A Survey of Solver-Related Geometry and Meshing Issues

James Masters[1], Derick Daniel[2]
Arnold Engineering Development Complex, Arnold Air Force Base, TN 37389

Jared Gudenkauf[3]
Jacobs ESSSA Group, Huntsville, AL 35808

David Hine[4]
NAVAIR, Patuxent River, MD 20670

Chris Sideroff[5]
Applied CCM, Ottawa, Canada

There is a concern in the computational fluid dynamics community that mesh generation is a significant bottleneck in the CFD workflow. This is one of several papers that will help set the stage for a moderated panel discussion addressing this issue. Although certain general "rules of thumb" and a priori mesh metrics can be used to ensure that some base level of mesh quality is achieved, inadequate consideration is often given to the type of solver or particular flow regime on which the mesh will be utilized. This paper explores how an analyst may want to think differently about a mesh based on considerations such as if a flow is compressible vs. incompressible or hypersonic vs. subsonic or if the solver is node-centered vs. cell-centered. This paper is a high-level investigation intended to provide general insight into how considering the nature of the solver or flow when performing mesh generation has the potential to increase the accuracy and/or robustness of the solution and drive the mesh generation process to a state where it is no longer a hindrance to the analysis process.

Nomenclature

A	=	Area
A^*	=	Nozzle throat area
C_L	=	Lift coefficient
C_D	=	Drag coefficient
C_M	=	Moment coefficient
k	=	Turbulent kinetic energy
P	=	Pressure
R	=	Universal gas constant
t	=	Time
T	=	Temperature
β	=	Crank-Nicolson coefficient
ε	=	Turbulence dissipation
δ	=	Discrete equivalent of the gradient or divergence operator
ρ	=	Density
μ	=	Viscosity
ψ	=	Compressibility (1/RT)

[1] Engineer/Scientist, Integrated Test and Evaluation Branch, Senior Member AIAA
[2] Engineer/Scientist, Integrated Test and Evaluation Branch, Member AIAA
[3] Computational Fluid Dynamicist, Fluid Dynamics Branch, Member AIAA
[4] Aerospace Engineer, Applied Aerodynamics and Store Separation Branch, AIAA Member
[5] Principal Director, Applied CCM, Member AIAA

I. Introduction

According to NASA's CFD Vision 2030 Study, mesh generation and adaptivity continue to be significant bottlenecks in the computational fluid dynamics (CFD) workflow and, as more capable HPC hardware enables higher resolution simulations, fast reliable mesh generation and adaptivity will become more problematic[1]. Based on this assessment, and acknowledging the fact that the computational mesh, along with the geometry that it represents, has a considerable impact on the quality, stability, and amount of resources required to complete numerical simulations,[2] a meshing 2030 subcommittee was formed as part of the AIAA Meshing Visualization and Computational Environments (MVCE) Technical Committee. The subcommittee is moving forward by planning several activities in the upcoming years leading to a series of Geometry and Mesh Generation Workshops, the first of which (Mesh-W1) will coincide with AIAA's 2017 Aviation conference and be held in conjunction with the 3rd AIAA High Lift Prediction Workshop. The workshop will focus on examining the current state of geometry and meshing technologies and making progress toward the goals outlined in the NASA 2030 study. Recognizing the importance of making the workshop relevant to the greater aerospace community, it is the intent of the subcommittee to continue partnering with other technical workshops as is being done with the High Lift.

The desire to optimize the relevance and usefulness of the meshing workshop has led to the planning of several lead-up events including a conference session discussing open issues in meshing where several papers will be presented followed by a moderated panel discussion. the work presented herein will explore a range of solver or flow regime related geometry and meshing issues and will be presented as part of an over-arching theme with other subtopics such as the community's reaction to the mesh assessment laid out in the CFD Vision 2030 study and an overview of the NASA Common Research Model (CRM) from a geometry handling perspective. The papers are meant to set the stage for a moderated panel discussion, the objective of which is to provide a technical interchange between simulation practitioners, researchers, and software developers. It is hoped that the discussion and the sharing of ideas will help to illuminate the path for geometry and mesh generation research and tool development to meet the 2030 CFD vision goal of fully automated mesh generation as an invisible component of a simulation based design program.

It is easy to discern when a mesh is "bad" (negative volumes, negative or very high skew) but it is often much harder to know when a mesh is "good." When generating a computational mesh for a given CFD analysis, the nature of the flow, as well as the geometry, necessarily determines the nature of the mesh. Experience working with computational meshes and the aerodynamic solutions generated on them has led to certain generic "rules of thumb" such as trying to advance mesh elements orthogonally away from viscous surfaces or refining meshes in regions where strong gradients (both in the solution and in the geometry) are known to exist. However, an inadequate degree of consideration is often given to the particular solver or range of flow regimes on which the mesh will eventually be utilized. Insight, which can be gained from considering the solver and flow regimes prior to generating meshes, has the potential to increase the efficiency of the mesh generation process and make it increasingly invisible to the end-user.

Several situations are explored herein related to increasing accuracy, robustness and/or efficiency of the CFD analysis process based on insight into the mesh/solver interactions. These include looking at how geometric curvature gradients can be efficiently resolved without sacrificing solution quality as well as how sharp corners in a geometry can potentially affect the mesh and subsequently the solver. Also examined is how the effectiveness of meshing the near-body region, especially in the viscous boundary layer, can be improved based on insight into the solver (e.g. is the solver node-centered or cell-centered). Finally, areas such as what impact does a mesh have on a solution if a flow is compressible vs. incompressible or hypersonic vs. supersonic, transonic or subsonic will be examined and certain solver specific mesh checking techniques will be explored.

II. Solver Specific Mesh Checking

A. Background

When creating an unstructured mesh with most mesh generators, the user needs to be aware that while the mesh generator may not have produced errors, geometric issues may still arise that specific flow solvers might not be able to handle. Determining if a flow solver can handle a mesh can require experience with the specific solver. Since there is no set standard for a priority mesh quality metrics across mesh generation tools, the creators of the solver Loci/Chem have developed a mesh checker that contains some good criteria for determining the overall mesh quality; this tool is called "vogcheck". Three mesh generators (Pointwise, VGRID, and AFLR3D) will be used to

generate unstructured meshes on an axisymmetric nozzle geometry and the meshes will be checked with "vogcheck" and run with the CFD flow solver Loci/Chem.

The "vogcheck" code is a utility built by the developers of the CFD code Loci/Chem which does four separate grid quality checks and then gives an overall mesh quality assessment. Loci/Chem can only read in a “vog” file, but it comes with converters for several different mesh types, thus the name vogcheck. The four quality checks are 1) Convexity, 2) Volume ratio between neighboring cells, 3) Angle between face normal and cell centroid, and 4) Face twist of non-triangular faces. The overall mesh quality using these four checks is then reported as either Excellent, Good, Poor, Marginal, or Unusable.

The first quality check of convexity is either a pass or a fail and if this check fails the mesh is unusable; no modern mesh generation software should produce a mesh that fails this check. The definition of a convex polygon is one that has all interior angles less than or equal to 180^0 and therefore the interior normal of all the edges points toward the center. Examples of convex (Fig. 1) and non-convex (Fig. 2) polygons are illustrated below.

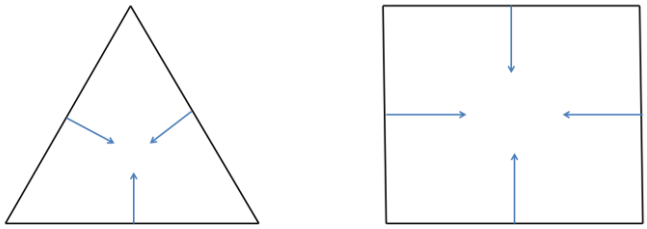

Fig. 1 Example of a Convex 3 and 4 Sided Polygon

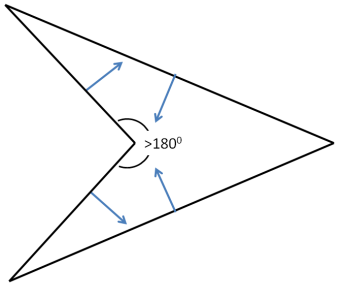

Fig. 2 Example of Non-Convex 4 Sided Polygon

Because a triangle can only be defined as convex, all positive volume cells that are created from triangles are convex and these cells would therefore have no problem passing the convexity criterion. One issue that does exist with building a grid that contains prismatic layers is that usually, due to geometric constraints, the number of prismatic layers generated off each surface triangle is variable and when this occurs a pyramid is required to transition the quadrilateral face of the prism to a tetrahedral cell; this is the location where a non-convex cell is usually generated. Should this prism-to-pyramid transition occur at a hip, or saddle point, in the geometry a non-convex pyramid might be placed that has a base that is similar to the polygon shown in Fig. 2.

The second quality check that occurs in "vogcheck" is called "Max Cell Volume Ratio," which is the volume mesh ratio between the two cells sharing a common face; where the ideal value of no greater than 10 is recommended and a maximum value of larger than 100 will affect the overall grade of the mesh. Large jumps in volume ratio usually occur in viscous layers of a poorly built mesh (i.e., not using enough layers to smoothly transition into the inviscid region of the mesh) or when a grid contains very small, high aspect-ratio cells ("sliver cells") in the volume; sliver cells are present in the VGRID mesh built for this paper. Another location where large volume ratio cells occur is when building the prismatic boundary layers and a pyramid is required for the transition. The reason that large volume ratios affect solution quality is that in the discretization equation[3] for a cell, the distance to the neighboring cell is in the denominator and when this volume ratio is large the flux from that cell decreases resulting in the slowing down of information from that cell[4].

The third quality check measure is called "Max Cell Face Angle" and is defined as the angle of the face normal to the cell centroid; the ideal value is less than 100^0 and a poor cell has an angle greater than 150^0. To illustrate this angle, Fig. 3shows two neighboring triangles, 1-2-3 and 2-4-3. The centroid of each triangle is marked with the red

circle (point 5 and 6) and the middle of the line segment between the two triangles with a blue diamond (point 7). The angle that vogcheck is measuring are the two angles defined as 6-5-7 and 5-6-7, respectively; only the maximum of the two is reported for each face. The angles that vogcheck would have measured for these two triangles is 52^0 and 8^0, respectively.

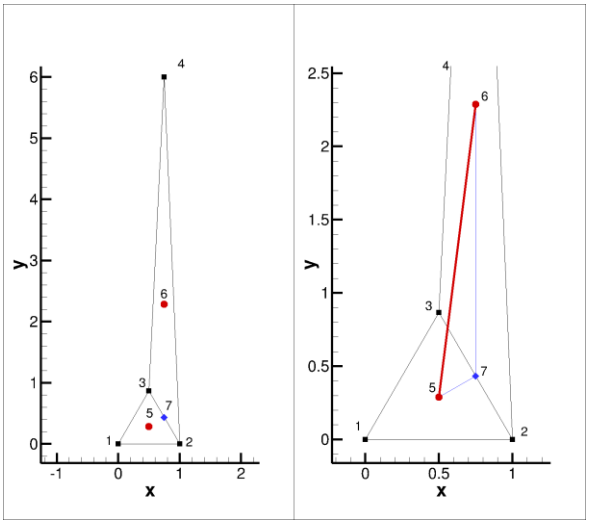

Fig. 3 Triangles illustrating a large "max cell face angle"

The final check that is performed measures the "Max Twist" and "Max Shear Twist" and pertains to cells containing 5 or more sides (e.g., prisms, pyramids, or hexahedral cells) where a twist value of less than 10% is desirable. Some face twist is unavoidable when building viscous layers on complex geometries but in the geometry presented in this section very little face twist was generated.

A. Mesh and Case Setup

A nozzle geometry will be tested using different CFD codes and different mesh software. The nozzle geometry contains a constant area plenum leading to a throat with an A/A* = 2 yielding a Mach Number of approximately 2.2; The exit of the nozzle was extended about two nozzle lengths to form a constant area section to highlight boundary layer growth and shock formation down the nozzle. The complete nozzle is illustrated in Fig. 4.

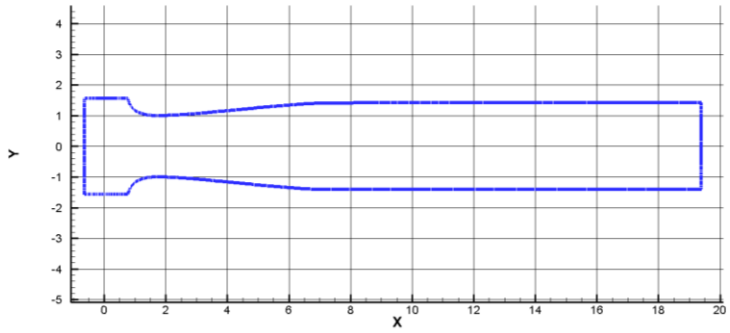

Fig. 4 Complete nozzle geometry

Grid solvers that will be highlighted in this section are Pointwise, AFLR3D, and VGRID. A total of 4 different surface grids were created (3 in Pointwise and 1 in VGRID). The combination of surface generator, boundary layer type, and volume generator used to build each of the 7 grids is displayed in Table 1.

Table 1: Grid Generation Method for Nozzle Grids

Grid	Surface	Surface Generator	Boundary Layer	Volume Generator
Volume Grid 1	Surface Grid 1 (hexes)	Pointwise	Hexahedron	Pointwise
Volume Grid 2	Surface Grid 2 (triangles)	Pointwise	Tetrahedron	Pointwise
Volume Grid 3	Surface Grid 2 (triangles)	Pointwise	Prisms	Pointwise
Volume Grid 4	Surface Grid 3 (triangles)	Pointwise	Tetrahedron	Pointwise
Volume Grid 5	Surface Grid 3 (triangles)	Pointwise	Tetrahedron	AFLR3D
Volume Grid 6	Surface Grid 3 (triangles)	Pointwise	Prisms	AFLR3D
Volume Grid 7	Surface Grid 4 (triangles)	VGRID	Tetrahedron	VGRID

1. Surface Meshes Generated with Pointwise

The first surface grid that was created was the structured grid illustrated in Fig. 5. The grid spacing was varied over the geometry to cluster cells in areas where high gradient flow is expected (e.g., the plenum, beginning of the bell mouth, and the throat); The number of points at the different radial stations is constant with 101 points.

The second surface grid was constructed using the exact same surface spacing that Pointwise used when making Surface Grid 1 and simply switching the solver to unstructured mode. As Fig. 6 shows, simply switching the solver from structured to unstructured using the same spacing produced an unstructured grid that is not as uniform. The unstructured surface solver blended the finer radial spacing into the coarser lengthwise spacing quite rapidly in the diverging section of the nozzle (i.e., after the throat); the decay value used in Pointwise was set to 0.99 (least decay) but the mesh still coarsened beyond what was desired.

As will be demonstrated later on, Surface Grid 2 is usable. However, the transition from the small radial cells to the larger lengthwise cells can be improved by either increasing the radial spacing or reducing the lengthwise spacing. This is what was done for Surface Grid 3, where the radial spacing was increased to match that of the lengthwise spacing and the spacing in the constant area section was halved.

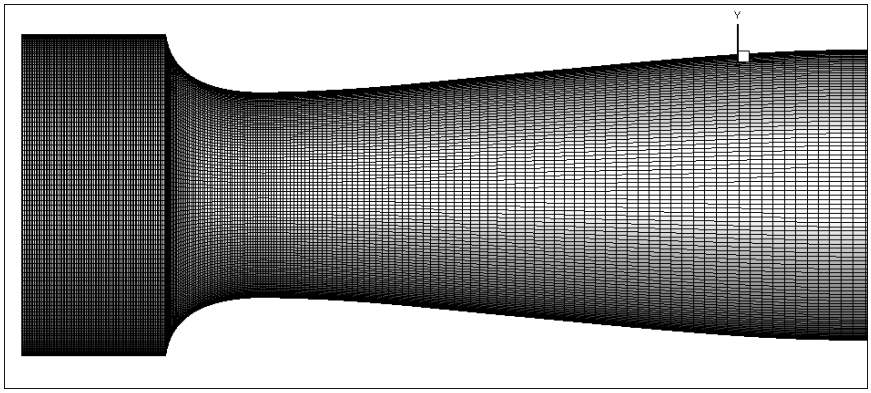

Fig. 5 Pointwise Surface Grid 1

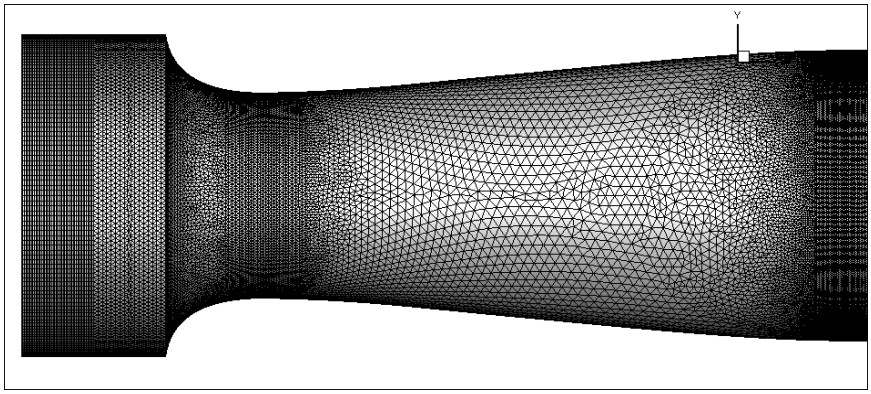

Fig. 6 Pointwise Surface Grid 2

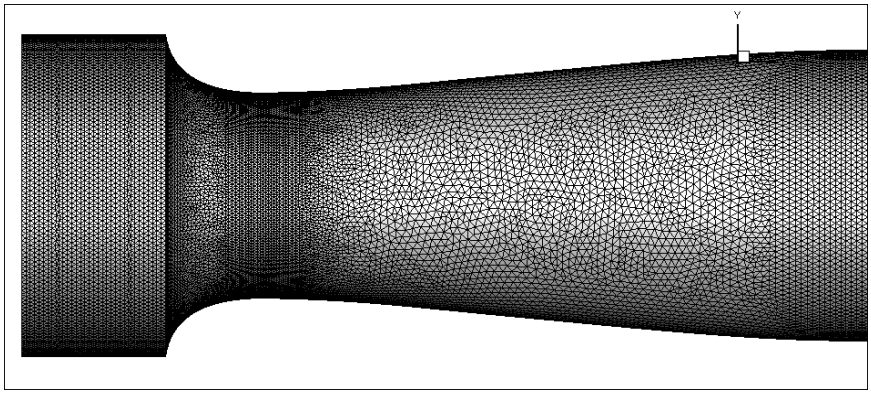

Fig. 7 Pointwise Surface Grid 3

2. Surface mesh using VGRID

Unlike Pointwise which is an all-in-one tool in which the user builds the geometry then the surface and volume grids, VGRID requires the use of another software package to generate the input file. For the nozzle geometry, the IGES surfaces created in Pointwise were imported into GRIDTOOL and the model was built for running with VGRID. VGRID uses a source model where sources placed around the geometry are used to solve for a distribution of points on the geometry. For this particular geometry, the type of sources that were used are volume sources which VGRID uses to apply the specified spacing inside the volume source. To build the surface grid, the spacing parameters used for Pointwise Surface Grid 3 were used resulting in the surface grid shown in Fig. 8.

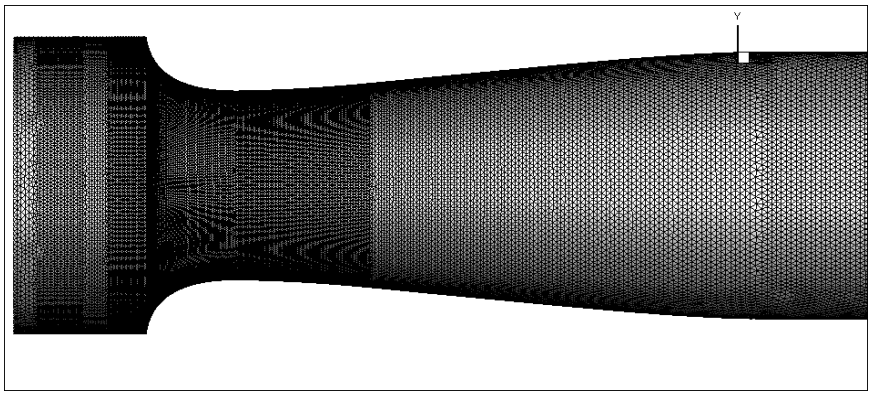

Fig. 8 Surface Grid for VGRID Mesh (Surface Grid 4)

3. Volume Meshes Generated with Pointwise

Each of the 7 volume meshes discussed in this section were built such that they are as similar as each grid solver would allow. When creating the viscous layers, all of the solvers required both the initial spacing and growth rate and these were set at 0.00001 and 1.2, respectively. Each of the solvers also has the feature where the viscous layers were built until the last layer approached the size of the inviscid volume cells at which point the layer was stopped. This leads to different numbers of viscous layers throughout the Mesh. For Volume Grids 1 through 3, the viscous layers were built as hexahedron and then transitioned to a tetrahedron inviscid volume. Fig. 9 shows the volume grids at a slice at X=17".

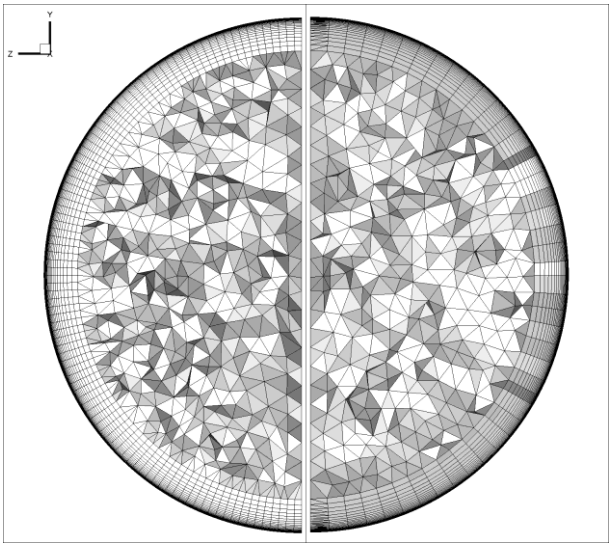

Fig. 9 Crinkle Cut at X = 17" for Volume Grid 1 (left) and for Volume Grid 2 and 3 (right)

Because there is a structured (hex) surface between the viscous layers and tetrahedron volume for Volume Grid 1, a layer of pyramids was required. This appears to have generated some irregularly shaped cells resulting in the Marginal rating in the "vogcheck" statistics shown in Table 2. This rating is based on the fact that the maximum volume ratio and maximum cell face angle are both higher than the suggested values. It is not surprising that because Volume Grid 2 and 3 are basically the same volume grid that their overall rating is Good. However, notice that the rest of the "vogcheck" statistics are different. Since Grid 2 was not converted to prisms, the cell count is nearly double and cell volume ratio is about half from the converted grid. However, converting the grid to prisms did lower the maximum cell face angle a few degrees and the grid is half the size so it should run about twice as fast. Also notice that in Figure 38 that Pointwise could not combine the layers on the plane of symmetry.

Volume Grid 4 was built to be similar to Grid 2 except with a different surface; one that had more consistent cell spacing in the radial and lengthwise directions but was finer in the constant area exit section. The "vogcheck" rating for this grid also came out as Good and has very similar statistics compared with Volume Grid 2.

Table 2: Vogcheck Statistics For Grids 1-4

Volume Grid	1	2	3	4
Cell Count	3,381,981	7,125,984	3,731,161	11,562,717
Max Cell Volume Ratio	157.85	16.85	35.30	16.85
Min Cell Volume	9.48e-12	1.88e-12	3.99e-12	1.88e-12
Max Cell Face Angle	119.69	102.04	97.07	108.87
Max Twist	0.05	457e-13	3.90e-2	4.14e-13
Max Shear Twist	0.02	9.94e-11	2.28e-2	6.96e-11
Rating	Marginal	Good	Good	Good

4. Volume Meshes Generated with AFLR3D and VGRID

Volume Meshes were built from Pointwise Surface Grid 3 using AFLR3D. Like the test for Pointwise Grids 2 and 3, Grids 5 and 6 were built using all tetrahedron (Grid 5) and combining the layers into prism (Grid 6). A Volume mesh was also built using VGRID. A cross-section at X=17" of the VGRID Volume Grid (Volume Grid 7) is shown in Fig. 10.

As

Table 3 shows, the VGRID grid is determined to be unusable because of the maximum cell volume ratio and the maximum cell face angle being above the threshold for a grid. An examination of volume mesh reveals clusters of cells with large volume ratios scattered throughout the entire grid. Further inspection of these cell clusters shows the areas of high volume ratio coincide with the areas of large cell face angles. In Fig. 11, one of the clusters of cells

with volume ratio greater than 1000 is illustrated and what can be seen is that small cells are being inserted on the corners of large cells to create the grid.

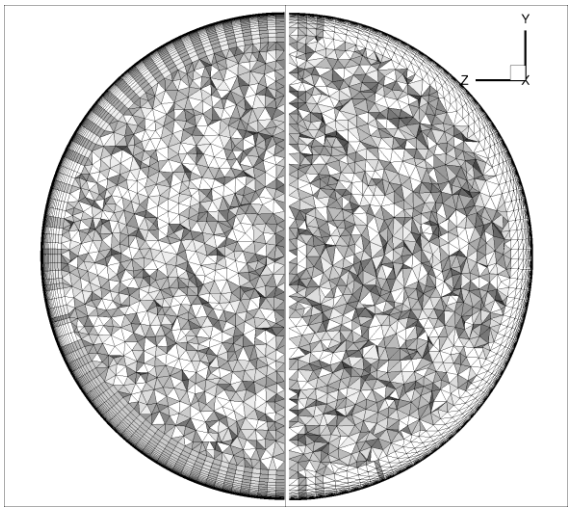

Fig. 10 Crinkle Cut at X = 17" for AFLR and VGRID Mesh (Volume Grid 5 & 7)

Upon inspection of the two volume grids illustrated in Fig. 39, above, that the it appears that either AFLR or VGRID did not grow the viscous layers properly. But this is not the case, the AFRL solver first inserts 5 identical thickness layers and then starts growing the layers at the specified rate. The effect of this option in AFLR will be explored in the future.

Table 3: Vogcheck Statistics for Grid 5-7

Volume Grid	5	6	7
Cell Count	11,799,988	5270637	10,279,351
Max Cell Volume Ratio	5.60	6.63	112,382
Min Cell Volume	1.84e-12	5.54e-12	3.75e-13
Max Cell Face Angle	76.33	61.75	175.915
Max Twist	5.27e-13	6.60e-2	3.47e-13
Max Shear Twist	1.05e-10	4.28e-2	5.17e-11
Rating	Excellent	Excellent	Unusable

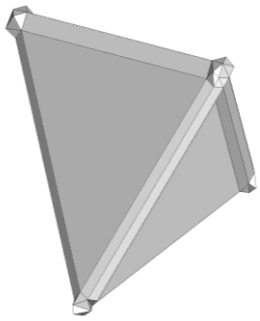

Fig. 11 Iso-Surface of cells with volume ratio >1000 using VGRID

B. CFD Solver Discussion and Results

For all cases, the inflow boundary condition was set to the same total pressure and temperature and a supersonic extrapolation boundary was set for the exit. The CFD flow solver Loci/Chem was run using the Spalart-Allmaras turbulence model and the inviscid Roe flux scheme assuming a perfect gas with a constant gamma of 1.4. The solution was initialized with a density of one atmosphere, a temperature of 298 K, and velocity of 300 m/s.

Unfortunately, some the grids could not start the nozzle flow properly from the initial condition without using an adaptive flux scheme, which switches to a more dissipative HLLE flux scheme when certain conditions are met in the flow field. Most notably was Grid #6 that required being run only with the HLLE flux scheme and then switched to Roe and Grid #7 which would only run using HLLE. Running only with HLLE has been advised against due to the dissipative nature of the scheme and therefore the results of Grid #7 will not be presented in this paper.

A common way to monitor a CFD run is to plot how the density, energy and momentum residuals are changing as a function of iteration. In Fig. 12 below, the log of the density residual for all Loci/Chem runs are plotted. Each of the Loci/Chem simulations were run for a total of 4000 iterations but using the log of the residuals as a basis of judgment for convergence, all of the runs were completely converged after only 500 iterations. Please note the difference in the residual for Grid #6 after 1000 iterations when it was switched from the HLLE to Roe flux scheme.

There is a direct correlation between the quality of a mesh (as rated by Vogcheck) and the amount by which the residual drops. Volume Grid 1, which was rated as marginal has the worst residual with the exception of volume Grid 3, which is known to be of low quality due to the large mesh size decay on the surface. As shown in Fig. 13, the solutions demonstrate similar characteristics so it appears that the primary effect that the quality of mesh has is on the robustness of the solution process. A future effort will involve running on an additional code (such as USM3D) to determine if the same behavior is experienced across multiple CFD solvers.

Although Vogcheck is a solver-specific tool, it is easy to see that its utility could be expanded to other codes as well. Also, an automated meshing process could potentially take advantage of the rating system employed by vogcheck to generate multiple types of meshes and compare them prior to running a CFD code.

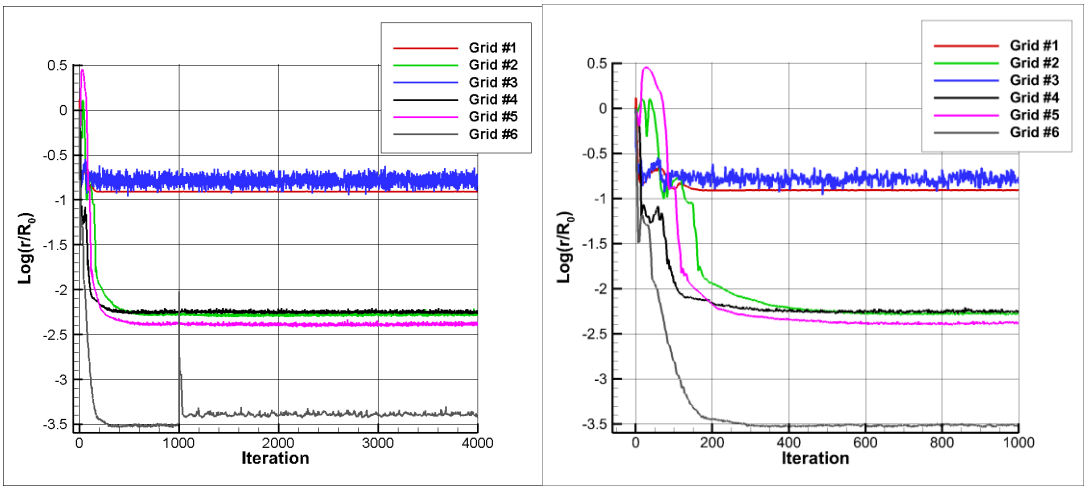

Fig. 12 Log of density residual for Loci/Chem runs

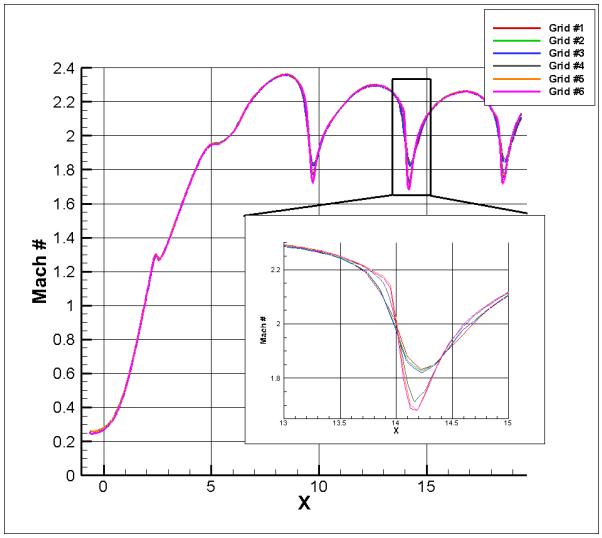

Fig. 13 Centerline Mach number for Loci/Chem simulations

III. Conclusions

The investigations described herein looked at several seemingly disparate flow analysis situations that were all, in fact, linked together by virtue of a strong dependence on a process that could be strengthened through a tighter coupling between the mesh generation process and the solver on which the mesh would eventually be utilized.

In the study looking at the effects of high-aspect-ratio cells on leading and trailing edges, it was found that although isotropic meshes are generally considered higher quality, situations exist where less isotropic meshes actually perform better. For very common geometric entities such as wings, a future meshing capability might be to recognize the entity and automatically generate an appropriate mesh.

Certain situations were also found where a mesh might be appropriate for one solver paradigm but not another. Convergence difficulties associated with node-centered solutions on grids with a polar singularity were linked to a degradation of nodal gradients and "fat cells" where the nominal "center" of the cell is not actually at the center of mass were linked to solution errors. In both of these cases, some knowledge of the solver could decrease the likelihood of generating a mesh that would cause complications down the road.

A study investigating the orthogonality of a mesh for an aerodynamic body at relatively low speeds provided insight that that gains made to the automation of the meshing process appear to, in many cases, outweigh the gains that would be experienced by subjecting the process to a hard rule that some very small level of non-orthogonality be maintained. However, it was also found that cases involving hypersonic flows were much more sensitive to near body mesh orthogonality. This sharply illustrates an example where decisions in the meshing process could be directly influenced by the aerodynamics of a particular case.

Hypersonic computations performed on several mesh variations showed that if orthogonal mesh layers do not exist at the location where a bow shock occurs, accurate heat flux prediction is difficult. However, viscous forces are less sensitive to this mesh/shock interaction. If only drag forces are of concern, meshing requirements can be less stringent than if heat flux is of concern, Additionally, incorporating compressible flow theory into the meshing process could provide an automated meshing algorithm with information that would assist in making sure that the bow shock on a hypersonic body was within an orthogonal mesh region.

The discussion outlined in this paper began with an acknowledgement that it is easy to discern when a mesh is "bad" but often much harder to know when a mesh is "good." The Loci/Chem mesh tool Vogcheck employs a philosophy that helps determine when a mesh is actually "good" not just usable. There was found a direct correlation between the quality of a mesh (as rated by Vogcheck) and the amount by which the residual of a CFD solution drops. Although Vogcheck is a solver-specific tool, it is easy to see that its utility could be expanded to other codes as well. In order to reach the goals laid out in the NASA CFD Vision 2030 study, it may be necessary for future mesh generation tools to be provided with additional information from solver developers. Mesh generation tools already have some knowledge of a given solver in order to define the appropriate mesh format so a mechanism for information exchange is already in place in many cases. Future work should include exploring ways that flow and solver specific information can be harnessed during meshing to increase the robustness, accuracy and efficiency of the over-arching analysis process.

Acknowledgments

Portions of the material presented in this paper received funding provided by the Arnold Engineering Development Complex and the CREATE™-AV Element of the Computational Research and Engineering for Acquisition Tools and Environments (CREATE™) Program sponsored by the U.S. Department of Defense HPC Modernization Program Office.

References

1. Slotnick, J., Alonso, J., et al, "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," NASA CR-2014-128178, March 2014.
2. Chawner, J., Dannenhoffer, J., et el, "The Path to and State of Geometry and Meshing in 2030: Panel Summary," 22nd AIAA Computational Fluid Dynamics Conference 22-26 June 2015, Dallas, TX, AIAA 2015-3409.
3. Patankar, S., "Numerical Heat Transfer and Fluid Flow", Hemisphere Publishing Corp, 1980
4. Ruf, J. and Lin, J., "Effect of Miss-Match of Cell Volumes in Transition From the Boundary Layer Cells to Unstructured Mesh on Computational Fluid Dynamics Results", NASA-ER-42-16-018, February 2016.
5. Ruf, J. and Lin, J., "Effect of Miss-Match of Cell Volumes in Transition From the Boundary Layer Cells to Unstructured Mesh on Computational Fluid Dynamics Results", NASA-ER-42-16-018, February 2016.