations performed with `OVERFLOW` for this case did not exhibit strong iterative convergence. While the flow fields for each grid resolution set up to be qualitatively reasonable, the L^2 -norm of the right-hand side showed approximately 0.5-orders of residual drop relative to the initialization from a much lower Mach number. By comparison, the inviscid solutions showed approximately 13 orders of magnitude residual drop. Consequently, the `OVERFLOW` solutions were not deemed to be suitable for comparison for the viscous case.

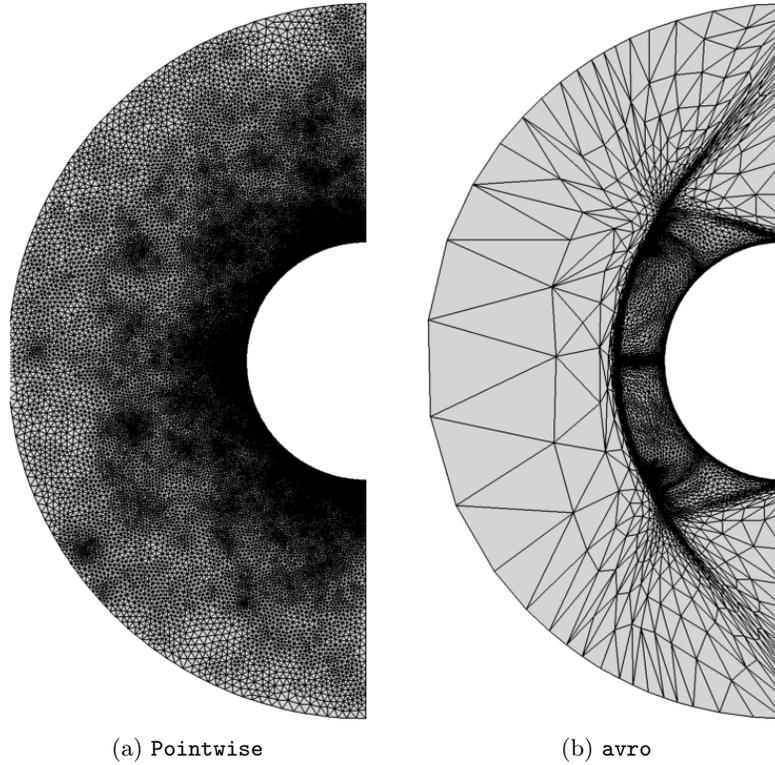


Figure 6: Final “Gnoffo cylinder” adapted meshes with comparable target DOF counts

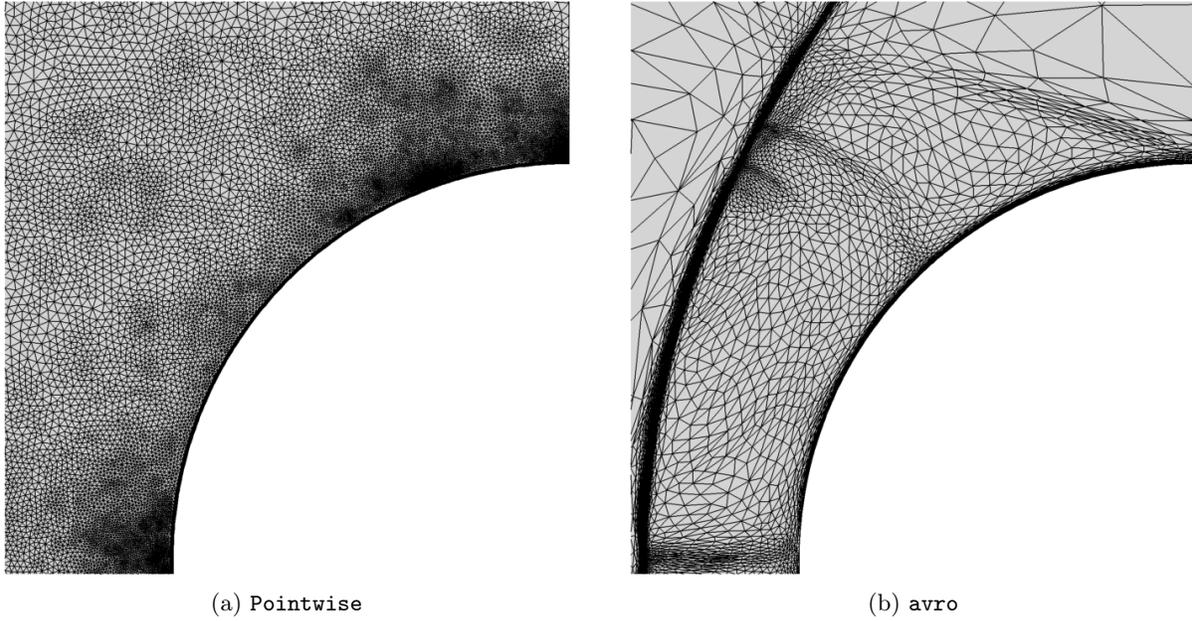


Figure 7: Closeup of “Gnoffo cylinder” final adapted meshes

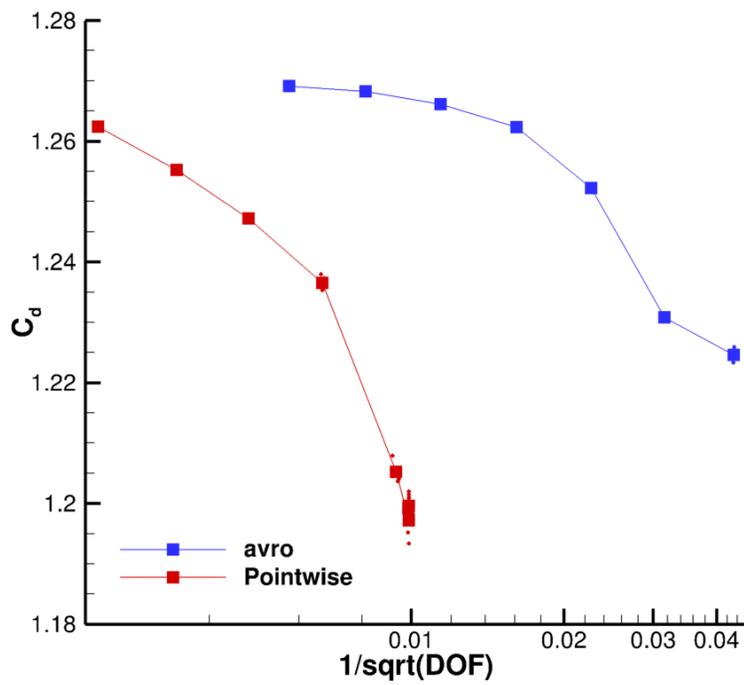


Figure 8: Grid convergence with mesh adaptation of the “Gnoffo cylinder” integrated drag coefficient

6 Conclusion and Future Work

The comparison between feature-based mesh adaptation (heuristic shock fitting) and output-based mesh adaptation for an inviscid cylinder at hypersonic conditions revealed important insights about the nature of the flow and what is required to obtain the highest-quality solution. First, the shock-fitting approach tracks the shock feature in the entire domain, for better or for worse. Shock fitting does improve the apparent accuracy of the integrated drag coefficient and improve the resolution of the shockwave itself, but it also refines regions that have no influence on the output functional.

The output-based adaptation was observed to be much more targeted in its refinement regions. The methods that permit anisotropic cells responded to the shockwave, but also reduced resolution outside of the domain of dependence. In a stark contrast with the shock-fit grids, the drag output functional led to a strong emphasis being placed on the stagnation streamline, and to a lesser extent, the sonic line. Although very influential to the flow field (as confirmed by the drag adjoint), they lack distinct features that would trigger a feature-based adaptation method.

Although deeply converged viscous solutions for the “Gnoffo cylinder” are unavailable from **OVERFLOW** at this time, the partial results from **SANS** and **avro** show that output-based mesh adaptation responds just as well to the boundary layer as it does to the shock wave at hypersonic conditions. Future work will focus on obtaining converged **OVERFLOW** solution to provide the basis of comparison against more-industrial approaches.

7 Acknowledgments

This work was supported by Cadence Design Systems through contracts with the University of Tennessee Knoxville and Massachusetts Institute of Technology.

References

- [1] Steven Bova, Ryan Bond, and Benjamin Kirk. “Stabilized Finite Element Scheme for High Speed Flows with Chemical Non-Equilibrium”. In: *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, 2010. DOI: doi:10.2514/6.2010-1560.
- [2] Peter A. Gnoffo. *Computational Fluid Dynamics Technology for Hypersonic Applications*. AIAA 2003-2829. 2003.
- [3] Arthur C. Huang et al. *A Variational Multiscale Method with Discontinuous Subscales for Output-Based Adaptation of Aerodynamic Flows*. AIAA 2020-1563. Jan. 2020. DOI: 10.2514/6.2020-1563.
- [4] Arthur C. Huang. “An Adaptive Variational Multiscale Method with Discontinuous Subscales for Aerodynamic Flows”. PhD thesis. Department of Aeronautics and Astronautics: Massachusetts Institute of Technology, Feb. 2020.
- [5] Marshall C. Galbraith, Steven R. Allmaras, and David L. Darmofal. *A verification driven process for rapid development of CFD software*. AIAA 2015-0818. Jan. 2015. DOI: 10.2514/6.2015-0818.
- [6] Satish Balay et al. “Efficient Management of Parallelism in Object Oriented Numerical Software Libraries”. In: *Modern Software Tools in Scientific Computing*. Ed. by E. Arge, A. M. Bruaset, and H. P. Langtangen. Birkhauser Press, 1997, pp. 163–202. DOI: https://doi.org/10.1007/978-1-4612-1986-6_8.
- [7] Satish Balay et al. *PETSc/TAO Users Manual*. Tech. rep. ANL-21/39 - Revision 3.16. Argonne National Laboratory, 2021.
- [8] Satish Balay et al. *PETSc Web page*. <https://petsc.org/>. 2021.
- [9] Robert H. Nichols and Pieter G. Buning. “User’s Manual for OVERFLOW 2.3, Version 2.3d”. In: *NASA Langley Research Center* (February 2021).

- [10] Robert W. Tramel, Robert H. Nichols, and Pieter G. Buning. “Addition of Improved Shock-Capturing Schemes to OVERFLOW 2.1”. In: *10th AIAA Computational Fluid Dynamics Conference*. San Antonio, TX: AIAA Paper 2009-3988, June 2009. DOI: 10.2514/6.2009-3988.
- [11] Masayuki Yano. “An Optimization Framework for Adaptive Higher-Order Discretizations of Partial Differential Equations on Anisotropic Simplex Meshes”. PhD thesis. Department of Aeronautics and Astronautics: Massachusetts Institute of Technology, June 2012.
- [12] Hugh Alexander Carson. “Provably Convergent Anisotropic Output-Based Adaptation for Continuous Finite Element Discretizations”. PhD thesis. Department of Aeronautics and Astronautics: Massachusetts Institute of Technology, Feb. 2019.
- [13] Hugh A. Carson et al. *Mesh optimization via error sampling and synthesis: An update*. AIAA 2020-0087. Jan. 2020. DOI: 10.2514/6.2020-0087.
- [14] Hugh A. Carson et al. “Anisotropic mesh adaptation for continuous finite element discretization through mesh optimization via error sampling and synthesis”. In: *Journal of Computational Physics* 420 (2020), p. 109620. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2020.109620.
- [15] R. Becker and R. Rannacher. “An Optimal Control Approach to A Posteriori Error Estimation in Finite Element Methods”. In: *Acta Numerica*. Ed. by A. Iserles. Vol. 10. Cambridge University Press, 2001, pp. 1–102. DOI: 10.1017/S0962492901000010.
- [16] Adrien Loseille and Frédéric Alauzet. “Optimal 3D highly anisotropic mesh adaptation based on the continuous mesh framework”. In: *Proceedings of the 18th International Meshing Roundtable*. Springer Berlin Heidelberg, 2009, pp. 575–594. DOI: 10.1007/978-3-642-04319-2_33.
- [17] Adrien Loseille and Frédéric Alauzet. *Continuous mesh model and well-posed continuous interpolation error estimation*. INRIA RR-6846. 2009, p. 54.
- [18] Adrien Loseille, Alain Dervieux, and Frédéric Alauzet. *On 3-d goal-oriented anisotropic mesh adaptation applied to inviscid flows in aeronautics*. AIAA 2010-1067. 2010. DOI: 10.2514/6.2010-1067.
- [19] Philip Claude Caplan. “Four-dimensional Anisotropic Mesh Adaptation for Spacetime Numerical Simulations”. PhD thesis. Department of Aeronautics and Astronautics: Massachusetts Institute of Technology, June 2019.
- [20] Carmen-Ioana Ursachi et al. “Output-Based Adaptive Reynolds-Averaged Navier–Stokes Higher-Order Finite Element Solutions on a Multielement Airfoil”. In: *AIAA Journal* 59.7 (2021), pp. 2532–2545. DOI: 10.2514/1.J059968.
- [21] M.S. Jefferies J.P. Steinbrenner N.J. Wyman and S.L. Karman. *Implementation of a Size Field Based Isotropic Hex Core Mesh Algorithm*. AIAA Paper 2020-1408. 2020.
- [22] Thorsten Hansen Nicholas J. Wyman Paul Galpin and Georg Scheuere. *Robust, Efficient and Accurate Mesh Adaptation for Turbomachinery CFD Simulation*. AIAA Paper 2020-3688. 2020.
- [23] Philip Claude Caplan, Robert Haimès, and David L. Darmofal. *Extension of Local Cavity Operators to 3d + t Spacetime Mesh Adaptation*. AIAA Paper 2019-1992. 2019.
- [24] Thierry Coupez. “Génération de Maillage et Adaptation de Maillage par Optimisation Locale”. In: *Revue européenne des éléments finis* 9 (4 2000), pp. 403–423. DOI: 10.1080/12506559.2000.10511454.
- [25] Adrien Loseille, Frédéric Alauzet, and Victorien Menier. “Unique Cavity-Based Operator and Hierarchical Domain Partitioning for Fast Parallel Generation of Anisotropic Meshes”. In: *Computer-Aided Design* 85 (Apr. 2017). 24th International Meshing Roundtable Special Issue: Advances in Mesh Generation, pp. 53–67. DOI: 10.1016/j.cad.2016.09.008.
- [26] Robert Haimès and Mark Drela. *On The Construction of Aircraft Conceptual Geometry for High-Fidelity Analysis and Design*. AIAA Paper 2012-683. 2012.
- [27] Todd Michal and Joshua Krakos. *Anisotropic Mesh Adaptation Through Edge Primitive Operations*. AIAA Paper 2012-159. 2012.
- [28] Frédéric Alauzet. “A Changing-Topology Moving Mesh Technique for Large Displacements”. In: *Engineering with Computers* 30.2 (Apr. 2014), pp. 175–200. DOI: 10.1007/s00366-013-0340-z.

- [29] Lori A. Freitag and Carl Ollivier-Gooch. “Tetrahedral Mesh Improvement Using Swapping and Smoothing”. In: *International Journal for Numerical Methods in Engineering* 40 (1997), pp. 3979–4002. DOI: 10.1002/(SICI)1097-0207(19971115)40:21<3979::AID-NME251>3.0.CO;2-9.
- [30] Robert Haimes and John F. Dannenhoffer III. *EGADSlite: A Lightweight Geometry Kernel for HPC*. AIAA Paper 2018–1401. 2018.