Further Extension and Validation Of A Parallel Unstructured Mesh Adaptation Package

Peter A Cavallo*
Combustion Research and Flow Technology, Inc. (CRAFT Tech)
6210 Keller's Church Road, Pipersville, PA 18947
cavallo@craft-tech.com
Phone: 215-766-1520 / Fax: 215-766-1524

and

Matthew J. Grismer†
Air Force Research Laboratory
Wright-Patterson AFB, OH 45433
matthew.grismer@wpafb.af.mil

A parallel tetrahedral mesh adaptation code is expanded to treat general, mixed-element unstructured meshes comprised of any combination of basic element types. Emphasis is placed on developing conforming mesh modification methods that are solver-independent. Specific developments include the implementation of a treatment for viscous, high aspect ratio near wall tetrahedra, and cell subdivision methods for prismatic, hexahedral, and pyramid cells. Rebalancing of the adapted grid, and particularly issues associated with processor assignment of parent/child cell sets, is addressed. Validations are performed for decomposed, mixed-element meshes using cell-vertex and cell-centered unstructured solvers. The resulting parallel adaptation package is a powerful, versatile tool for obtaining grid-converged, steady state results, and may readily be applied to other unstructured flow solvers.

I. Introduction

Three-dimensional numerical simulations on parallel computing platforms have become commonplace in recent years. A number of structured and unstructured grid solvers for fluid dynamics and finite element analysis have been extended to operate in parallel environments. Not surprisingly, parallel adaptive mesh refinement methods have also received a fair amount of attention[1-8]. For fixed boundary applications, these include treatments for embedded quadrilateral[1] and Cartesian meshes[2], *r*-refinement of block structured grids[3], cell subdivision of unstructured triangular[4] and tetrahedral meshes[5,6], octree-based tetrahedral grids[7], and unstructured hexahedral grids[8]. Previous work by the first author[9,10] demonstrated parallel mesh coarsening and refinement methods for tetrahedral meshes, with a focus on moving boundary applications with deforming cells. Parallel adaptation of a computational mesh offers several improvements over traditional, serial mesh refinement schemes. These advantages include a reduction in memory requirements for large, multi-million cell models, efficiency gains from the use of multiple processors, and a closer coupling with the flow solution process, removing the need for separate preprocessing and repartitioning steps.

Accurate resolution of viscous phenomena such as boundary layers and shear layers requires the use of high aspect ratio cells. While it is possible to compute such flows using anisotropic tetrahedral meshes, mixed-element meshes are more advantageous from the standpoint of efficiency[11]. A single hexahedron or prism can occupy the same volume as several tetrahedra with fewer edges and faces, resulting in fewer flux calculations. With the

43rd Aerospace Sciences Meeting and Exhibit, Jan. 10-13 2005, Reno, NV

* Senior Research Scientist, AIAA Senior Member.
† Senior Research Aerospace Engineer, AIAA Senior Member.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2005		**00-00-2005 to 00-00-2005**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Further Extension and Validation of A Parallel Unstructured Mesh Adaptation Package	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Combustion Research and Flow Technology Inc (CRAFT Tech),6210 Keller's Church Road,Pipersville,PA,18947**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**		**15**	

increased flexibility in geometric modeling and resolution offered by mixed-element meshes, adaptation methods for such meshes are of interest.

Mixed-element unstructured mesh adaptation has been investigated by several researchers[12-15]. Parthasarathy and Kallinderis[12] explored subdivision and redistribution methods for tet/prism grids. The prism subdivision was restricted to vertical refinements propagating from the tet/prism interface down to the wall, producing additional prisms. The first author also considered such an approach in previous work[13]. Biswas and Strawn[14] devised a conforming mesh adaptation method for unstructured hexahedral grids, in which child cells of tetrahedra, pyramids, or prisms were introduced to close hanging nodes between levels of refinement. Mavriplis[15] implemented a variety of cell subdivision techniques for each element type in a conforming approach for fully mixed-element meshes. Each of these prior research efforts focused on single CPU operations. The implications of these methods for parallel operation and load rebalancing were not addressed.

This paper presents recent developments in the adaptation of viscous, unstructured multi-element grids in parallel computing environments, where the mesh is partitioned by domain decomposition. Attention is restricted to conforming mesh modifications, to produce adapted grids that are not solver-dependent. Parallel cell subdivision methods are developed for mixed topology meshes with hexahedral and prismatic regions, and includes an interface pyramid refinement scheme. Purely tetrahedral viscous meshes with high aspect ratio near wall cells are now treated. A load balancing strategy is implemented which permits multiple adaptation passes by preventing child cells to reside on separate processors, ensuring a consistent restart process for each partition. These advances are demonstrated on practical applications of interest with cell-vertex and cell-centered unstructured solvers

II. Mesh Adaptation Methods

A. Overview

The unstructured mesh adaptation package used in this work is *CRISP CFD*® [9], a mesh modification and quality improvement code for three-dimensional mixed-element unstructured meshes. Meshes comprised of tetrahedral, prismatic, and hexahedral regions may be readily modified to generate more accurate flow solutions through local refinement and coarsening. In moving body applications, these coarsening and refinement methods may also be employed to accommodate boundary motion.

An estimate of the solution error is obtained to drive mesh adaptation. The method currently employed[9] is based on forming a higher order approximation of the solution at each mesh point using a least squares approach. The difference between the higher order reconstruction from incident nodes and the current solution forms the error measurement. If the current mesh is sufficiently fine to support the spatial variation in the solution, the estimated error will be low, allowing coarsening to take place. Conversely, a high degree of error indicates additional refinement is needed. This approach has proven successful in a variety of applications and is capable of detecting shear layers, separation, vortical flows, and weak gradients in coarse regions, as well as shocks and expansions.

Mixed-element meshes are treated as two regions, a tetrahedral region and a prism/hex region. Current mesh generators capable of creating mixed-element meshes, such as *GRIDGEN*® [16] and *VGRIDns*[17], form the unstructured mesh in a manner that produces distinctly separate tetrahedral, prismatic, and hexahedral topologies. The tetrahedral region is treated using a Delaunay refinement and edge collapse methods, while the prism/hex region is refined using cell subdivision methods. It should be noted that while *VGRIDns* is a tetrahedral mesh generator, the near wall viscous cell layers it creates can be combined into prisms through a post-processing step.

Upon completing the mesh modifications, the solution is interpolated onto the adapted mesh using a grid-transparent procedure based on nearby point clouds[18]. The method is founded on the use of volume coordinates of a containing tetrahedron as a set of basis functions. Using an efficient octree search procedure, four close nodes in the original mesh are identified that form a tetrahedron containing the new point in the adapted grid. These nodes may belong to any type of element, and the search is independent of the underlying cell topology. The use of volume coordinates guarantees the boundedness of the reconstructed values. Higher-order reconstruction is possible by adding a quadratic correction to the linear interpolation[18].

B. Coarsening and Refinement of Tetrahedra

The tetrahedral region of the grid is locally refined by means of a constrained Delaunay refinement algorithm combined with a circumcenter point placement strategy. Any inconsistency between the circumradius of a tetrahedron and some desired point spacing triggers the point insertion procedure. This iterative cell refinement is repeated until the cell circumradii are consistent with the prescribed point spacing. Coarsening of the tetrahedral region is also permitted through an edge collapse procedure. In regions where the grid is distorted or where solution

errors are negligible, edges may be selected for removal. All cells incident to the deleted edge are removed from the mesh, the adjacent cells are redefined, and the two nodes of the edge are collapsed to a single vertex. This procedure is applicable for boundary edges as well as interior edges.

C. Refinement of Hexahedra

The hexahedral cell refinement method implemented in *CRISP CFD*® is based on the method devised Biswas and Strawn[14]. Each hexahedral cell is subdivided to generate new hexahedra, in 2:1, 4:1, or 8:1 patterns, as shown in Figure 1. The center vertex pattern of Figure 1(d) is used to terminate the propagation of refinement patterns by inserting a vertex at the cell center, creating six pyramids that may then be further subdivided to conform to the surrounding mesh. This procedure creates child cells of tetrahedra and pyramids, which are introduced between successive regions of refinement to close the grid and remove hanging nodes. In subsequent adaptation passes, these child cells are removed, the parent hexahedra are restored for further subdivision, and the cell splitting and child cell creation process is repeated. In addition to these patterns, a hexahedral cell may be subdivided into three prisms.

D. Prism Refinement

The subdivision method for prism cells borrows from that employed for hexahedra, and is more general than the approach previously explored for tet/prism grids[12,13]. Each of the prism subdivision patterns depicted in Figure 2 has been incorporated. The refinement of the three vertical edges of the cell, shown in Figures 2(a), (d), and (e), requires the introduction of child elements to close hanging nodes. The center vertex pattern accomplishes this task by splitting the original prism into two tetrahedra and three pyramids. Each of these child cells is then further subdivided into additional tetrahedra or pyramids as necessary, based on how each of the five faces of the original prism is marked.

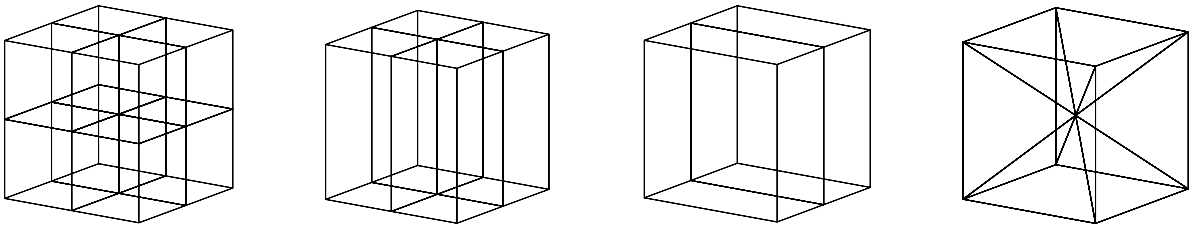

a) 8:1 b) 4:1 c) 2:1 d) Center vertex

Fig. 1. Subdivision of hexahedral cells.

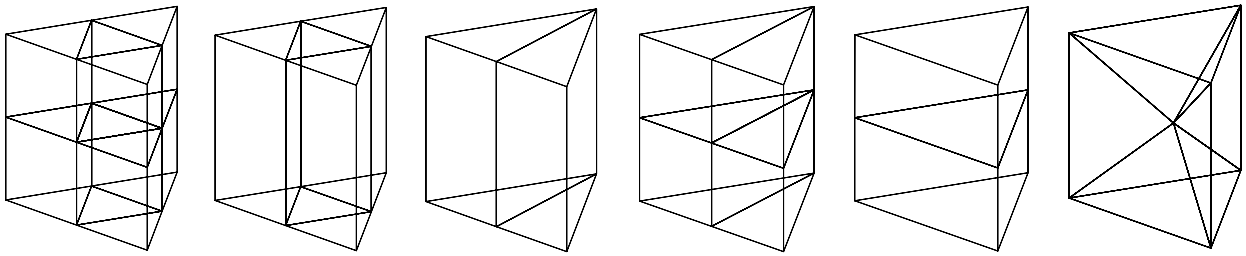

a) 8:1 b) 4:1 vertical c) 2:1 vertical d) 4:1 horizontal e) 2:1 horizontal f) Center vertex

Fig. 2. Prism subdivision patterns.

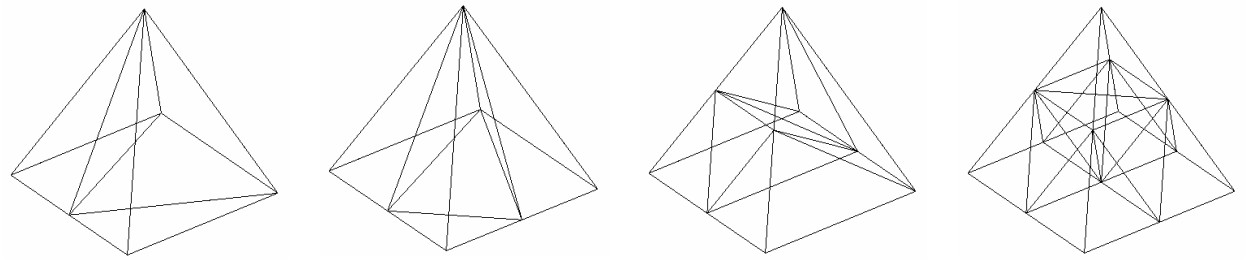

a) 1P→3T b) 1P→4T c) 1P→2P+4T d) 1P→4P+8T

Fig. 3. Subdivision of interface pyramids.

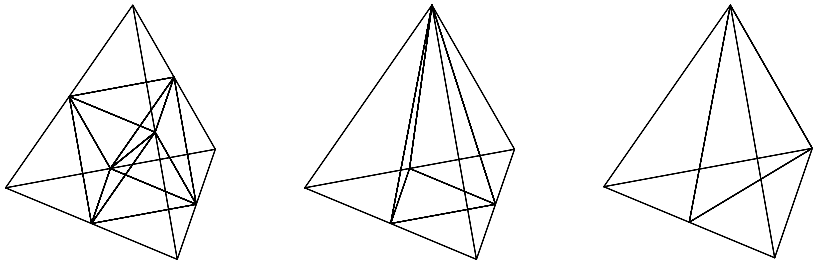

a) 8:1 b) 4:1 c) 2:1

Fig. 4. Subdivision of tetrahedral cells.

E. Interface Pyramids

Proper subdivision of any interface pyramids bridging tetrahedra with either prisms or hexahedra requires the pyramid base edges to be split in the same manner as the adjacent prism or hex cells. The pyramids are then subdivided into more pyramids and tetrahedra based on predefined patterns, illustrated in Figure 3. In subsequent passes, child pyramids formed by 2P+4T and 4P+8T splits may be subdivided further. The tetrahedra created by these patterns may be lumped with the rest of the tetrahedral region. Child tetrahedra formed by 3T and 4T splits however must be kept separate from the Delaunay refinement and coarsening processes. As the pyramid refinement progresses, the base and vertical edges are marked so that the incident tetrahedra may also be split using patterns. The incident tetrahedra are split 2:1, 4:1, or 8:1, as in Figure 4, to conform to the pyramids. Any incident boundary triangles are also subdivided in the process.

F. Viscous Tetrahedral Meshes

Viscous meshes comprised purely of tetrahedra, such as those generated using an advancing layers method like *VGRIDns* [17], are treated in two stages. The high aspect ratio, boundary layer tetrahedra are protected from the normal coarsening and Delaunay refinement methods applied to the isotropic, inviscid region of the mesh, and are instead treated in a separate cell subdivision procedure.

The boundary layer tetrahedra are first identified by computing a parameter β, which is a function of the cell aspect ratio AR and distance from the wall S.

$$\beta = \frac{AR/\overline{AR}}{S/\overline{S}} \tag{1}$$

In Eqn. (1), $\overline{AR}$ and $\overline{S}$ are the average values computed for all cells in the mesh. Including the wall distance serves to filter out high aspect ratio cells or slivers in the field that are not part of the boundary layer grid. The parameter β decays as one leaves the boundary layer region, distance to the wall increases, and the cells become more isotropic. The boundary layer cells are selected as the set of all cells where β is above a given threshold. In practice, selecting a threshold of 5 to 10 for β is sufficient to identify the boundary layer cells. Once identified, these cells are protected from the coarsening and refinement processes and are treated in a separate subdivision step. These cells are split in 2:1, 4:1, or 8:1 patterns according to their edge markings. These subdivision patterns are depicted in Figure 4. Using this approach, the structure of the boundary layer mesh is readily retained.

Figures 5 and 6 depict isosurfaces and contours of this quantity for a viscous grid. The decay in β is observed as one leaves the boundary layer region, as the cells become more isotropic and the distance to the wall increases. Figure 7 illustrates the improvement in viscous mesh adaptation for a finned missile. This case is a purely tetrahedral mesh decomposed on 4 CPU's. In Figure 7 the original approach is presented on the left and the new approach on the right. When the viscous cells are not identified and protected, the boundary layer structure is clearly lost as these "poor quality" high aspect ratio cells are removed through the Delaunay refinement and smoothing processes.

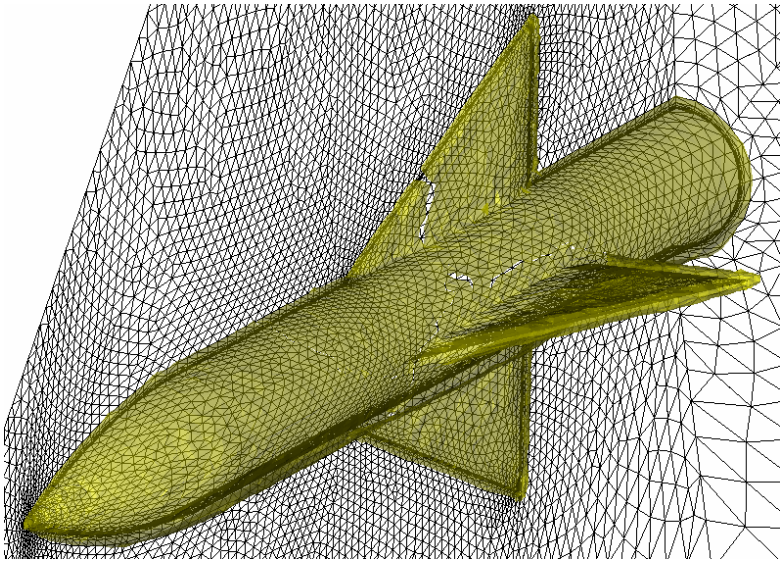

Fig 5. Isosurface of β=5 for viscous tetrahedral mesh of a finned missile.

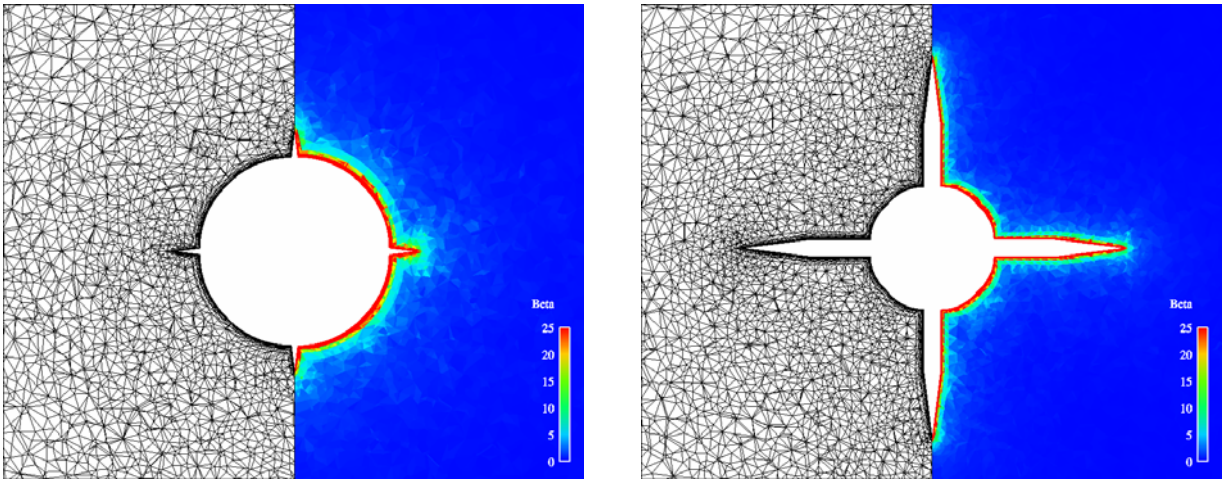

a) x=25 b) x=35

Fig. 6. Selection of boundary layer cells using β.

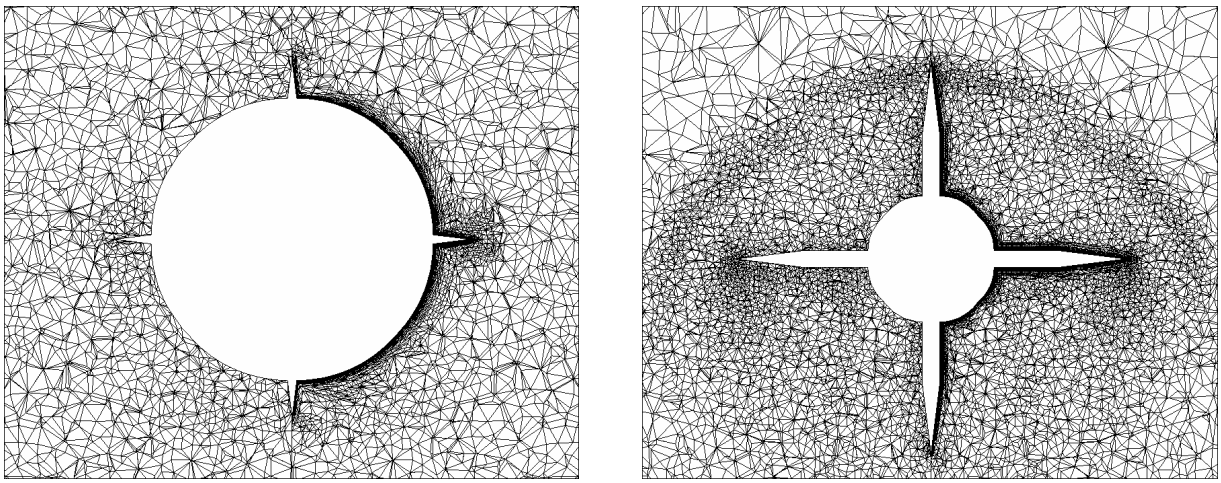

a) x = 25 b) x = 35

Fig. 7. Mesh adaptation without (left) and with (right) boundary layer cell treatment.

G. Parent Cells and Restart Data

The edge-based refinement procedures require restart data in order to restore parent cells for further subdivision in subsequent adaptation attempts. Parent cell definitions and their child cell numbers and cell types are stored for subsequent refinement passes. This parent/child cell data is formed for each parent cell type. In a fully mixed-element mesh one may have child cells of various types associated with tetrahedra, pyramids, prisms, and hexahedra. In parallel, a separate set of restart data is created for each processor after the mesh is rebalanced.

III. Parallel Implementation

A. Interprocessor Boundary Treatment

In parallel, each processor operates on its own partition, concurrent with and independent of the others. Therefore the first issue that arises in parallel adaptation is how to treat the interprocessor boundaries. In the tetrahedral region, since coarsening is potentially taking place, such changes cannot be readily synchronized among the processors. Rather than modify these faces, the interprocessor boundaries are shifted using a cell migration technique[9]. The interprocessor faces and adjacent cells then become interior faces and interior cells, which may be readily modified through a second adaptation pass. In the second pass only the former interprocessor boundary region needs to be coarsened, refined, or smoothed, as the remainder of the mesh is already consistent with the prescribed point spacing.

Figure 8 illustrates the two-pass approach for solution-based coarsening and refinement of supersonic flow entering a duct. The original mesh partitions, shown in Figure 8(a), are independently adapted to produce the mesh of Figure 8(b). Note that the interprocessor boundaries are not modified, which leaves a region of the mesh that still requires adaptation. Several layers of cells are migrated from the right processor to the left, as seen in Figure 8(c). The interprocessor boundary is now to the right of its original location. A second coarsening and refinement pass treats the former interprocessor faces and adjacent cells, producing the final adapted grid of Figure 8(d).

Parallel adaptation of the prismatic and hexahedral regions of the mesh is more straightforward. Since these elements are refined through subdivision patterns, a conforming mesh across interprocessor boundaries is ensured by exchanging nodal tags for each of the interprocessor edges marked for refinement and recomputing the patterns. This exchange guarantees that adjacent cells are subdivided in a consistent manner across processors. Figure 9 demonstrates this pattern matching on an originally hexahedral mesh, where it may be seen that adjacent cells on the two partitions are modified in a consistent manner.

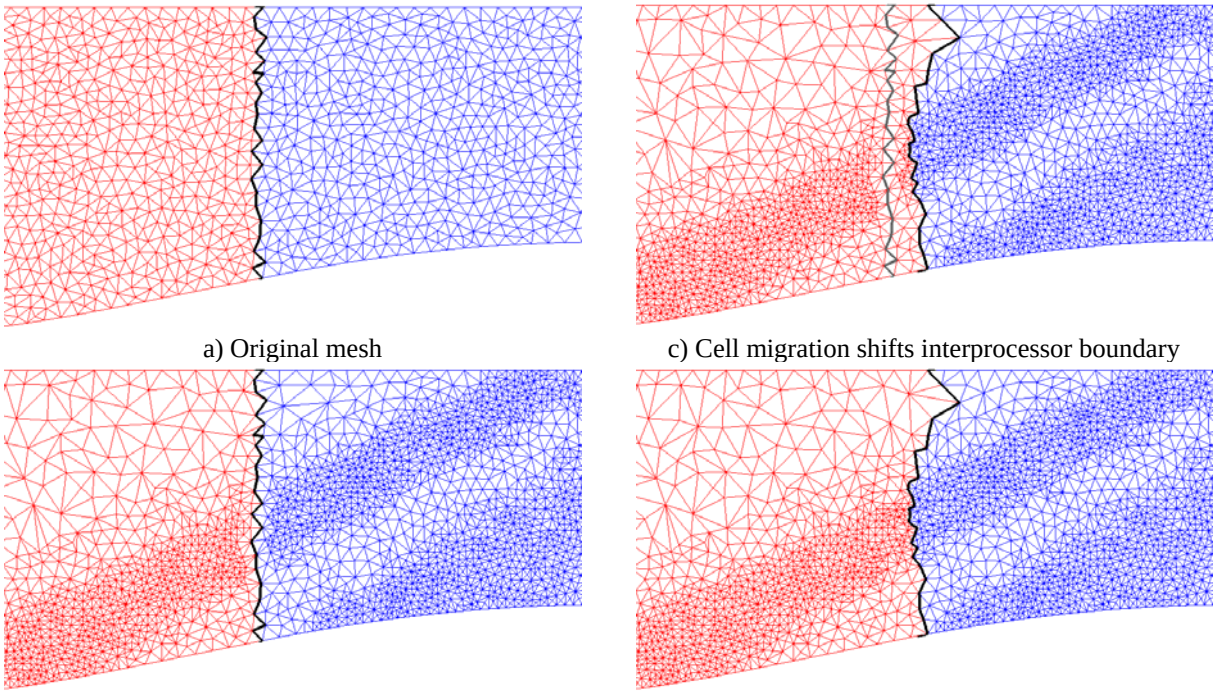

a) Original mesh

c) Cell migration shifts interprocessor boundary

b) First adaptation pass

d) Second adaptation pass

Fig. 8. Two-pass approach for parallel coarsening and refinement of tetrahedral meshes.

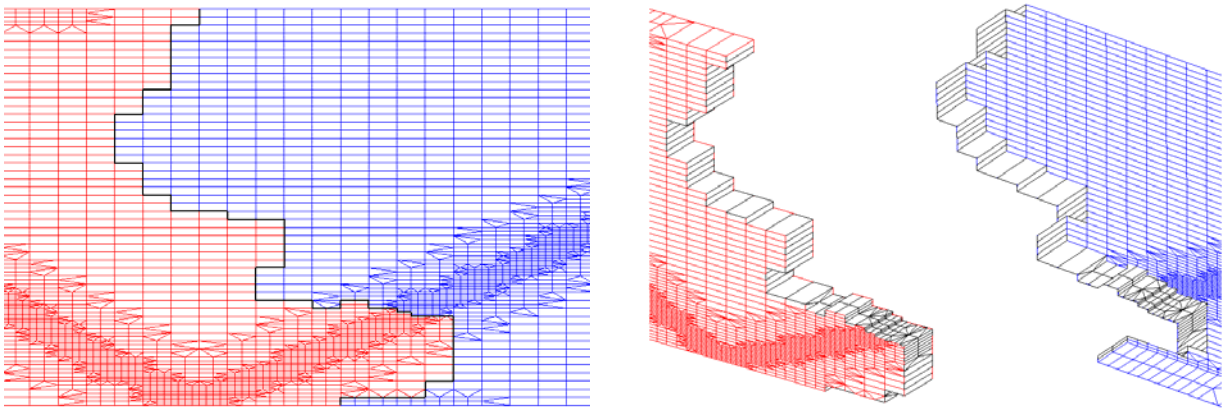

a) Side view b) Expanded view of each boundary

Fig. 9. Interprocessor boundary treatment for edge subdivisions.

B. Load Rebalancing

An adapted grid is inherently unbalanced, with mesh refinement taking place on certain partitions and coarsening taking place on others. Re-establishing load balance is desirable for steady-state problems, and essential for transient simulations, particularly after several adaptations. Load rebalancing is accomplished using the *ParMETIS* package[19] developed at the University of Minnesota.

Special care must be exercised when rebalancing the refined cells in the prism and/or hex regions of the grid. Figure 10 shows how *ParMETIS* may rebalance an adapted hexahedral mesh. The new partitions may be defined such that child cells of a given parent lie on different processors, as highlighted in Figure 10 by the dotted circle. In this instance, the parent cell may not be restored for further subdivision, and the refinement process cannot proceed. This issue is remedied by assigning all child cells belonging to the same parent to the same processor, after their initial assignments are determined by *ParMETIS*. Figure 11 illustrates an example of hexahedral refinement and rebalancing with this new constraint applied. Examining the interprocessor region, it may be seen that the partition boundary does not cut across buffer cells.

The key issues associated with rebalancing grids adapted using pattern schemes and transition elements are the assignment of child cells and formation of consistent restart data for subsequent adaptations. To properly restore a parent cell of any type, tetrahedron, pyramid, prism, or hexahedron, all child elements of that cell must reside on the same processor. After the initial cell assignments are obtained from *ParMETIS*, the child cell list is examined for each parent cell, and all child cells are assigned to the same processor as the first child in the list. The processor assignment for the parent is set in the process.

Assembling the restart data for each processor is accomplished by referring to a global list of parent cells and their children. After all cell assignments are complete, each processor forms the new set of parent cells by searching the global list for parents assigned to it. In this manner, parent cells may shift across processors during the rebalancing procedure along with their children. This is performed for each parent cell type: tetrahedra, pyramid, prism, and hexahedra.

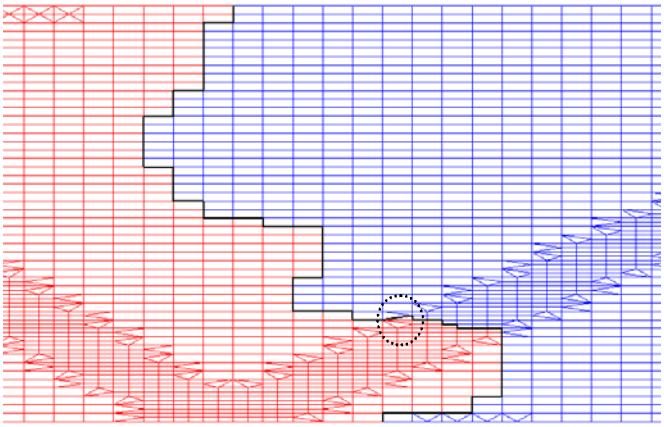

Fig. 10. Partial reassignment of child elements during rebalancing.

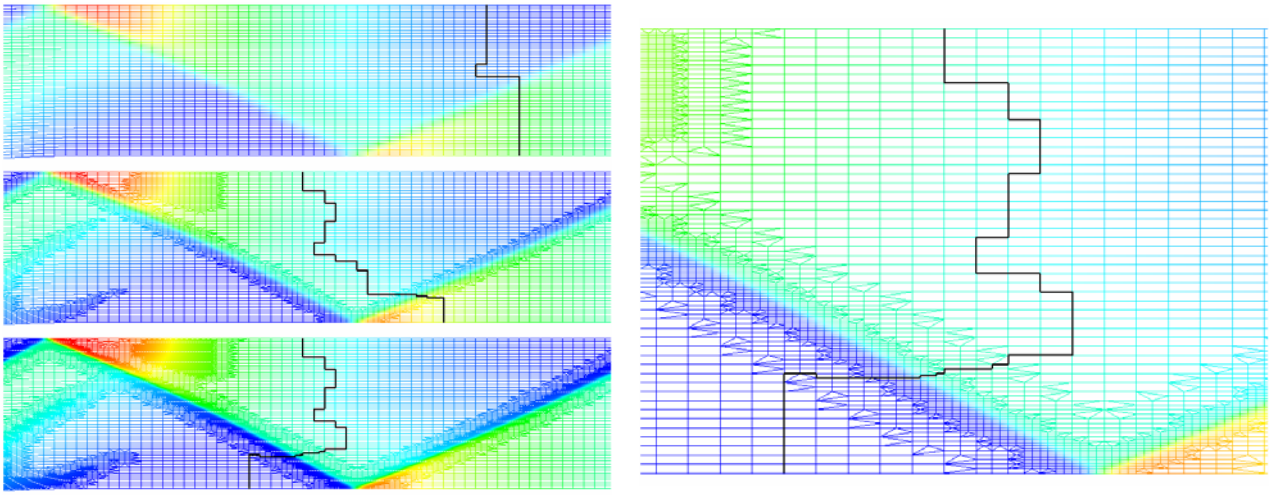

a) Multiple adaptations with load rebalancing b) Close-up of interprocessor boundary

Fig. 11. Parallel adaptation of an unstructured hexahedral mesh.

IV. Flow Solution Methods

A number of adaptation methods reported in the literature rely on embedded cell techniques where hanging nodes are permitted between levels of refinement. These non-conforming methods are designed for use with a specific solver, that is typically cell-centered or finite-element. For such solvers a list of face pairs facilitates the flux calculation between levels of refinement on the adapted grid. The mesh adaptation methods implemented in *CRISP CFD*® are designed to be solver-independent, enforcing a conforming mesh. The resulting cell and boundary face definitions may be converted into various data structures, such as edge lists or face-cell adjacency arrays, as required. Validations presented in the next section are performed using two parallel multi-element unstructured Navier-Stokes solvers, to demonstrate the flexibility of the parallel mesh adaptation package for edge-based and cell-centered solvers.

A. CRUNCH CFD®

The first flow solver used in this work is *CRUNCH CFD*® [20,21]. *CRUNCH CFD*® is a parallel, implicit solver for 3-D chemically reacting turbulent real gases and dynamic domains. The basic numerical framework of the *CRUNCH CFD*® code is a finite volume higher-order Roe/TVD scheme in which the flow variables are defined at the vertices of the mesh. An edge-based data structure is employed wherein a polyhedral control volume is constructed from the union of all cells incident to a given node, and the control volume faces are associated with each edge. The inviscid flux calculation proceeds by looping over all edges in the mesh, and is grid-transparent, while a cell-based method is employed to compute the flowfield gradients at the control volume faces for evaluating the viscous fluxes[9]. Turbulence modeling is provided by a k-ε model with various near-wall, compressibility, and EASM extensions. *CRUNCH CFD*® has also been validated for propulsive simulations involving gas phase chemical reactions.

B. AVUS

The Air Vehicles Unstructured Solver (*AVUS*)[22], formerly known as Cobalt_{60}, is an Euler/Navier-Stokes code developed in the Computational Sciences Branch of the Air Vehicles Directorate. AVUS is an unstructured, cell-centered, finite-volume, Godunov-type solver that uses least-squares gradient reconstruction and limiting for second-order spatial accuracy, and a second-order, point-implicit time integration. It handles two and three dimensions, arbitrary cell types, and has been efficiently parallelized using Message Passing Interface (MPI). It implements the Spalart-Allmaras (S-A) one-equation turbulence model, and the Menter shear-stress transport (SST) two-equation turbulence model for turbulent flow calculations (as well as others). *AVUS* has been verified on a variety of cases ranging from the exact Riemann problem to the supersonic, high angle-of-attack, flow over a missile with fins. It has also been applied successfully to complex, real-world problems.

V. Applications

A. Afterburning Missile Exhaust

This demonstration considers afterburning of the exhaust from a generic finned missile with forward strakes. The turbulent, reacting simulations are performed using *CRUNCH CFD*®. A simplified chemistry model assuming infinitely fast reaction rates is applied, in which the various exhaust species are lumped as three species of "fuel", "oxidizer", and "products". Figure 12 depicts the missile geometry along with selected stations downstream of the nozzle exit plane where the mesh and solution will be assessed. Hexahedra are used to encompass the anticipated exhaust region, while prism cells are extruded from the missile surface to provide boundary layer resolution. The prism cells easily accommodate the geometric complexity of the fins and strakes. Tetrahedra are employed to fill the remainder of the domain. Figure 13 illustrates the mesh topology employed in constructing the initial grid. The purple, orange, and green sections of the mesh denotes the tetrahedral, prismatic, and hexahedral regions, respectively. Note the exhaust grid downstream of the missile also contains a small prism section. The mesh is decomposed on 16 processors.

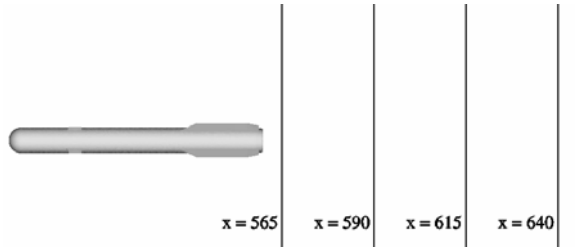

Fig. 12. Missile geometry and plume stations.

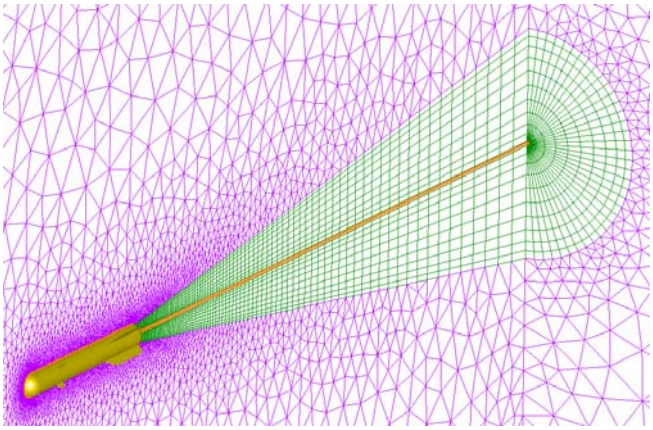

Fig. 13. Overview of missile mesh topology.

In this scenario, the missile is at 20 degrees angle of attack in a Mach 0.95 freestream. At this angle of attack, the plume is deflected significantly, to the point where it falls outside the hexahedral region, becoming very diffuse. This traversal across cell topologies is evident in Figure 14, which presents contours of product species mass fraction. Once the mesh is adapted, the flame region of the shear layer is thinner and well defined. Mesh induced diffusion is reduced as indicated by both the streamwise penetration of the jet and the deflection into the tetrahedral region. Figure 15 further illustrates the improvement with mesh adaptation. At each of the selected axial stations, mesh refinement better defines the cross-sectional shape of the deflected jet. Furthermore the refinement is seamless across the tetrahedral and hexahedral cell topologies

Table 1 summarizes the changes in mesh size with each adaptation. There is a considerable increase in the number of cells and vertices, largely owing to the subdivision procedures for the prisms and hexahedra. The marked increase in the number of tetrahedra and pyramids is predominantly due to the introduction of child cells for hanging node closure. While the current simulation was retained on 16 processors, it is possible to increase the number of processors employed as the mesh grows in size. It is also possible that a mesh movement (or *r*-refinement) scheme combined with the current methods could serve to better optimize the mesh, providing increased resolution while reducing the number of cells refined through subdivision. Such studies could be explored in future work.

Table 1. Adaptation statistics for missile exhaust case, 16 processors.

	Original Grid	First Adaptation	Second Adaptation
Vertices	226,569	601,876	1,849,645
Tetrahedra	272,984	1,145,087	2,903,460
Pyramids	3,600	206,624	618,964
Prisms	267,178	391,878	1,180,356
Hexahedra	36,000	142,427	576,409

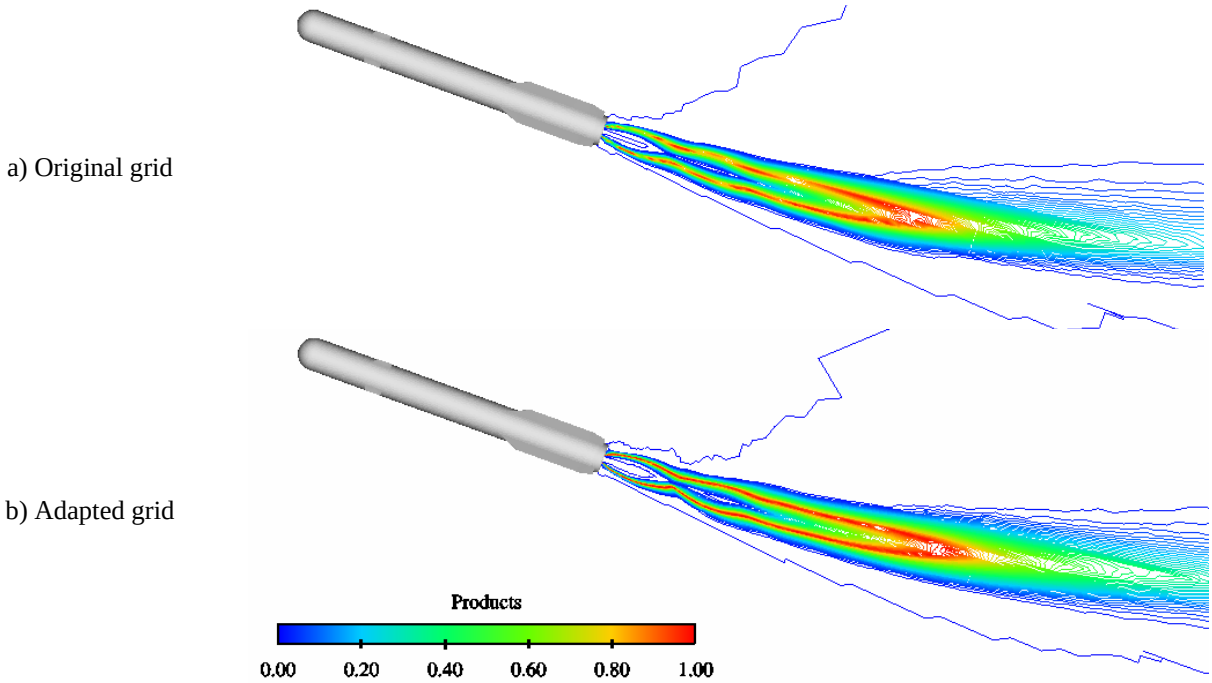

Fig. 14. Missile exhaust product distribution in the symmetry plane.

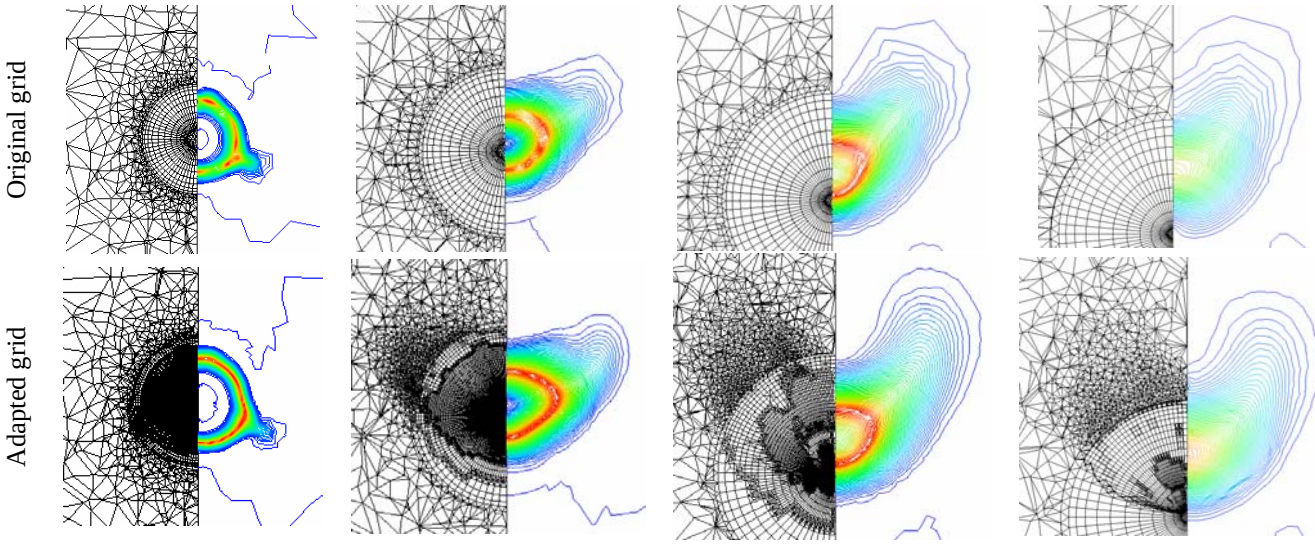

Fig. 15. Product species comparison at selected axial locations.

Figure 16 depicts the changes in the symmetry plane mesh and interprocessor boundaries. Two levels of refinement are seen in the hexahedral region of the mesh, along with refinement across topologies into the tetrahedral region as the exhaust plume is deflected by the freestream. Also note the coarsening of the tetrahedral mesh above and below the missile. The interprocessor boundaries, shown in red, are shifted each time the mesh is rebalanced. While a fairly straightforward geometry and flowfield, this case successfully demonstrates the parallel refinement and rebalancing methods described earlier.

a) Original grid

b) Adapted grid

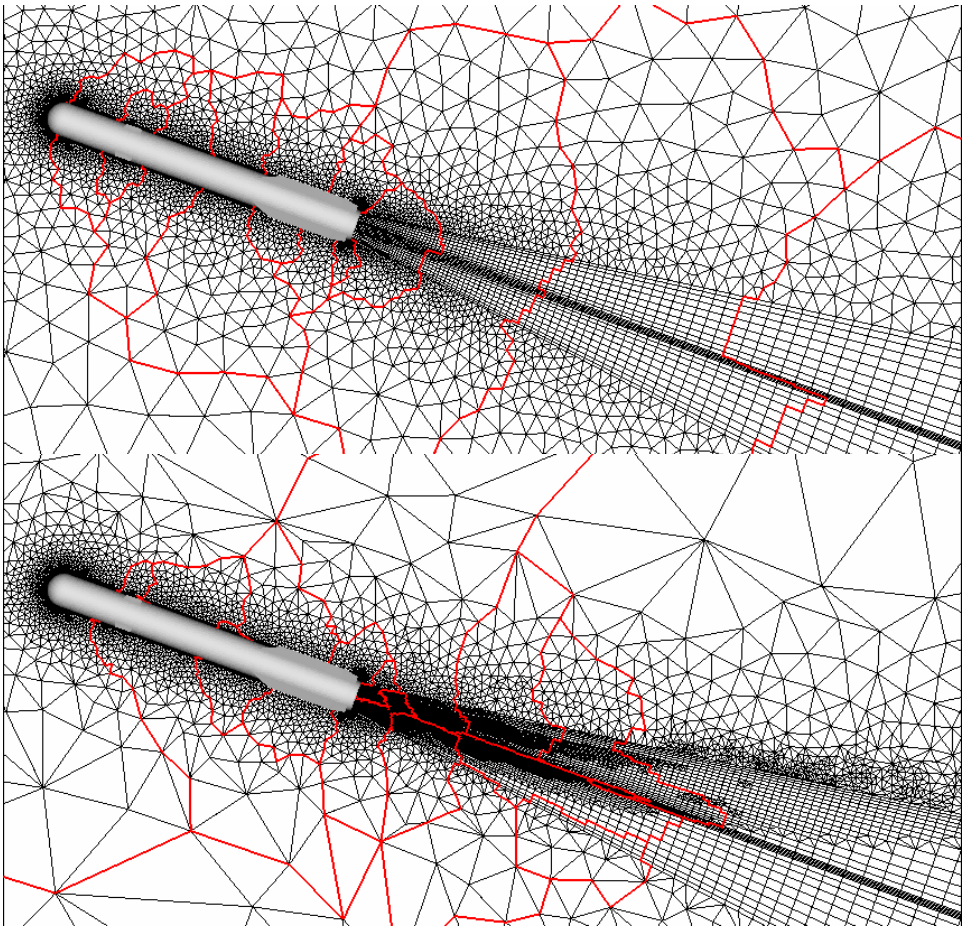

Fig. 16. Symmetry plane mesh and interprocessor boundaries for missile exhaust.

B. Hydrogen Flow Control Valve

The final demonstration is a gaseous hydrogen valve, depicted in Figure 17. In this scenario, the plug is set at a 52% open position, and an inlet pressure of 4400 psi is applied. The pressure drop across the valve is 2600 psi, which produces an underexpanded, annular jet at the nozzle throat between the plug and valve seat. The complex flowpath leads to formation of a large secondary flow region in the exhaust duct downstream of the valve seat. Turbulent simulations were performed using *CRUNCH CFD*®.

Figure 18 illustrates the mixed-element mesh topology. The purple, orange, and green sections of the mesh denote the tetrahedral, prismatic, and hexahedral regions, respectively. Hexahedral domains were constructed around the plug and in the narrow clearances of the seat region. The inlet and exhaust pipes were modeled with extruded hexahedra and prism cells. These portions of the mesh were connected to the domains in the plug and seat region using tetrahedra. The mesh was decomposed on 64 processors.

Adaptation has a notable impact on the computed flowfield, evident in the 3-D streamribbons shown in Figure 19. The streamribbons are colored by the local Mach number. With the original mesh, nozzle flow entering from above is immediately turned to the bottom of the duct, producing a tight vortical flow, while a second, weaker vortical structure forms downstream in the exhaust duct. The flowfield upstream of the valve seat is rather incoherent. In contrast, the solution on the adapted grid exhibits several differences. Upstream of the seat, the flow negotiates the plug more effectively, entering the throat region more directly. This, combined with added resolution in the annular jet region, contributes to the greater degree of penetration observed. The vortical flow in the exhaust duct is also altered, producing a more coherent, tighter vortex in the upper section of the duct, and relaxing the lower vortex where the flow first impinges on the wall. Figure 20 further illustrates the improvement in jet definition with mesh adaptation. A slice through the mesh at the y=0 plane is depicted, along with Mach number contours. After two adaptations, five distinct shock cells are observed in the annular jet. On the original mesh, these structures are diffuse and even exhibit some unsteady character.

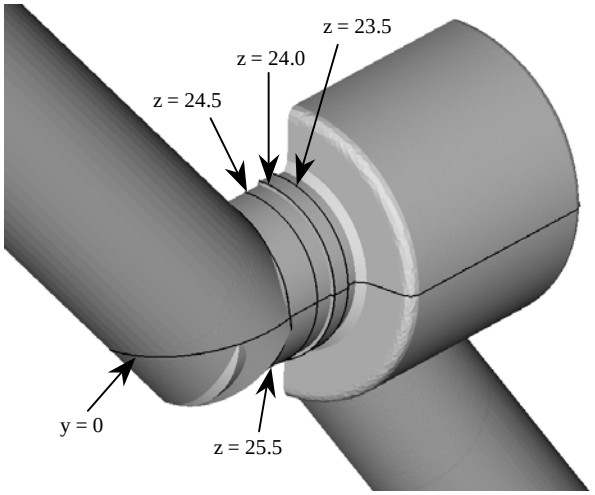

Fig. 17. Outer surface of gaseous hydrogen valve depicting selected slice.

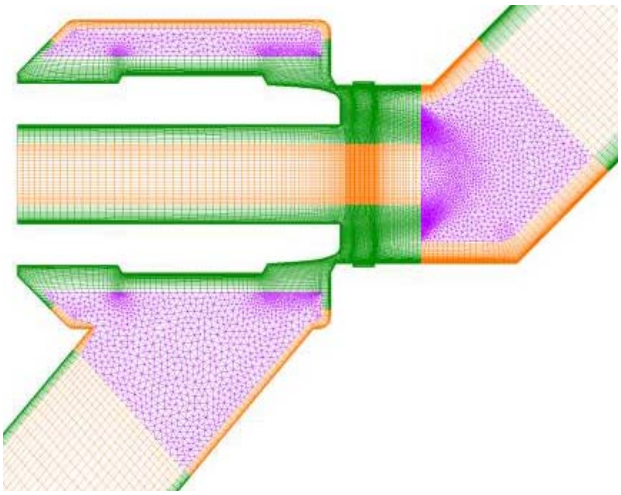

Fig. 18. Symmetry plane with overview of mixed-element mesh topology.

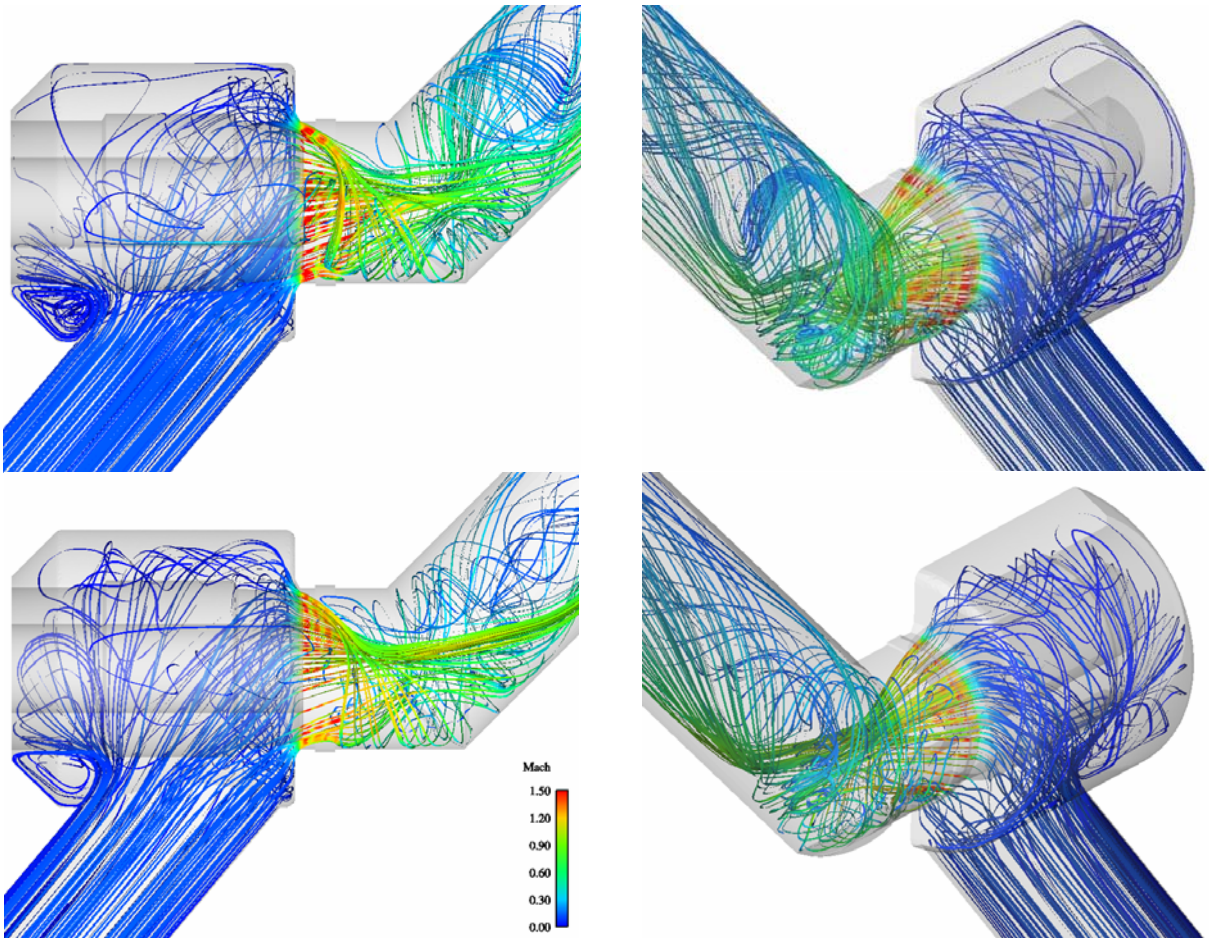

a) Side view b) Isometric view

Fig. 19. Streamribbon comparisons for original (top) and adapted (bottom) meshes.

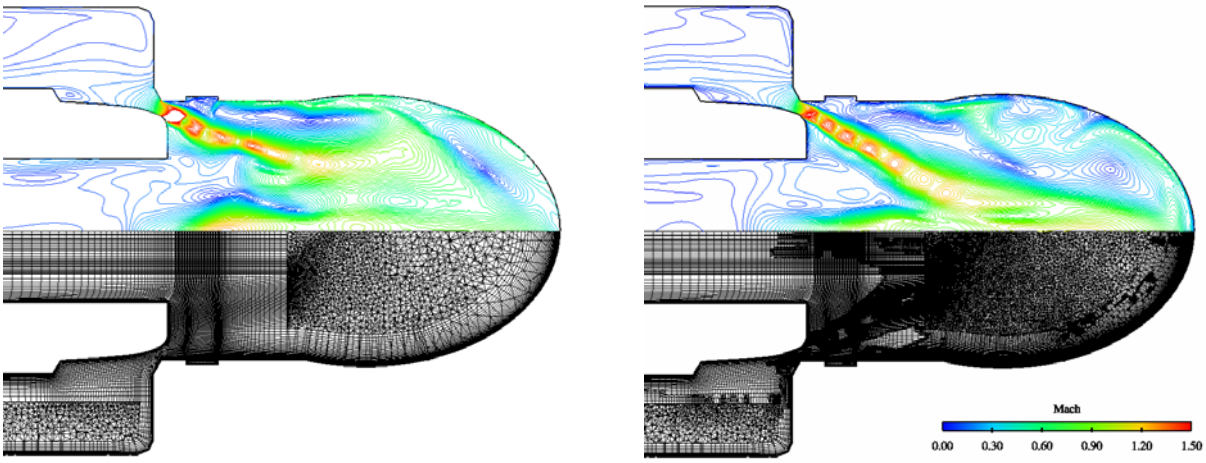

a) Original grid b) Adapted grid

Fig. 20. Mach number comparison at y=0.

Table 2 summarizes the mesh statistics for this case. Note the increase in the number of tetrahedra and pyramids, which is partly due to hanging node closure for the prisms and hexahedra. Each adaptation required less than one hour of CPU time on an IBM SP with 375 MHz Power3 processors. Currently the main contributors to CPU cost include the cell migration procedure, data structure formation and compression, global grid recomposition, and load rebalancing. CPU costs should lessen as these methods are further refined. The rebalancing of the adapted mesh among the 64 processors is illustrated in Figure 21, along with the surface mesh. The interprocessor boundaries, shown in red, are shifted to maintain an approximately equal number of cells on each processor as the mesh is adapted. This case represents a practical application of the parallel methods in *CRISP CFD*® for a complex mesh of significant size. A mixed-element grid comprised of 3.3 million vertices and 4 million cells was readily treated.

Table 2. Adaptation statistics for hydrogen valve case, 64 processors.

	Original Grid	First Adaptation	Second Adaptation
Vertices	1,338,624	3,318,622	4,821,025
Tetrahedra	828,820	5,248,062	6,639,388
Pyramids	10,650	820,213	1,291,500
Prisms	437,626	878,241	1,560,404
Hexahedra	936,290	1,727,572	2,486,150

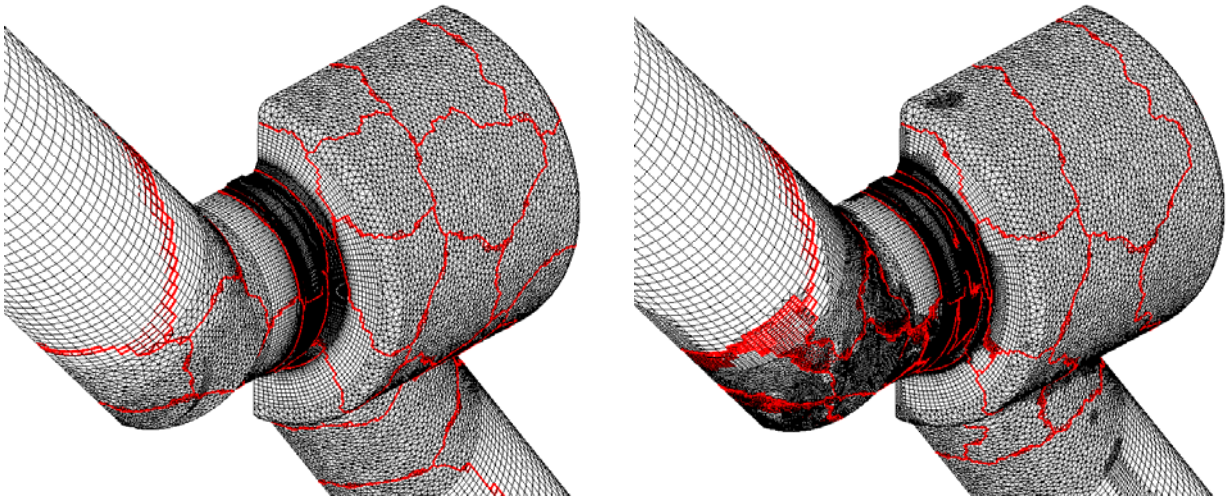

a) Original grid b) Adapted grid

Fig. 21. Surface mesh and interprocessor boundaries for hydrogen valve.

VI. Conclusions

A parallel, unstructured mesh adaptation package, *CRISP CFD*®, has been extended to permit repeated refinement of mixed element meshes comprised of any combination of tetrahedra, pyramids, prisms, or hexahedra. These recent developments in parallel, adaptive methods expand the ability of unstructured Navier-Stokes solvers to provide rapid, accurate detailed analyses and/or design assessments. Repeated adaptation allows the analyst to obtain grid-converged results for complex flowfields where the locations of relevant structures of interest are not known *a priori*. Validation studies were performed on selected problems illustrating the utility of the package for different flow solvers in practical steady-state applications on multi-million cell models employing up to 64 processors. In addition to *CRUNCH CFD*® and *AVUS*, the code presently supports *USM3D*[23] and the *FAST*[24] format. Filters for additional unstructured grid solvers may be readily developed for use with *CRISP CFD*®.

Continued research focuses on transient applications. Specific areas of research include maintaining control over the size of the mesh, and developing error control strategies for automated adaptation. Constraining the mesh size is advantageous for unsteady problems where adaptation may be performed dozens of times, and an approximately constant CPU time per iteration is desirable. Automating the decision to adapt the grid is also an important element of coupled flow solution and adaptation. In transient solutions a user-specified frequency would be both problem-dependent and inadequate in cases where changes in the solution may be rapid. Various candidate strategies are currently being explored.

Acknowledgements

Timothy Baker of Princeton University developed and provided the solution interpolation method used in this work. The Air Force Research Laboratory Computational Sciences Branch at Wright-Patterson AFB funded this research under contract F33615-02-C-3215, Phase II SBIR. Support for the hydrogen valve study was provided by NASA Stennis Space Center under contract NNS04AA08C, with Peter Sulyma as technical monitor. Jeremy Shipman and Bruce Crawford of Combustion Research and Flow Technology provided the initial mesh and solution for the valve validation case.

References

[1] De Keyser, J., and Roose, D., "Run-Time Load Balancing Techniques for a Parallel Unstructured Multi-Grid Euler Solver with Adaptive Grid Refinement", *Parallel Computing*, Vol. 21, pp. 179-198, 1995.

[2] MacNeice, P., Olson, K.M., Mobarry, C., de Fainchtein, R., and Packer, C., "PARAMESH: A Parallel Adaptive Mesh Refinement Community Toolkit", *Computer Physics Communications*, Vol. 126, pp. 330-354, 2000.

[3] Soni, B.K., Koomullil, R., Thompson, D.S., and Thornburg, H., "Solution Adaptive Grid Strategies Based on Point Redistribution", *Computer Methods in Applied Mechanics and Engineering*, Vol. 189, pp. 1183-1204, 2000.

[4] Jones, M.T., and Plassman, P.E., "Parallel Algorithms for Adaptive Mesh Refinement", *SIAM Journal on Scientific Computing*, Vol. 18, No. 3, pp. 686-708, 1997.

[5] Barry, W.J., Jones, M.T., and Plassman, P.E., "Parallel Adaptive Mesh Refinement Techniques for Plasticity Problems", *Advances in Engineering Software*, Vol. 29, No. 3-6, pp. 217-225, 1998.

[6] Oliker, L., Biswas, R., and Gabow, H.N., "Parallel Tetrahedral Mesh Adaptation with Dynamic Load Balancing", *Parallel Computing*, Vol. 26, pp. 1583-1608, 2000.

[7] Flaherty, J.E., Loy, R.M., Shephard, M.S., Szymanski, B.K., Teresco, J.D., and Ziantz, L.H., "Adaptive Local Refinement with Octree Load Balancing for the Parallel Solution of Three-Dimensional Conservation Laws", *Journal of Parallel and Distributed Computing*, Vol. 47, pp. 139-152, 1997.

[8] Niekamp, R., and Stein, E., "An Object-Oriented Approach for Parallel Two- and Three-Dimensional Adaptive Finite Element Computations", *Computers and Structures*, Vol. 80, pp. 317-328, 2002.

[9] Cavallo, P.A., Sinha, N., and Feldman, G.M., "Parallel Unstructured Mesh Adaptation for Transient Moving Body and Aeropropulsive Applications", AIAA Paper 2004-1057, 42nd Aerospace Sciences Meeting and Exhibit, Reno, NV, January 5-8, 2004.

[10] Cavallo, P.A., and Sinha, N., "An Adaptive Unstructured Mesh Approach for Aircraft/Store Compatibility Assessment," Proceedings of NATO Symposium RTO-AVT-108, Williamsburg, VA, June 7-10, 2004.

[11] Shipman, J.D., Cavallo, P.A., and Hosangadi, A., "Efficient Simulation of Aircraft Exhaust Plume Flows using a Multi-Element Unstructured Methodology", AIAA Paper 2001-0598, 39th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 8-11, 2001.

[12] Parthasarathy, V., and Kallinderis, Y., "Adaptive Prismatic-Tetrahedral Grid Refinement and Redistribution for Viscous Flows", *AIAA Journal*, Vol. 34, No. 4, pp. 707-716, April 1996.

[13] Cavallo, P.A., and Baker, T.J., "Efficient Delaunay-Based Solution Adaptation for Three-Dimensional Unstructured Meshes", AIAA Paper 2000-0809, 38th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 10-13, 2000.

[14] Biswas, R., and Strawn, R.C., "Tetrahedral and Hexahedral Mesh Adaptation for CFD Problems", *Applied Numerical Mathematics*, Vol. 26, pp. 135-151, 1998.

[15] Mavriplis, D.J., "Adaptive Meshing Techniques for Viscous Flow Calculations on Mixed Element Unstructured Meshes", NASA CR-201675, 1997.
[16] Gridgen, Grid Generation Software Package, Ver. 15, Pointwise, Inc., Ft. Worth, TX, 2004.
[17] Pirzadeh, S., "Progress Towards a User-Oriented Unstructured Viscous Grid Generator", AIAA Paper 96-0031, 34th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 15-18, 1996.
[18] Baker, T.J., "Interpolation from a Cloud of Points", Proceedings of the 12th International Meshing Roundtable, Santa Fe, NM, September 14-17, 2003.
[19] Schloegel, K., Karypis, G., and Kumar, V., "A Unified Algorithm for Load-balancing Adaptive Scientific Simulations", Technical Report, 00-033, University of Minnesota, Department of Computer Science and Engineering, 2000.
[20] Hosangadi, A., Lee, R.A., Cavallo, P.A., Sinha, N., and York, B.J., "Hybrid, Viscous, Unstructured Mesh Solver for Propulsive Applications", AIAA Paper 98-3153, 34th Joint Propulsion Conference, Cleveland, OH, July 13-15, 1998.
[21] Hosangadi, A., Lee, R.A., York, B.J., Sinha, N. and Dash, S.M., "Upwind Unstructured Scheme for Three-Dimensional Combusting Flows", *Journal of Propulsion and Power*, Vol. 12, No. 3, pp. 494-503, May-June 1996.
[22] Strang, W.Z., Tomaro, R.F., and Grismer, M.J., "The Defining Methods of $Cobalt_{60}$: A Parallel, Implicit, Unstructured Euler/Navier-Stokes Flow Solver", AIAA Paper 99-0786, 37th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 11-14, 1999.
[23] Frink, N.T., "Tetrahedral Unstructured Navier-Stokes Method for Turbulent Flows", *AIAA Journal*, Vol. 36, No. 11, pp. 1975-1982, November 1998.
[24] Walatka, P.P., Clucas, J., McCabe, R.K., Plessel, T., and Potter, R., "FAST User Guide", NAS Technical Report RND-93-010, June 1993.