

PARALLEL ADAPTIVE HIGH-ORDER CFD SIMULATIONS CHARACTERIZING SOFIA CAVITY ACOUSTICS

MICHAEL F. BARAD, CHRISTOPH BREHM,
CETIN C. KIRIS, AND RUPAK BISWAS

NASA Ames Research Center
Moffett Field, CA 94035, USA
web page: <http://www.nas.nasa.gov/>
e-mail: michael.f.barad@nasa.gov, christoph.brehm@nasa.gov,
cetin.c.kiris@nasa.gov, and rupak.biswas@nasa.gov

Key words: Higher Order, Immersed Boundary, Parallel CFD applications

Abstract. This paper presents large-scale MPI-parallel computational fluid dynamics simulations for the Stratospheric Observatory for Infrared Astronomy (SOFIA). SOFIA is an airborne, 2.5-meter infrared telescope mounted in an open cavity in the aft fuselage of a Boeing 747SP. These simulations focus on how the unsteady flow field inside and over the cavity interferes with the optical path and mounting structure of the telescope. A temporally fourth-order accurate Runge-Kutta, and spatially fifth-order accurate WENO-5Z scheme was used to perform implicit large eddy simulations. An immersed boundary method provides automated gridding for complex geometries and natural coupling to a block-structured Cartesian adaptive mesh refinement framework. Strong scaling studies using NASA's Pleiades supercomputer with up to 32k CPU cores and 4 billion computational cells shows excellent scaling. Dynamic load balancing based on execution time on individual AMR blocks addresses irregular numerical cost associated with blocks containing boundaries. Limits to scaling beyond 32k cores are identified, and targeted code optimizations are discussed.

1 Introduction

Stratospheric Observatory for Infrared Astronomy (SOFIA) is an airborne, 2.5-meter infrared telescope mounted in a large open cavity in the aft fuselage of a Boeing 747SP, as shown in Figure 1a. The aircraft typically flies high in the atmosphere, between 12 to 14 kilometers (39,000 to 45,000 feet), reducing atmospheric distortion and enabling high quality images. SOFIA is a joint research program between NASA and the German Aerospace Center (DLR). In this work, we focus on 1) noise generation and flow physics in the open cavity where the telescope is mounted, and 2) assessment of solver parallel performance for up to 32k cores for this multi-scale, demanding application with a complex geometry.

Higher-order accurate implicit large eddy simulations (ILES) using NASA's Launch Ascent and Vehicle Aerodynamics (LAVA) CFD solver [3] were performed. These simulations become computationally very expensive due to the requirement of accurately

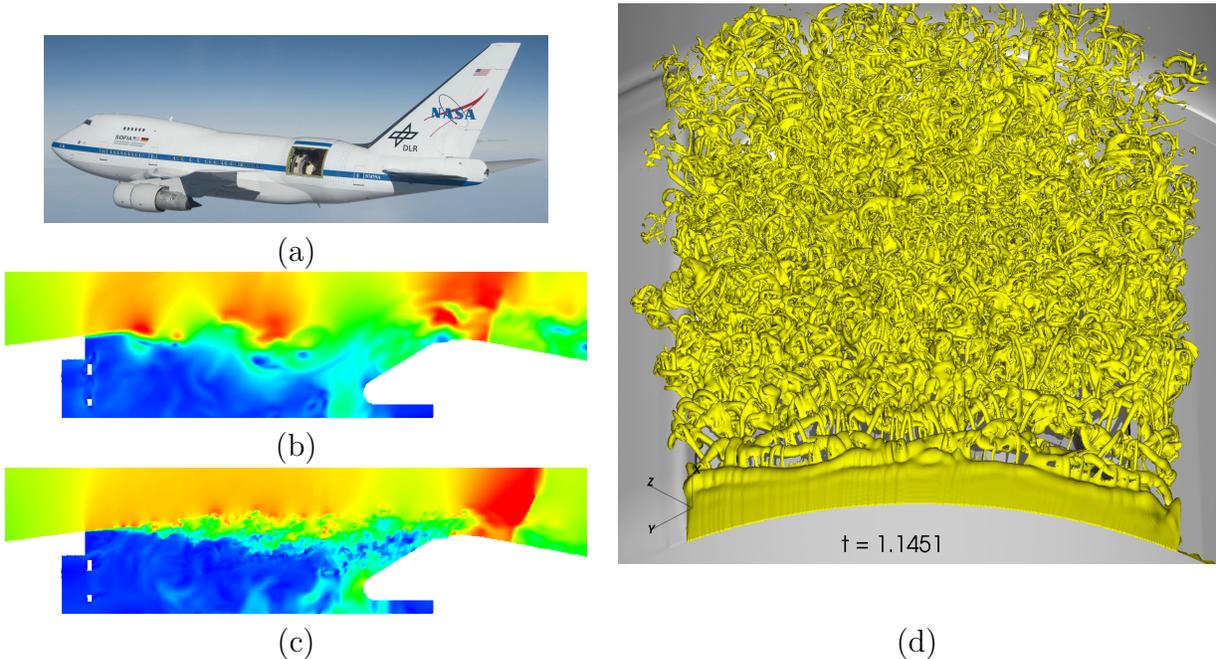


Figure 1: (a) SOFIA aircraft and aft open cavity where a 2.5-meter infrared telescope is mounted. Slices colored by Mach number showing sensitivity of flow to mesh resolution: (b) coarse grid, (c) fine grid. (d) vorticity iso-surfaces in the shear layer at the cavity opening.

capturing a wide range of physically and temporally relevant scales. Numerical methods with high spectral accuracy are required to efficiently simulate these types of problems. The computation of transonic flows is further complicated by the occurrence of shocks which generate discontinuities in the flow field. In order to handle these discontinuities, we apply state-of-the-art higher-order shock capturing schemes. LAVA’s immersed boundary method was used to circumvent the very time-consuming (e.g. structured curvilinear) volume mesh generation process for the complex SOFIA geometry. The current ghost cell based immersed boundary method is a natural fit with the block-structured Cartesian adaptive mesh refinement framework.

The current simulations focused on how the unsteady flow field inside and over the open cavity interferes with the optical path and mounting of the telescope. Results include the identification of two important mechanisms potentially affecting image clarity: (1) density variations in the shear layer leading to image distortion, and (2) vibrations of the telescope structure perturbing the optical path, hence, reducing image quality. The “secondary” breakdown process of the shear layer involves the interaction and generation of vortical flow structures with significant density fluctuations disturbing the optical path.

2 Parallelization Strategy

The parallelization strategy currently relies on a pure MPI-based communication pattern based on an SPMD paradigm. The flow domain is decomposed into a hierarchy of refinement levels, and on each level the domain is decomposed into boxes. We impose that the boxes are defined on a box-index-space lattice for each level. Boxes are distributed

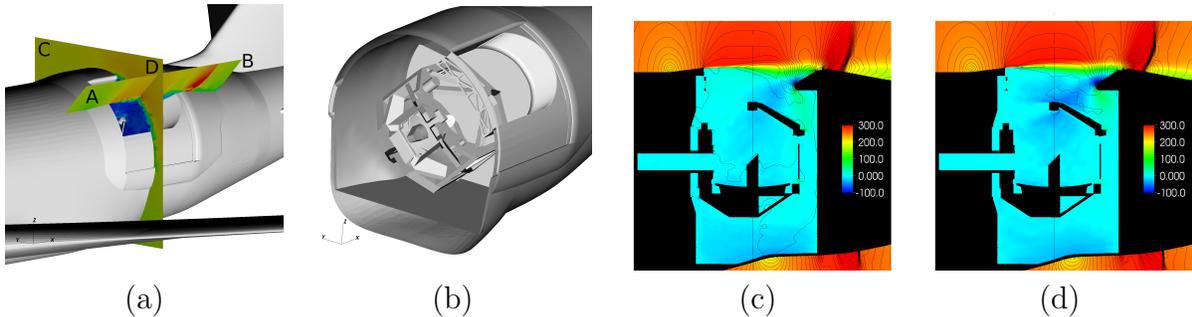


Figure 2: Two slices: A-B tilted at 45 degrees, and C-D whose normal is parallel to freestream flow direction. Slices are colored by Mach number used for post-processing analysis. (a) shows opaque walls, while (b) shows the telescope/cavity with hidden upstream parts. (c) and (d) present mean flow slices A-B colored by streamwise velocity, with pressure contour lines showing sensitivity to mesh resolution. (c) coarse grid, (d) fine grid.

to processors using a weighted locality preserving knapsack algorithm, where weights are either uniform per box, or dynamically computed using cumulative execution time for each box on the lattice. Dynamic load balancing is used to address irregularities in the performance per box, typically caused by the immersed boundary procedure. All boxes contain three ghost cells per side due to the WENO-5Z scheme. At box interfaces for each level, we do traditional ghost cell exchanges. At coarse-fine interfaces we interpolate to the ghost cells.

Complex geometry is defined using water tight surface triangulations. Triangulations are read from disk at the start of the simulation by each MPI rank, and multi-resolution distance function and binning algorithms are utilized for distance and in/out queries required by the immersed boundary method. Each MPI rank contains the data necessary to independently respond to geometry queries, resulting in a completely parallel immersed boundary method.

3 Results

3.1 Cavity Flow Physics

CFD based analysis is now presented with the aim of identifying and quantifying acoustic noise generation mechanisms, including shear layer, flow impingement. The unsteady pressure field in and around the cavity was analyzed by using disturbance flow fields, spectral analysis, and proper orthogonal decomposition (POD not shown here).

Figure 2a defines two planar slices used in the remainder of the paper: A-B tilted at 45 degrees, and C-D where the normal is parallel to freestream flow direction. Figure 2b shows details of the telescope in the SOFIA cavity. Most grid points are clustered in the shear layer region where the flow unsteadiness originates. The effect of the grid resolution on the flow field is demonstrated in Figures 2c-d.

Mach number contours for the coarse mesh Figure 1b (~ 170 million cells), and fine mesh Figure 1c (~ 300 million cells) visualize the transonic flow field around the cavity opening. This is characterized by the highly unsteady shear layer and the interaction of

it with the aft aperture as well as with a re-compression shock located just downstream of the cavity opening. The simulation with refined mesh (Figure 1c) shows an earlier breakup of the shear layer and finer scale vortical structures in comparison to the coarse grid simulation (Figure 1b). The interaction between the unsteady shear layer and the re-compression shock is reduced as a consequence of the smaller vortical flow structures.

The time-averaged colored streamwise velocity contours and solid pressure contour lines for the fine and coarse grids can be compared in Figure 2. At first glance, the general mean flow field features appear very similar for both simulations. A closer look reveals some differences in the mean flow distortions of the velocity and pressure fields. The mean pressure distribution along the shear layer is slightly different, and the downwash at the shear layer impingement point at the aft aperture is increased for the fine grid simulation. Moreover, the fine grid allows for a better resolution of the shock structure.

The iso-surfaces of vorticity in Figure 1d provide an overview of the general features and complexity of the flow field. At the inception of the shear layer, the spanwise roll-up of the vortex sheet, characteristic of the primary instability, i.e., Kelvin-Helmholtz instability, can be observed. Shortly downstream a spanwise modulation occurs introducing 3D flow structures. Subsequent breakdown of coherent structures results in fine scale turbulent-like flow features. Note that at these flow conditions a grid resolution for direct numerical simulations (DNS), where all temporal and spatial scales are resolved down to the Kolmogorov length scale, would require orders of magnitude more grid points.

3.2 Spectral Analysis

In contrast to small differences in mean flow distortion, the differences in the unsteady part of the solution $u' = u - \bar{u}$ is very significant. Figure 3 shows the breakdown of large coherent structures on the fine mesh which reduces the total power spectral density (PSD) values. A sound pressure level of 120-130dB was determined at the rear wall of the cavity for the fine simulation. These values are consistent to the results reported by Cummings [2]. It is well known that not resolving the unsteady flow features will lead to highly energetic spanwise coherent flow structures. This strongly coherent flow field tends to cause elevated acoustic noise levels in and around the cavity. The fine grid solution provides an improved mechanism for the energy cascade where the spanwise coherence is reduced and fine scale turbulent like flow features can be observed. It must be pointed out that telescope image distortion depends strongly on the actual shear layer thickness. In the large eddy simulation, the spreading rate of the shear layer in turn depends strongly on the eddy viscosity. Unsteadiness in the flow field needs to be well resolved to accurately model the shear layer mixing and provide precise predictions of the eddy viscosity. As a side note, Reynolds-averaged Navier-Stokes (RANS) simulation strategies are computationally far less expensive than the current ILES approach. RANS approaches perform very well in predicting eddy viscosities for attached wall bounded flows but for separated flows and shear layers, reliable data cannot be expected. This is why we chose the more expensive ILES approach. In addition, RANS approaches do not provide insight about the occurrence of unsteady cavity modes discussed below.

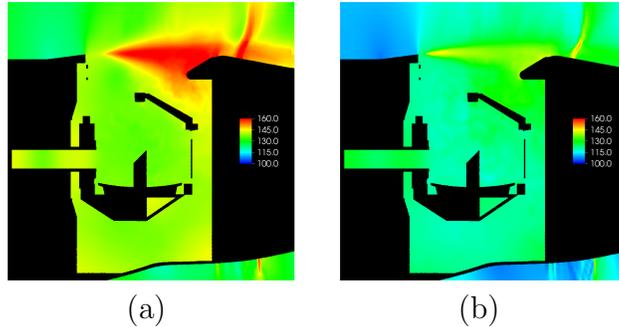


Figure 3: Slice A-B (see Figure 2a) colored by total power spectral density (PSD) [dB] showing sensitivity of mean flow to mesh resolution. (a) coarse grid, (b) fine grid.

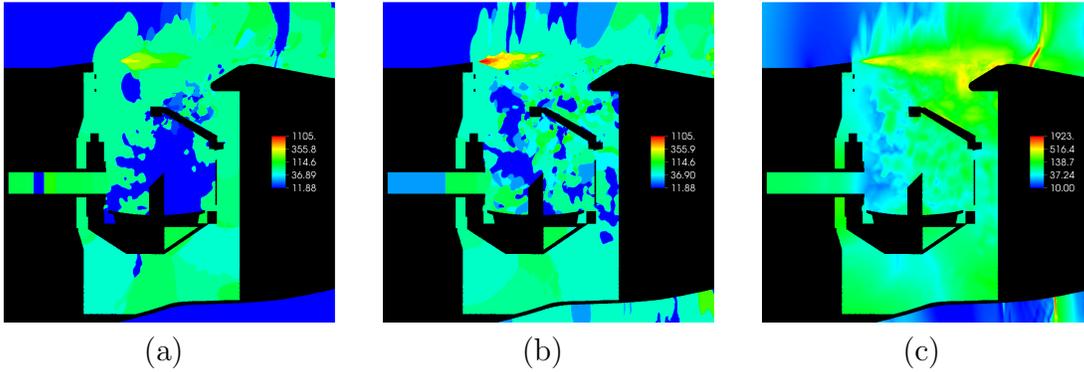


Figure 4: Slices A-B as in Figure 2a colored by spectral frequency [Hz] at peak amplitude showing sensitivity of mean flow to mesh resolution. (a) coarse grid, (b) fine grid. (c) Fine grid slice colored by peak spectral amplitude, showing locations with dominant frequencies.

The spectral characteristics of the unsteady flow and the origin of these peaks in the pressure Fourier spectra is analyzed by applying discrete Fourier transforms in time to the data collected on the two sampling planes shown in Figure 2a. Figures 4a and 4b display the frequency at the peak amplitude in the Fourier spectrum for the coarse and fine grids. The highest frequency at the inception of the shear layer is associated with the primary instability. Even though the boundary layer at the beginning of the cavity opening is turbulent, a “secondary” transition process occurs when the shear layer forms [4, 5]. The coarse grid simulation fails at predicting the initial growth rate of the shear layer instability and its subsequent breakdown to turbulence. For the coarse grid simulation, the peak frequency is significantly lower and the peak occurs further downstream. Thus, the shear layer breakdown process is delayed. Nearly uniform low peak frequencies are observed everywhere inside the cavity. The acoustic near field of the shear layer shows larger peak frequencies for the coarse grid simulation in comparison to the fine grid results. The peak amplitudes in Figure 4c are largest at the shocks and in the shear layer region. The peak amplitude inside the cavity is an order of magnitude lower than in the shear layer but significant enough to cause small vibrations of telescope

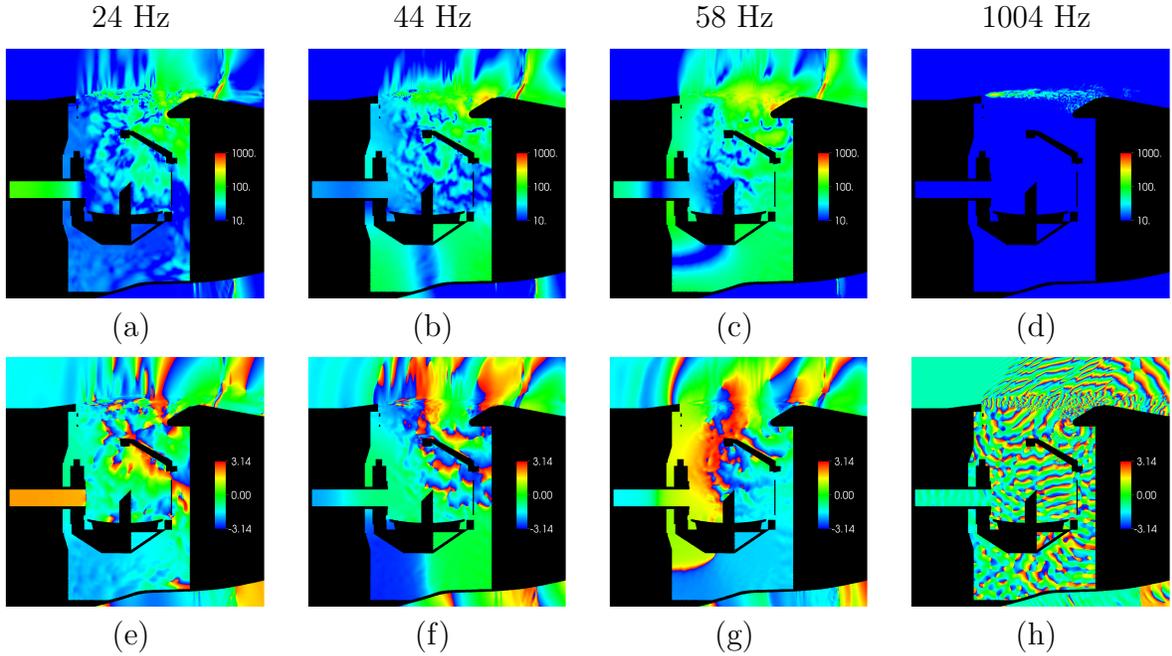


Figure 5: Phase analysis for pressure at different frequencies. Top row: spectral amplitude [Pa] at (a) 24 Hz (b) 44 Hz (c) 58 Hz (d) 1004 Hz; Bottom row: phase plots [rad] at (e) 24 Hz (f) 44 Hz (g) 58 Hz (h) 1004 Hz.

holding structure. Moreover, these pressure fluctuations may be close to a resonance frequency of the telescope structure.

Figure 5 presents temporal Fourier analysis results. Amplitude and phase plots in plane A-B are shown for selected frequencies. At 24Hz large pressure amplitudes can be observed in the Nasmyth tube. The phase plot indicates that the pressure oscillation is in phase throughout the tube. The standing wave with a frequency of 24Hz is perfectly tuned to the length of the tube. For all three lower frequencies (24Hz-58Hz), no wave-like phase pattern can be observed due to the large wave lengths at these frequencies. At 44Hz and 58Hz, large amplitudes can be observed inside the cavity. The phase plots are not straight forward to interpret since the cut planes only provide a 2D view of a 3D spectral mode. For 44Hz, there appears to be a phase shift from the lower left corner to the lower right corner. This pressure mode can induce a left-to-right rocking motion of the telescope holding structure that must be controlled with structural damping. The interaction of the shear layer with the aft aperture is picked up in the 44Hz and 58Hz modes. The phase plot for 44Hz also points towards a source location where the shear layer impinges onto the end of the cavity opening. The iso-contour lines of constant phase display a radial pattern away from the source location. At high frequency 1004Hz, the highest fluctuation amplitude is located in the shear layer. The frequency of 1004Hz was chosen since the most amplified initial disturbance growth was obtained around this particular frequency. Initially a very coherent pattern can be identified inside the shear layer by the planar wave fronts in the phase plot. The short wave-length is related to the convection

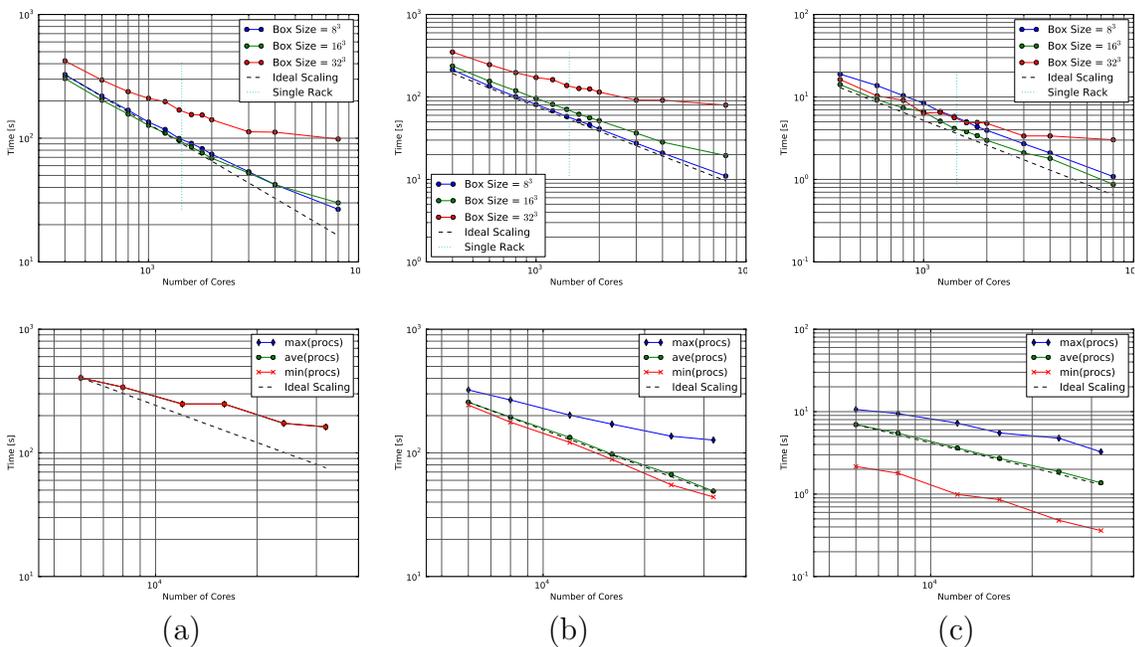


Figure 6: Strong parallel scaling study for 25 time steps: top row is ~ 300 million cells (fine mesh), max time over all procs shown; bottom row is 4.4 billion cells (super-fine mesh), min/ave/max time over all procs shown for box size of 32^3 . (a) Entire RK-solve including communication, (b) WENO-5Z flux and divergence operators, (c) immersed boundary conditions.

speed and high frequency. Various acoustic wave patterns are visible in the phase plots. The large unsteady coherent flow structures generate acoustic noise which propagates away from the shear layer. The sonic line separates the silent region upstream of the beginning of the cavity from the acoustic near field of the shear layer and the cavity flow. The interaction of the shear layer with aft aperture also generates some high frequency pressure oscillations. The acoustic wave-like phase pattern quickly loses amplitude as the waves are reflected and interfere with others inside the cavity.

3.3 Solver Parallel Performance

The analysis of the flow field in the previous section has highlighted the importance of providing sufficient grid resolution such that the wide range of physically relevant temporal and spatial scales can be resolved. Even with very efficient higher-order numerical schemes, significant resources are required to conduct such high fidelity CFD analysis. In the current section, we want to demonstrate two key capabilities of the LAVA solver: (1) scalability can be achieved for up to at least tens of thousands of processors and (2) the ability to dynamically adjust to the non-uniform per box timings of the immersed boundary method. To address the first objective, strong scaling studies were performed for both fine and super-fine meshes with the highly complex SOFIA geometry with multi-levels of resolution. All runs were conducted using NASA's Pleiades supercomputer. The Pleiades nodes used are composed of Intel Xeon E5-2680v2 processors (Ivy Bridge). The

Ivy Bridge memory hierarchy is as follows: L1 instruction cache: 32 KB, private to each core; L1 data cache: 32 KB, private to each core; L2 cache: 256 KB, private to each core; L3 cache: 25 MB, shared by 10 cores in each socket; and DDR3 FB-DIMMs main memory: 32 GB per socket, total of 64 GB per node. Sockets are connected with a 32 GB/s link. Each node (20 cores) was assigned to 20 MPI ranks, and internode communication is performed over InfiniBand, with all nodes connected in a partial 11D hypercube topology.

The fine mesh strong scaling study was performed using three different box sizes, 8^3 , 16^3 , and 32^3 and 7 levels of refinement, and 400-8000 cores. At 8^3 the tagging resulted in 392,461 total boxes or 200,940,032 cells, and 1,076,912,984 total cells including 3-ghost cells on all sides of the box (i.e. 14^3). At 16^3 the tagging resulted in 59,904 total boxes or 245,366,784 cells, and 637,857,792 total cells including 3-ghost cells on all sides of the box (i.e. 22^3). At 32^3 the tagging resulted in 10,362 total boxes or 339,542,016 cells, and 568,583,664 total cells including 3-ghost cells on all sides of the box (i.e. 38^3). At 32^3 there are not enough boxes / work for the large processor count.

The top row of Figure 6 presents the results of the strong scaling study on the fine mesh. Three components of the solver were characterized, (a) entire Runge-Kutta time integration step including communication, (b) WENO-5Z flux and divergence operators, and (c) immersed boundary treatment. The entire Runge-Kutta time-integration step includes four substeps for the fourth-order accurate time integrator. At each substep, the solver proceeds as follows: first, the convective fluxes are computed at each node utilizing Lax-Friedrichs flux vector splitting. The fluxes at the half points are then obtained via WENO reconstruction from the fluxes at nodes. The second-order viscous terms are evaluated in conservative form where we obtain the viscous fluxes at the half points employing classical finite difference evaluation. Finally, at each computational cell the discrete divergence operator is applied to the convective and viscous fluxes previously obtained at the half-points. After each step, block, coarse-fine interface, and domain ghost cells need to be filled. For coarse-fine grid interfaces an interpolation procedure is needed to fill the coarse and fine grid neighboring ghost cells. The WENO-5Z flux reconstruction procedure consists of three basic steps:(1) forward characteristic transform of convective flux, (2) computation of smoothness indicator and nonlinear weights for three candidate stencils, and (3) backward characteristic transform. The WENO-5Z flux reconstruction is applied to individual grid lines to keep the memory local in cache. The immersed boundary method is extremely powerful because it provides the capability of computing flows around highly complex geometries without significant manual mesh generation. From a programming perspective however, it is a non-uniform part of the code. The current implementation of the ghost cell method is first-order accurate at the wall. Each IB ghost-cell needs to be visited and an immersed boundary ghost cell procedure is applied. For higher-order accurate immersed-boundary schemes, as the one currently being developed in LAVA [1], the cost of the immersed boundary treatment increases. Hence, processors with immersed boundary ghost cells or irregular points incurs a larger overhead. In order to load balance the work on the processors, we utilize dynamic load

balancing where the block assignment per processors is re-evaluated after a number of time-steps based on recorded execution times for each box, where the boxes are on a lattice based uniform box-index-space per level.

The fine mesh scaling study shows that 8^3 boxes are preferred for the WENO-5Z flux and divergence operators (Figure 6b top row) due to locality in the memory hierarchy. This trend is not the same for the immersed boundary conditions (Figure 6c top row), where 16^3 is preferred. Note the 32^3 flat-line in Figure 6c top row, is due to not enough boxes/work available. Figure 6a top row, which shows the performance for all operations during 25 Runge-Kutta steps scales very well for the smaller sized boxes. No significant dependence was found when more than one rack was used (dotted blue line in Figures).

A super-fine mesh was generated by refining the shear layer region with an additional refinement level, and a strong scaling study was conducted using 6000-32k cores. For this super-fine mesh, smaller sized boxes were attempted (8^3 and 16^3), which would yield more boxes per level and improved load balancing, but memory was not sufficient due to the volume mesh meta-data [6]. Mesh meta-data compression and/or hybrid MPI/OpenMP (anticipated future work) should allow for smaller boxes, and therefore improved scaling. Boxes were fixed at 32^3 and the tagging resulted in 136,603 total boxes or 4,476,207,104 cells, and 7,495,679,816 total cells including 3-ghost cells on all sides of the box (i.e. 38^3). Figure 6 bottom row presents the strong scaling study for the 8 level mesh. Similar trends are observed as compared to the fine mesh study. Due to the multi-scale nature of the refinement hierarchy there are many more boxes on the finer levels for this application. The 8 level computations have the following number of boxes for each level: [8, 8, 12, 52, 474, 6623, 3660, 125766]. In the current algorithm, these boxes are distributed to processors per level, which for portions of the algorithm can lead to inter-level load imbalances, as seen in Figure 6 bottom row, when the number of boxes per level is non-uniform, e.g. 8 boxes on coarsest level vs 125766 on finest level.

On the fine mesh, another strong scaling study was performed illustrating the impact of timer based “dynamic” load balancing during adaptive mesh refinement as compared to uniform or “static” loads. Dynamic load balancing is based on timers for each box in the AMR hierarchy, weighting box loads by the accumulated time. This is possible due to the fixed-box-size box layout used in these simulations, where the boxes are stored on a lattice in a box index-space for each level. Dynamic load balance timings were recorded for all box operations. Static load balancing is based on a uniform weighting of the box loads. For this scaling study the simulation was run for 100 time steps, with re-gridding and load balancing occurring every 5 steps. Figure 7 shows an improvement for the immersed boundary condition routine due to dynamic loads. No significant improvement was observed for the remainder of the routines as the load is well balanced with naive uniform load balancing (see Figure 6), given a sufficient workload.

4 Conclusion

The paper presents a higher-order implicit large eddy simulation of a large telescope mounted in a large open cavity in the aft of a Boeing 747SP. Employing conventional

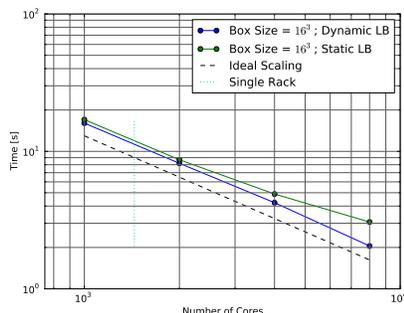


Figure 7: Strong parallel scaling study for 100 time steps, adaptively regridding every 5 steps, comparing dynamic and static load balancing schemes. Simulation is for fine mesh, with max time over all procs shown. Timings are for immersed boundary conditions routines only, using box size of 16^3 .

CFD approaches, significant time must be spent to generate a computational mesh which allows for high quality CFD results. The current approach which utilizes an automatically generated block-Cartesian data layout, allows for fast data access, and which promises fast execution times.

5 Acknowledgments

We gratefully acknowledge the SOFIA project for access to CAD models and conditions. C. Brehm supported under contract STC-NNA10DF26C with NASA.

REFERENCES

- [1] C. Brehm and H. Fasel. A novel concept for the design of immersed interface methods. *Journal of Computational Physics*, 242(0):234 – 267, 2013.
- [2] S. B. Cumming, M. Smith, L. Cliatt, M. Frederick, and F. by Institutions. Aerodynamic and acoustic flight test results for the stratospheric observatory for infrared astronomy. 2013.
- [3] C. Kiris, M. Barad, J. Housman, E. Sozer, C. Brehm, and S. Moini-Yekta. The LAVA Computational Fluid Dynamics Solver. In *AIAA Paper 2014-0070, Jan 13-17, National Harbor, Maryland*, pages 1–43, 2014. AIAA–2014–0070.
- [4] M. Shur, P. Spalart, and M. Strelets. Noise Prediction for Increasingly Complex Jets. Part 2: Applications. *International Journal of Aeroacoustics*, 4(4):247–266, 2005.
- [5] M. Shur, P. Spalart, and M. Strelets. Noise Prediction for Increasingly Complex Jets. Part I: Methods and Tests. *International Journal of Aeroacoustics*, 4(3):213–246, 2005.
- [6] B. Van Straalen, P. Colella, D. T. Graves, and N. Keen. Petascale block-structured AMR applications without distributed meta-data. In *Euro-Par 2011 Parallel Processing*, pages 377–386. Springer, 2011.