




Parallel Geometry-Based Mesh Adaptation for Supersonic Jet Acoustic Analysis

Chase Ashby, Gerrit-Daniel Stich, Jeffrey Housman,
Jared Duensing


NASA Ames Research Center

November 2024


Computational fluid dynamics enables the study of fluid flows using HPC




Physical space discretization results in meshes



Cartesian




Curvilinear




Voronoi

Mesh adaptation reduces discretization errors


- Adaptation methodologies:
 - r-adaptation: redistribute points (fixed degrees of freedom)
 - h-adaptation: add or remove degrees of freedom
 - p-adaptation: increase polynomial order



Initial grid




h-adaptation




r-adaptation


Launch, Ascent, and Vehicle Aerodynamics (LAVA)



Mesh adaptation has been successfully applied to supersonic applications



Mesh adaptation has been utilized to enhance shock resolution to address solution sensitivities



CFD is assisting in the design of new quiet propulsion systems

- NASA's Commercial Supersonic Technology project aims to develop propulsion systems that surpass current noise regulations
- Dual-stream nozzle configurations are a key focus area
- A variety of configurations are being tested, some of which include lobed mixer designs







Image credit: Timothy Sandstrom




Body-fitted grids for the m5 lobed mixer are difficult to generate and increase computational costs




Immersed boundary methods avoid the need to be body-fitted for complex geometries




Adaptation indicators are computed to target geometric boundaries via wall-distance calculations




Geometry-based mesh redistribution provides near-wall resolution



Geometry-based mesh redistribution provides near-wall resolution



Geometry-based adaptation is applicable to non-smooth surfaces



The solution procedure involves repeated parallel interpolation of data


Nonlinear Solution Procedure

- 1: Initialize global vectors
- 2: **while** $\|\mathbf{R}_{\text{adapt}}\|_2 > \text{tolerance}$ **do**
- 3: Solve linear system for Newton step
- 4: **while** Armijo-Goldstein condition not met **do**
- 5: Update search direction
- 6: Update coordinates
- 7: Interpolate indicator
- 8: Update control functions
- 9: **end while**
- 10: **end while**


Potential bottleneck




LAVA uses a tile-based partitioning strategy




Grids are divided into sets
of tiles with fixed size



Tiles are distributed using ParMetis




Two processors



Five processors


Parallel interpolation leverages AD-tree data structures

- Interpolation requires:
 - Containing cells
 - Trilinear interpolation
- Bounding boxes are computed for each 'owned' portion of the tiles
- A hexahedral bounding box AD-tree is generated using the global set of tile bounding boxes
 - Each processor retains ownership of the global tile bounding box tree
- Given a search point, we can quickly determine which tile its containing cell belongs to



Computed bounding boxes (black) for each tile

Strong scaling test on overlapping structured grids



Parallel interpolation algorithm demonstrates strong scalability


MPI Ranks	Threads per Rank	Wall-Clock Time (s)	Ideal Speedup	Speedup
1	1	151.35	1.00	1.00
1	2	75.49	2.00	2.00
2	2	37.94	4.00	3.99
2	4	18.99	8.00	7.97
4	4	9.52	16.00	15.90
4	8	4.80	32.00	31.51
8	8	2.58	64.00	58.69
8	16	1.30	128.00	116.73
16	16	0.66	256.00	230.40
16	32	0.39	512.00	390.93
32	32	0.27	1024.00	564.02
32	64	0.10	2048.00	1460.81
64	64	0.06	4096.00	2764.96

HPC Architecture:


Aitken supercomputer at NASA Ames Research Center

AMD Rome nodes:

Two 64-core AMD EPYC 7742 Rome processors
(2.25 GHz) and 512 GB of memory.



Adaptation is enabled by immersing the m5 lobed mixer in a cylindrical multi-block grid



Adaptation Computing Resources

- Number of grid nodes: 34,319,600
- HPC Architecture:
 - 20 Skylake nodes
 - 2 MPI ranks per host
 - 20 threads per rank
- Wall-clock time:
 - 8 minutes 53 seconds

HPC Architecture:


Electra supercomputer at NASA Ames Research Center
Skylake Nodes

Two 20-core Xeon Gold 6148 sockets (2.4 GHz)
and 192 GB of memory.




<https://www.nasa.gov/hecc/resources/electra.html>


Adaptation indicator is constructed to highlight near-wall region




Near-wall resolution improved without sacrificing speed



Adaptation improves the numerical representation of the geometric boundary




Uniform





Adapted

Noise spectra accuracy is enhanced



Farfield directivity compares well to body-fitted approach



Conclusions and Future Work

- Geometry-based mesh redistribution enables higher quality immersed boundary grid systems
- Geometry-based mesh redistribution offers an alternative to structured overset gridding methodologies for complex geometries
- Results obtained for jet acoustic analysis of a dual-stream nozzle comparable favorably to traditional body-fitted approaches
- Future work will focus on the following aspects of the adaptation:
 - Controllable wall resolution
 - Automatic surface alignment
 - Combined geometry- and adjoint-based adaptation



Acknowledgements

- Funded by NASA's Aeronautics Research Mission Directorate (ARMD) Commercial Supersonic Technologies (CST) and Transformational Tools and Technologies (TTT) projects.
- Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.
- All LAVA team members



Questions?

chase.p.ashby@nasa.gov
jared.c.duensing@nasa.gov