

Coupled RANS/LES for SOFIA Cavity Acoustic Prediction

Stephen Woodruff*

NASA Dryden Flight Research Center, Edwards, CA 93523-0273

A fast but accurate approach is described for the determination of the aero-acoustic properties of a large cavity at subsonic flight speeds. This approach employs a detached-eddy simulation model in the free-shear layer at the cavity opening and the surrounding boundary layer, but assumes inviscid flow in the cavity and in the far field. The reduced gridding requirements in the cavity, in particular, lead to dramatic improvements in the time required for the computation. Results of these computations are validated against wind-tunnel data. This approach will permit significantly more flight test points to be evaluated computationally in support of the Stratospheric Observatory For Infrared Astronomy flight-test program being carried out at NASA's Dryden Flight Research Center.

Nomenclature

SOFIA	Stratospheric Observatory For Infrared Astronomy
LES	Large-Eddy Simulation
RANS	Reynolds-Averaged Navier Stokes
DES	Detached-Eddy Simulation
PSD	Power Spectral Density
C_{DES}	DES coefficient modulating blending between RANS and LES
u	Velocity tangent to wall in turbulent boundary layer
u_∞	Velocity tangent to wall outside of boundary layer
y	Coordinate normal to wall in boundary layer
δ	Boundary-layer thickness
δ^*	Displacement boundary-layer thickness ($\delta^* = \int_0^\infty (1 - u/u_\infty) dy$)
θ	Momentum boundary-layer thickness ($\theta = \int_0^\infty (u/u_\infty)(1 - u/u_\infty) dy$)
n	Parameter in power-law boundary-layer profile $u(y) = u_\infty(y/\delta)^{1/n}$.

I. Introduction

An effort has been undertaken to provide timely computational cavity aero-acoustic support for the Stratospheric Observatory For Infrared Astronomy (SOFIA) flight-test program at NASA's Dryden Flight Research Center. This effort includes assessing cavity acoustics (including mitigation methods) at nominal and non-nominal conditions prior to flight tests, rapidly responding to issues raised during the flight-test program and using the computational results to assist with interpretation of the flight-test data. This is also a valuable opportunity to compare the results of aeroacoustic cavity computations using hybrid RANS/LES techniques with full-scale flight data, and to make improvements in the computational techniques based on those comparisons.

The SOFIA project is the latest in a series of aircraft-mounted optical observatories which require the fuselage be open to the external airflow in order to provide an unobstructed viewing field. Typical missions will involve cruising for extended periods at Mach 0.85 at around 40,000 ft. SOFIA is distinguished by the significantly larger size of the cavity opening than previous observatories and it is this feature that

*Aerospace Engineer; Current Address: Computational Aerosciences Branch, NASA Langley Research Center, MS 128, Hampton, VA, 23681. Non-Member AIAA.

makes computational assessments of its aero-acoustic behavior particularly important, as well as particularly difficult.

A full assessment of SOFIA cavity acoustics requires computations at many different flight conditions. In addition to the planned, nominal, flight conditions, it is necessary to get data at non-nominal flight conditions that might occur in emergencies: the possibility of the cavity door getting stuck at nominal altitude, for example, requires checking acoustics at flight conditions all the way down to landing. Performing this assessment requires a reasonably rapid computation of each case, as does responding in a timely manner to analysis needs during flight testing. The weeks such computations have typically taken in the past is too long and so a more rapid approach had to be developed. As described below, this has been accomplished through the use of a highly efficient structured-grid code and by focusing on the significant physical effects so that computational effort is not wasted on flow features not relevant to the desired results. Extensive wind-tunnel test results are available for detailed validation of this approach.¹⁻³

As the goal of the project is to understand the effect of cavity acoustics on flight safety and mission success, the project is fundamentally concerned with the unsteady aeroacoustics of the cavity and of the aircraft as a whole. Capturing these unsteady processes accurately requires accurate modeling of the unsteady motion of the turbulent shear layer at the cavity opening and thus a hybrid RANS/LES approach is called for: a RANS (Reynolds-Averaged Navier-Stokes) approach would not provide sufficient detail about the unsteady motion and a full LES (Large-Eddy Simulation) approach would be far too computationally intensive. Currently, the most commonly used hybrid RANS/LES approach is DES (Detached-Eddy Simulation), which seeks to resolve large-scale eddies in free shear layers explicitly, but relies on RANS modeling for smaller-scale boundary-layer turbulence. The DES technique of Strelets,⁴ based on the Menter SST two-equation turbulence model,⁵ is used in the present computations.

Cavity acoustics is an active research topic in a number of areas, including automotive design and aircraft internal weapons bays. The approach developed here should be of value in these areas as well, particularly when the cavity is deep and a significant portion may be treated inviscidly.



Figure 1. SOFIA telescope cavity. The red cover of the primary mirror may be seen near the bottom of the cavity; otherwise, the telescope structure occupies a minimal amount of the cavity volume. The large lip on the aft portion of the cavity doorway (in the right foreground of the picture) is designed to promote a more stable reattachment of the free shear layer over the cavity doorway.

II. Computational Approach and Grid Development

Typical computations of cavity acoustic problems have taken several days to several weeks on clusters of four to several dozen processors.⁶⁻⁸ The cavity on the SOFIA aircraft, shown in Fig. 1, is significantly larger than those of earlier studies, which increases the computational complexity of the problem by at least three or four orders of magnitude. The frequency of the lowest resonance of the SOFIA cavity, estimated at




Figure 2. Coarse mesh. (View from back of aircraft looking forward.)

2-3 Hz, is lower by a factor of ten than that of previous studies, requiring that computations be continued for at least ten times as long. Other modes scale similarly. Gridding the larger cavity requires one hundred to one thousand times as many cells, also increasing the computational time. These estimates are borne out in the SOFIA cavity computations of Schmid and colleagues.⁹⁻¹²

It would be impossible to meet the goals of the project if the computations took this long. Computations were thus planned at two levels. The first level involves a computational domain restricted to the cavity and its immediate vicinity, a simplified geometry, and flow modeling restricted to only what's necessary to capture the essential physics (DES-modeled turbulent flow in the shear layer, inviscid flow elsewhere). The second level consists of full-aircraft unsteady simulations with full turbulence physics. This approach provides for rapid computation of cases both before and during flight testing, as well as more detailed computations for validation and to address specific issues that may arise. The purpose of the present paper is to report on the development of the first-level approach and its validation with experiment.

Insights gained from SOFIA wind-tunnel studies and the computations of Schmid and colleagues⁹⁻¹² provide the basis for the simplifying assumptions employed in the present approach. First, it is noted that test results¹⁻³ show the dominant frequencies are below about 200 Hz. This implies that the relevant wavelengths are of the order of a meter and greater; thus, fine details of the cavity geometry are unlikely to affect the behavior of aero-acoustical phenomena at these frequencies. In particular, of the components of the telescope structure, only the primary mirror is likely to play any significant role. Experiments,¹ in which removal of the telescope from the model cavity resulted in only slightly higher noise levels, confirm this. In the present computations, the cavity geometry is simplified by the elimination of fine details and the telescope structure is entirely neglected.

Second, Schmid and colleagues^{9,12} have shown that the frequency and shape of vibrational modes within the cavity are predicted nearly as well by a linear acoustic wave computation as they are by DES computations. This indicates a full viscous solution with turbulence modeling in the cavity is unnecessary and computational time may be saved by solving only the Euler equations and (more significantly) by not having to grid the cavity to resolve shear layers.

As a consequence of these two insights, the geometrically simplified cavity and the farfield are gridded sparsely, to a degree adequate only for an Euler computation, and the larger proportion of computational effort and grid points are focused on the fuselage boundary layer and the free shear layer in the cavity




Figure 3. Close-up of cavity of coarse mesh. (View from left and front of aircraft looking back.)

doorway, where the full viscous/DES equations are solved.

II.A. Grids

The initial, coarse, grid employed in this investigation is shown in Fig. 2. The 60-degree door position, for which the cavity opening is at its highest position on the fuselage, has been chosen for the present computations. The flow domain extends from just aft of the wing root to just forward of the vertical tail and consists only of the left side of the aircraft and the surrounding flow, from the fuselage to a cylinder of radius about 17.5 m. The external portion of the flow is covered by a single block with 129 points in the axial (x) direction, 97 points in the azimuthal direction and 97 grid points in the radial direction. The axial and azimuthal points are evenly spaced, except for mild clustering near the edges of the cavity doorway. The radial points are distributed so that the cell center of the first cell is approximately 1 mm above the fuselage surface and the next cell is 20% larger. This block is split into two blocks at the 37th radial point in order to facilitate the mixed-equation solution described below.

The cavity doorway, including the lip in the aft portion, is gridded with a single block with 49 points in each direction. Clustering in the axial and azimuthal directions are dictated by the external block; some clustering in the radial direction is done to mitigate the abrupt change in cell size where this block fits against the external block.

The cavity itself is meshed with a block for the upper cavity (the half-cylinder nearest the cavity opening) and a block for the lower cavity (the half-cylinder farthest from the cavity opening). In keeping with the discussion above, the cavity is meshed fairly coarsely, with the upper cavity block containing $55 \times 49 \times 49$ points and the lower cavity block containing $59 \times 29 \times 29$ points. A detail of the cavity and doorway meshing is given in Fig. 3.

As discussed in the following section, computations conducted with this grid suggested that increased shear-layer resolution in the streamwise and cross-stream directions would be desirable. A finer grid was thus constructed in which the exterior block was increased to $155 \times 145 \times 97$ grid points, the cavity doorway block was increased to $73 \times 73 \times 73$ grid points and the cavity blocks were increased to $83 \times 73 \times 73$ grid points and $89 \times 43 \times 43$ grid points. The cavity blocks were refined solely to maintain a conforming structured grid; if one were willing to introduce a patch interface between the cavity blocks and the cavity doorway block,




Figure 4. Close-up of cavity of finer mesh. (View from left and front of aircraft looking back.)

fewer points could probably be used in the cavity blocks. The cavity and doorway meshing for this grid is shown in Fig. 4.

II.B. Boundary Conditions

Inviscid wall conditions are applied in the upper and lower cavity blocks, in keeping with the Euler-equation solution to be performed in these blocks. The third Euler block, the portion of the external block outside the fuselage boundary layer, has freestream boundary conditions fixed at the reference conditions applied on its inflow and far-field faces, symmetry conditions on the vertical faces and a subsonic outflow condition applied on the outflow face. (The results computed to date have given no reason to believe that the use of symmetry conditions on the vertical faces is problematic, though this was not specifically tested in the experiments described below.)

Adiabatic, viscous, wall boundary conditions with turbulence wall functions are used on the fuselage surface. Turbulence wall functions were not found to work well on the walls of the cavity doorway itself, especially on the lip at the aft of the doorway, most likely due to the decidedly non-thin-shear-layer behavior that occurs there. After some experimentation, it was found that inviscid (slip) wall conditions were perfectly adequate for the forward and side walls, but the lip on the aft wall required full adiabatic viscous wall conditions.

Symmetry and outflow boundary conditions for the portion of the external block adjacent to the fuselage, containing the fuselage boundary layer, are the same as for the outer portion. The inflow boundary condition requires specification of the upstream boundary-layer profile, which is described in the following subsection.

II.C. Inflow Boundary Layer Specification

The critical nature of the upstream boundary layer to cavity acoustics is well-known and often cited in the literature (e.g., Ref. 9). This makes specification of the inflow boundary condition particularly important in the present computation; the fact that the computational domain inflow boundary is so close to the leading edge of the cavity only exacerbates the problem.

Initially, the inflow boundary condition was constructed from Prandtl's $u/u_\infty = (y/\delta)^{1/4}$ profile for a turbulent boundary layer. Turbulence model variables were set to constant values. Experiments with




Figure 5. PSD at forward bulkhead, coarse-fine grid comparison: 0.85 Mach, 40,000 ft. (Wind tunnel, red solid line; coarse-grid computation, green long dashes; fine-grid computation, blue short dashes.)

varying the boundary-layer thickness and turbulence levels at the upstream boundary were conducted to get some understanding of the sensitivity of the cavity acoustics to these factors. Variation of pressure amplitudes with boundary-layer thickness has been reported in the literature; it was found in the present experiments that the overall pressure amplitude could be controlled pretty much at will if one has sufficient control of the upstream boundary-layer thickness. Turbulence levels had an interestingly targeted effect: increased boundary-layer turbulence largely eliminated some resonance peaks in the pressure spectrum, but left others pretty much unchanged. (This effect has been observed in experiments with upstream boundary-layer modification devices.¹)

To determine the effect of more realistic inflow boundary-layer specifications, experiments were conducted in which a steady-state, two-dimensional, RANS flat-plate boundary computation was used to provide the inflow profile for the unsteady DES computation. The flat-plate boundary layer was allowed to develop for a streamwise distance approximately equal to the distance along the fuselage from the nose to the inflow boundary of the DES computational domain; the resulting two-dimensional density, velocity, energy and turbulence-variable profiles were mapped onto the inflow boundary of the DES grid. Angle of attack, free-stream turbulence levels and the angle of the upper boundary of the computational domain (to impose an adverse pressure gradient) were all used to provide different boundary-layer properties. This arrangement worked quite well and is not that much more complicated than the earlier approach once the transformation from two-dimensional to three-dimensional input file formats is set up. It is used in most of the fine-grid computations presented here.

Ultimately, it is expected that one would want to use a steady-state RANS computation of the full aircraft to provide the flow variables and turbulence quantities for the inflow boundary condition for the unsteady calculation. In the present situation, for example, the cavity opening has been placed right at the point on the fuselage where significant streamwise and azimuthal variations in the boundary layer take place due




Figure 6. PSD at forward bulkhead, effect of upstream boundary-layer profile on coarse-grid computation: 0.85 Mach, 40,000 ft. (Wind tunnel, red solid line; computation with $n = 5$, $\delta = 0.49$ m, boundary layer, green long dashes; computation with $n = 5$, $\delta = 0.95$ m, boundary layer, blue short dashes; computation with $n = 9$, $\delta = 0.95$ m, boundary layer, purple dots.)

to interactions with the tail planes. (This is also an argument for moving the DES grid inflow boundary upstream of its current location to provide more room for the boundary layer to sort itself out before hitting the cavity opening.) The simpler methods have been tried in order to begin to establish how simple one may go and still get reasonable results with the computational approach being developed here.

II.D. Code and Related Computational Matters

The VULCAN code¹³ for structured grids was chosen for this computation due to its excellent performance in other projects and the fact that a number of DES models are available. It also allows the viscous and turbulence stresses to be turned on or off in individual grid blocks, permitting the computation to be set up so that the Euler equations are solved everywhere but in the fuselage boundary layer and the free shear layer in the cavity opening, where the full viscous/DES equations are solved.

VULCAN provides several upwinding options and, given the sensitivity of LES acoustic problems to dissipation, experiments were carried out with different options to determine which gave the most accurate results. Surprisingly, the Roe scheme¹⁴ gave acceptable results, in spite of its relatively high dissipation. The low-Mach-number preconditioner¹³ was looked at carefully, in view of the low speeds in the cavity, and this option gave good results, but, in general, the best results were obtained with the Edwards low-dissipation LDFSS scheme.¹⁵ The smooth limiter¹⁶ option was used.

The computations are started by determining a steady-state solution, which is then used as the initial condition for the time-dependent solution. A dual time step time-advancement algorithm is employed, with a time step of 10^{-4} seconds. The residual was typically reduced by one order of magnitude in six or seven subiterations. This is a coarse time step, and one would generally prefer to see the residual reduced by at




Figure 7. (a) PSD at forward bulkhead, effect of upstream boundary-layer profile on fine-grid computation: 0.85 Mach, 40,000 ft. (Wind tunnel, red solid line; computation with flat-plate 0.45 m boundary layer, green long dashes; computation with flat-plate 0.70 m boundary layer, blue short dashes; computation with flat-plate 1.0 m boundary layer, purple dots.)

least two orders of magnitude, but experiments indicated that decreasing the time step and increasing the number of subiterations had only small effects on the power spectral density. Sticking with the coarse time step and smaller numbers of subiterations permitted more experiments to be conducted.

The DES scheme of Strelets⁴ is available in VULCAN and was employed in these computations. A value of $C_{DES} = 0.61$ was used for the constant modulating the cross-over from RANS behavior to LES behavior; this is the value determined by Strelets to give the correct behavior in a homogeneous-flow test case. (Here and elsewhere in this paper, C_{DES} refers to the parameter introduced by Strelets, not its reciprocal employed in the VULCAN implementation.)

In these computations, the coarse grid was run on 64 processors and a full computation, from determining the steady-state solution used as the initial condition, through evolution of the time-dependent solution to a statistically steady state and ending with at least one second of the statistically steady state for computing spectra, takes about 40 hours of wall-clock time, or about 2,500 cpu-hrs. An attempt to decompose the grid in a manner suitable for 128 processors was abandoned when the decomposition process split the grid into computationally problematic blocks only a few cells thick. This could no doubt be remedied, but long-term use of this coarse grid is not planned and reaching an absolute minimum in wall-clock time was not a priority. The finer grid was run on 128 processors and a full computation took about 34 hours of wall-clock time, or about 4,300 cpu-hrs.

Parallelization of the computation is handled by VULCAN, a well-established, highly developed, parallel research code. Computations with the coarse grid, running on 64 processors, consistently maintained a computational efficiency of 96% – 98%. The finer grid, running on 128 processors, generally maintained the same level of computational efficiency.




Figure 7. (b) PSD at forward bulkhead, effect of upstream boundary-layer profile on fine-grid computation: 0.85 Mach, 40,000 ft. (Wind tunnel, red solid line; computation with $n = 5$, $\delta = 0.95$ m, boundary layer, green long dashes.)

III. Results

The results of the present computations are described and discussed in the following subsections, with the twin goals of assessing the effectiveness of the computational approach outlined in this paper and of better understanding the dynamics of the SOFIA cavity aero-acoustics. The present approach is first validated by comparison with wind-tunnel test data, then the nature of its results is explored through examination of acoustic spectra at additional Mach-number conditions, and through examination of body-force spectra, streamwise vorticity and pressure contours at two qualitatively distinct resonance conditions.

III.A. Comparison with wind-tunnel tests.

Wind tunnel tests have been conducted on a number of occasions on the SOFIA configuration (including Refs. 1–3). While the focus of much of the testing was on the effect of the cavity on steady-state aerodynamic force properties, time-dependent acoustic data was also taken in several of the test programs at a variety of points in the cavity. As these provide the most detailed experimental information available about the time-dependent aspects of the cavity flow, this data is the primary means of validation of the present computations.

The present computations yield results similar to the wind-tunnel tests of Ref. 3, while the computations of Schmid and colleagues^{9–12} yield results similar to those of the wind-tunnel tests of Refs. 1, 2. The discrepancy between these two sets of computational and experimental results is as yet unexplained and will be returned to briefly at the end of this section.

The wind tunnel tests of Ref. 3 provide data for a 60-degree door angle at Mach numbers of 0.60, 0.80, 0.85 and 0.90; all of these cases were computed and the results are compared here with the wind-tunnel data for validation.




Figure 8. Upstream boundary-layer profiles: 0.85 Mach, 40,000 ft. (0.45 m boundary layer, red line; 0.70 m boundary layer, blue short dashes; 1.0 m boundary layer, purple dots; $n = 5$, $\delta = 0.95$ m, boundary layer, green long dashes.)

The case of Mach 0.85 and 40,000 ft., which approximates a cruise condition for a SOFIA mission, is employed here as a basic reference condition. The power spectral density at a point on the forward bulkhead about midway between the center and the doorway is shown in Fig. 5 for the wind-tunnel test and for computations using the coarse and fine grids. The frequencies of the first two peaks are adequately predicted by the coarse grid computation, but their amplitudes are not, with the first peak too small and the second peak too large. Extensive numerical experimentation, some of which is described below, revealed that changes in outer boundary conditions and other aspects of the numerical implementation had minimal effect on these results; the inflow boundary-layer specification did affect the peak amplitudes but do not change this basic behavior. The correct prediction of the frequencies indicates the acoustic modes in the cavity itself are well resolved, but the consistent mischaracterization of the amplitudes suggests the free shear layer, whose interaction with the cavity modes determines their amplitude, is not adequately resolved.

The fine grid is better, predicting the frequencies of the peaks accurately and predicting the amplitude of the dominant peak very closely as well. The increasingly inadequate attenuation of the dominant peak's harmonics as one goes up in frequency is not unexpected, since the higher frequencies and smaller length scales are less accurately resolved by any given grid and this situation is improved, but not cured, by making the grid finer. (An even finer grid was constructed and tested, but exhibited convergence problems which could not be addressed due to time constraints.)

It's also worth noting that the computed spectrum peaks are somewhat wider than the experimental peaks. Going from the coarse to the fine grid had little effect on the width of the computed peaks, indicating this is not a grid-resolution issue. It could be a time-resolution issue, however, arising from the much longer time-series records available from the experiments for computing spectra.

In Figure 5 and the other spectra shown in this paper, the range of frequencies from 10 Hz to 230 Hz was chosen to capture the dominant features of the spectra. As mentioned above, experiments and previous




Figure 9. PSD at forward bulkhead, effect of hybrid RANS/LES parameter C_{DES} : 0.85 Mach, 40,000 ft. (Wind tunnel, red line; computation with $C_{DES} = 0.61$, green long dashes; computation with $C_{DES} = 0.71$, blue short dashes.)

computations show the dominant acoustic modes have frequencies below about 200 Hz. While there are modes with frequencies below 10 Hz, the frequency resolution of the computational PSD's (1 Hz due to the one-second length of the time series), was found to be too coarse to identify these features with certainty. (Certainly longer records could have been produced with the present approach, and with less cost than would otherwise be possible. It was simply decided that the available resources could best be used on more experiments with shorter records.)

The computational PSD's are all computed using pressure time series extracted from the computation at 1 ms intervals. The length of each time series is one second. The experimental data,³ which is composed of multiple records much longer than could feasibly be obtained from the computations, was windowed and averaged; it was not clear the much shorter computational records would be amenable to these techniques and they have not been used. The one-second computational records were simply Fourier-transformed and the amplitude spectra computed.

The experimental data in these plots has been scaled to the flight conditions in a manner that seems to have become standard (see, for example, Refs. 1,2). Frequencies are scaled by the ratio of wind-tunnel to flight length scales (0.03 in this case) and the ratio of flight sound speed to wind-tunnel sound speed. Spectrum amplitudes (whose units are Psi^2/Hz) are scaled by the square of the ratio of flight dynamic pressure to wind-tunnel dynamic pressure and divided by the frequency scaling factor.

Numerical experiments with reduced time steps and increased subiterations per time step were conducted to check that the pressure spectra peaks were essentially unchanged. Use of the full viscous and turbulence equations in the entire flow domain was also investigated, with little change in the pressure spectra. (These latter computations were with the same grid as before, with inviscid (flow-tangency) boundary conditions at the walls of those domains formerly computed using the Euler equations, so wall shear layers were still not being resolved.)




Figure 10. PSD at forward bulkhead, fine-grid computation with two upstream profiles: 0.80 Mach, 40,000 ft. (Wind tunnel, red line; computation with 0.45 m boundary layer, green long dashes; computation with $n = 5$, $\delta = 0.95$ m, boundary layer, blue short dashes.)

Extensive numerical experimentation was performed to examine the effect of the upstream boundary-layer specification with the coarse grid and, to a lesser extent, with the fine grid. Fig. 6 illustrates some of the main conclusions drawn from these experiments. In this figure, coarse-grid computations for three different upstream boundary-layer profiles are compared with the wind-tunnel experiments. The velocity profiles are defined by the Prandtl power-law described in the previous section and the turbulence variables are given constant values that correspond to free-stream turbulence levels.

The original coarse-grid computation shown in Fig. 5 was computed using a boundary-layer profile defined by $n = 5$ and $\delta = 0.38$ m, corresponding to the 15-inch boundary layer measured on the full-sized SOFIA aircraft. Fig. 6 shows the PSD for a computation with the slightly thicker profile $\delta = 0.49$ m. Also shown are PSD's for profiles defined by $n = 5$, $\delta = 0.95$ m and $n = 9$, $\delta = 0.95$ m. Holding n constant and increasing the boundary-layer thickness is seen to decrease the overall amplitude of the peaks: the thicker free-shear layer that results is more effective at damping the cavity resonances. An equally thick boundary layer, but with an $n = 9$ profile, gives a PSD very similar to that of the thin boundary layer. The thinner $n = 5$ boundary layer and the $n = 9$ boundary layer have similar displacement thicknesses (0.082 m and 0.095 m, respectively) and momentum thicknesses (0.059 m and 0.078 m, respectively), in contrast to those of the thicker $n = 5$ boundary layer ($\delta^* = 0.16$ m, $\theta = 0.11$ m), emphasizing the utility of these parameters in correlating upstream boundary layer properties with the acoustic properties of the cavity.

These conclusions were carried over to the fine-grid computations, most of which were conducted with upstream boundary-layer specifications provided by the two-dimensional flat-plate computations described in the previous section. Fig. 7 (a) shows the results of a series of fine-grid computations conducted with progressively thicker upstream boundary layers, whose velocity profiles may be seen in Fig. 8. The 0.45 m boundary layer, from the two-dimensional flat-plate computation, was used in the original fine-grid compu-




Figure 11. PSD at forward bulkhead, fine-grid computation with two upstream profiles: 0.90 Mach, 42,000 ft. (Wind tunnel, red line; computation with 0.45 m boundary layer, green long dashes; computation with $n = 5$, $\delta = 0.95$ m, boundary layer, blue short dashes.)

tation shown in Fig. 5 and corresponds to the thin coarse-grid boundary layer of Fig. 6. The 0.70 m and 1.0 m boundary layers also come from flat-plate computations, in which the streamwise extent was increased and the upper boundary was angled outward to thicken the boundary layer and produce profiles intended to be more characteristic of the adverse-pressure-gradient conditions on the fuselage forward of the cavity.

The progressively thicker upstream profiles yield progressively improved PSD predictions. All the profiles yield approximately the same frequency and amplitude for the dominant peak in the spectrum. The two thinnest profiles give similar results for the second peak, but the thicker profile is somewhat better. All three flat-plate profiles give similar over-predictions of the third peak's amplitude; it is necessary to go to the power-law $n = 5$, $\delta = 0.95$ m profile (Fig. 7 (b)) to predict this peak correctly. The sensitivity of the PSD predictions to the upstream profile seems to increase with frequency: the lowest-frequency peak is fairly robust to changes in the boundary-layer profile, but the higher-frequency peaks are progressively more sensitive.

Fig. 8 suggests the failure of the flat-plate profiles to do better is due to inadequate mass and momentum defects; the largest, the 1.0 m profile, reaches $\delta^* = 0.10$ m and $\theta = 0.082$ m, which are still well below the $\delta^* = 0.16$ m, $\theta = 0.11$ m, values of the power-law profile.

The upstream profiles required to yield good agreement with the wind-tunnel PSD's for the higher-frequency peaks compare unfavorably with the upstream profiles measured in the wind tunnel^{2,3} or in flight,² which are comparable to the thinner profiles used in the present computations. This is at least partly due to the absence of the telescope in the computations. Experiments¹ show removal of the telescope yields slightly higher sound levels, and a follow-on study to this one does indicate the effect of obstacles in the cavity can be to attenuate harmonics. (The shorter wavelengths of the higher-frequency waves are more influenced by telescope-sized obstacles.)




Figure 12. PSD at forward bulkhead, coarse-fine grid comparison: 0.60 Mach, 25,000 ft. (Wind tunnel (60-degree door), red line; wind tunnel (20-degree door), green long dashes; coarse-grid computation, blue short dashes; fine-grid computation, purple dots.)

Aside from these considerations, there are indications in Fig. 7 that the use of the flat-plate computation to provide a more realistic upstream profile brings out some of the finer details in the PSD that were not present when the power-law profile and constant turbulence levels were used. The clearest example of this is the small side peak just to the left of the dominant peak: none of the power-law runs show evidence of this peak, but all of the runs using the flat-plate profiles do. On the other hand, a peak at 36-37 Hz is clearly defined in the power-law-profile result of Fig. 7 (b), but is, at best, less well-defined in the flat-plate-profile results of Fig. 7 (a). (The thickest of these profiles does the best.) This peak, also present in the experimental results, will be seen below to be the remnants of a low-speed resonant state.

The sensitivity of the results to the hybrid RANS/LES model constant C_{DES} was briefly evaluated. Strelets⁴ introduced this single additional constant in developing the DES version of Menter's SST turbulence model. (All other constants are the standard model values as implemented in VULCAN.) A homogenous flow calculation was used by Strelets to calibrate this constant to the value 0.61 and this value has been used in the computations discussed here. Fig. 9 shows what happens when the Mach 0.85 case is recomputed with C_{DES} increased to 0.71. This fairly large change (over 15%) does have a noticeable quantitative effect on the results, particularly at higher frequencies, but overall, a change in C_{DES} of this magnitude does not dramatically alter the predictions of the computation. Further evaluation of the effect of this parameter would be advisable, though, in order to get the best possible results from this approach.

Additional experimental data are available at Mach 0.80 and Mach 0.90. Corresponding flight conditions chosen for the computations were Mach 0.80 at 40,000 ft and Mach 0.90 at 42,000 ft. The Mach 0.80 case is shown in Fig. 10 and the Mach 0.90 case in Fig. 11. Computations were performed using both the 0.45 m upstream boundary-layer specification from the two-dimensional flat-plate computation (recomputed for the new Mach numbers and altitudes) and the $n = 5$, $\delta = 0.95$ m, power-law upstream boundary-




Figure 13. PSD at forward bulkhead, fine-grid computation with two upstream profiles: 0.60 Mach, 25,000 ft. (Wind tunnel (60-degree door), red line; wind tunnel (20-degree door), green long dashes; computation with 0.45 m boundary layer, blue short dashes; computation with $n = 5$, $\delta = 0.95$ m, boundary layer, purple dots.)

layer specification. The results for these conditions are quite similar to the Mach 0.85 case, as would be expected from the similar Mach number, and the agreement with experiment is similarly good. Coarse-grid computations were not made for these cases. The excessive amplitude of the higher-frequency peaks is becoming more noticable at 0.90 Mach, suggesting further refinement of the grid, the upstream boundary-layer, or both is required as the Mach number is increased.

Experimental data are also available for Mach 0.60 and are shown in Fig. 12. Unlike the test conditions discussed above, data is available for the 60-degree door angle at Mach 0.60 only with a boundary-layer rake mounted just upstream of the cavity opening. (Wind-tunnel measurements at the test conditions discussed previously were all done both with and without the rake. The measurements without the rake were used in the above comparisons.) Measurements at Mach 0.60 are available³ (without the rake) for the door at 20 degrees (meaning the cavity is rotated so the opening is oriented closer to the horizontal) and these are also shown in Fig. 12. A coarse-grid computation using the $n = 5$, $\delta = 0.38$ m upstream profile and a fine-grid computation using a similar 0.45 m flat-plate boundary layer are seen in the figure to give results very similar to the 20-degree door-angle measurements, with the fine grid predicting the peak amplitudes somewhat more accurately.

The similarity between the 20-degree door-angle measurements and the 60-degree door-angle computations is not so odd as it might at first seem. The restricted computational domain employed in the present approach ensures that the dominant means by which the overall flow influences the aero-acoustics of the cavity is through the upstream boundary-layer profile. The effect of the local surface geometry is secondary. The primary difference between a 60-degree door computation and a 20-degree door computation is thus not the problem geometry, but the upstream boundary-layer specification. The computations shown in Fig. 12 employ upstream profiles that are consistent with those seen when the door is in the lower position, so their




Figure 14. Computed PSD's at forward bulkhead, effect of increasing Mach number. (Mach 0.60, red line; Mach 0.70, green long dashes; Mach 0.76, blue short dashes; Mach 0.80, purple dots.)

PSD's agree with measurements taken at that condition. On the other hand, the PSD measurements taken with the rake in place are significantly influenced by the rake and its mount, which occupies one corner of the cavity doorway: comparison of measurements at other Mach numbers made with and without the rake show this clearly. The rake and its mount were not included in the computations, so the computational results do not reflect this influence.

The effect of upstream boundary-layer thickness on fine-grid computations at this flight condition is shown in Fig. 13. The 0.45 m flat-plate profile from Fig. 12 is repeated here for convenience so that it may be contrasted with the $n = 5$, $\delta = 0.95$ m power-law profile. The effect of these profiles on the computed PSD's are similar to those seen for the Mach 0.85 case in that the dominant peak is predicted almost equally well by both upstream profiles. In contrast to the Mach 0.85 case, the thicker profile increases the amplitudes of the peaks, particularly those of the first three harmonics.

That this case differs fundamentally from the three higher-speed cases discussed above is also indicated by the difference in frequency of the dominant peak. The high-speed cases all had dominant peaks at approximately 51–53 Hz; the present case has its dominant peak at 36–37 Hz. The wind-tunnel test reports do describe qualitatively different low-speed and high-speed resonance conditions and it appears the present computations have captured the low-speed resonance at the 0.60 Mach flight condition. The wind-tunnel measurements indicate that this is correct for a 20-degree door angle. In the discussion of additional results that follows, an attempt will be made to draw further conclusions about these different resonance conditions.

III.B. Additional cases between Mach 0.60 and Mach 0.80

In order to explore further the nature of the different resonance states captured in the Mach 0.60 computation and in the other computations, and the transition between them, cases were set up and computed for Mach




Figure 15. Computed force spectra on fuselage segment: 0.60 Mach, 25,000 ft. (Longitudinal force, red line; lateral force, green long dashes; vertical force, blue short dashes.)

numbers of 0.65, 0.70 and 0.76. The transition was found to occur between Mach 0.70 and 0.76, as shown in Fig. 14. The Mach 0.70 power spectrum is quite similar to that for the Mach 0.60 case (as is the Mach 0.65 spectrum, which is not shown) and the Mach 0.76 spectrum is quite similar to that for the Mach 0.80 case. There is a small peak in the Mach 0.76 spectrum at 36 – 37 Hz, indicating the lower-speed resonance is still present, but with a much smaller amplitude. A similar, but smaller, peak is visible in the Mach 0.80 case as well.

These observations are consistent with the view that a simple, low-speed resonance condition gives way, at some Mach number between 0.70 and 0.76, to a high-speed resonance condition and that there is a small range of Mach numbers where both exist simultaneously.

III.C. Force spectra

Examination of the unsteady forces exerted by this flow on the fuselage is valuable in order both to assess the possible impact on structural and flight dynamics and to identify differences between the resonance state computed at Mach 0.60 and at the other conditions. Figs. 15 and 16 show, for the Mach 0.60 and Mach 0.85 cases, the spectra for the longitudinal, lateral and vertical forces integrated over the fuselage section of the present computation. These forces are normalized by the amplitude of the time-averaged integrated-force vector.

The normalized peak magnitudes are less than 10^{-4} . This is reassuringly small, but it must be noted that local unsteady surface pressures could be significantly higher. It is also worth noting that when each integrated-force component is normalized by its own mean (rather than the three-dimensional force-vector amplitude used in the figures), the longitudinal and lateral forces reach amplitudes on the order of half a percent. If the unsteady amplitudes remain at such levels when the forces on the full fuselage are taken into




Figure 16. Computed force spectra on fuselage segment: 0.85 Mach, 40,000 ft. (Longitudinal force, red line; lateral force, green long dashes; vertical force, blue short dashes.)

account, it may be necessary to evaluate potential effects on flight dynamics.

The force spectra serve the further purpose of breaking the resonance peaks into three directional components, which allows more to be learned about the two resonance states present in the computations. Figs. 15 and 16 are distinctly different, in that the Mach 0.60 resonance is dominated by the longitudinal component and the Mach 0.85 resonance is much less so. In fact, the transverse components (lateral and vertical) dominate the longitudinal at nearly every peak except the main one. This clearly indicates that a much more complex unsteady mode is present at the higher Mach-number condition. Small peaks in the Mach 0.85 spectrum at the location of the main peak in the Mach 0.60 spectrum indicate the low-speed resonance condition is present to a small degree at the higher speed. (There is evidence of this peak in the pressure spectrum of Fig. 5, as well.) The converse is not true, suggesting that the more complex high-speed resonance mode is not present to any significant degree at low speeds. Note also the presence of an almost purely transverse mode at approximately 30 Hz in both Figs. 15 and 16; the longitudinally oriented pressure sensors from which the pressure spectra are derived pick up little, if any, evidence of this mode.

III.D. Streamwise vorticity in the free shear layer.

While the results shown so far describe the overall behavior of the resonance conditions through spectra of pressures and integrated forces, it is also valuable to look at the behavior of the free shear layer above the cavity opening since it plays a crucial role in determining which of the many possible acoustic modes will be attenuated and which will be amplified. It is by influencing this shear layer that the upstream boundary layer plays such a significant role in this problem.

Examination of the streamwise vorticity in the shear layer shows that the Mach 0.60 and Mach 0.85 cases are distinguished by significantly different levels of three-dimensionality. Fig. 17 shows streamwise




Figure 17. Streamwise vorticity contours in a $y - z$ plane located halfway between the leading and trailing edges of the cavity opening: 0.60 Mach, 25,000 ft. (Red indicates positive vorticity, blue indicates negative vorticity and the x axis is into the page.)




Figure 18. Streamwise vorticity contours in a $y - z$ plane located halfway between the leading and trailing edges of the cavity opening: 0.85 Mach, 40,000 ft. (Red indicates positive vorticity, blue indicates negative vorticity and the x axis is into the page.)

vorticity contours at Mach 0.60 in the $y - z$ plane at a point midway between the leading and trailing edges of the cavity opening. While there is vorticity in the cavity doorway and in the cavity itself, relatively little vorticity penetrates into the shear layer directly above the cavity doorway. The shear layer is predominantly two-dimensional and streamwise vortices have not formed in the layer to a significant extent. This shear-layer behavior is consistent with the relatively simple longitudinal cavity resonant state indicated by the pressure and force spectra.

In contrast, Fig. 18 shows vorticity contours for the Mach 0.85 case. Streamwise vortices are present throughout the free shear layer and are as intense as the eddies in the cavity itself. In this case, the shear layer is fully three-dimensional, with correspondingly more complex behavior that clearly influences the more complex resonant state in the cavity.




Figure 19. Cavity pressure isosurfaces, low-speed resonant condition: 0.60 Mach, 25,000 ft. (View is from side of cavity, rotated so that the cavity door is straight up. Freestream flow is from left to right.)

III.E. Pressure modes in the cavity.

Finally, pressure isosurfaces in the cavity are examined to identify the shape of the acoustic modes associated with each of the two resonance states observed in the present computations. Fig. 19 shows instantaneous pressure isosurfaces at the Mach 0.60 flight condition. The surfaces are parallel flat disks oriented normal to the longitudinal axis of the aircraft. This indicates the acoustic mode is a one-dimensional longitudinal wave oriented in the direction of the x axis: the x component of velocity has nodes at the fore and aft bulkheads and the gradient of the associated pressure wave has zeros at the bulkheads, which explains why the equally-spaced isosurfaces are blurred there.

The more complex acoustic mode of the high-speed resonance is shown (for the Mach 0.85 flight condition) in Fig. 20. This is a two-dimensional mode, in which each of the four quadrants are 180 degrees out of phase with its two neighbors. The view from the aft bulkhead, in Fig. 21, shows the essential two-dimensionality of this mode is destroyed only when the corner geometry of the cavity walls forces three-dimensional motion. In addition to its distinctly different shape, this mode is distinguished from the low-speed resonance mode by the streamwise variation its structure imposes on the shear layer in the doorway, which is consistent with the enhanced three-dimensionality of the shear layer in the high-speed resonance state discussed in the previous subsection.

These pressure isosurfaces seem to be consistent with the 35 Hz and 45 Hz acoustic mode shapes presented by Schmid and colleagues.^{9,12} They established that these acoustic modes were associated with the resonances at approximately 35 Hz and 47 Hz seen in their computations and in the wind-tunnel experiments of Refs. 1, 2. While the 35 Hz mode (identified here as the 36–37 Hz longitudinal, low-speed, resonance mode) is similar in all four investigations, the 47 Hz resonance is replaced in the present paper and in the wind-tunnel experiments of Ref. 3 by a 51–53 Hz resonance. The similarity in the acoustic mode shapes of the 47 Hz resonance and the 51–53 Hz resonance provides evidence that these resonances are, at least in part, physically similar; what it is that causes their differing frequencies and amplitudes is still unclear. (The side peak referred to in the discussion of Fig. 7(a) appears to be evidence of the 47 Hz resonance, so it may be more accurate to say the results of Ref. 3 and the present computations have a strong 51 Hz peak and a weak 47 Hz peak.)




Figure 20. Cavity pressure isosurfaces, high-speed resonant condition: 0.85 Mach, 40,000 ft. (View is from side of cavity, rotated so that the cavity door is straight up. Freestream flow is from left to right.)




Figure 21. Cavity pressure isosurfaces, high-speed resonant condition: 0.85 Mach, 40,000 ft. (View is from aft bulkhead looking forward, rotated so that the cavity door is straight up.)

IV. Conclusions

The finer-mesh computations took approximately 34 hours of wall-clock time on 128 processors. The coarse-grid computation took approximately 40 hours of wall-clock time on 64 processors. These times are significantly faster than those cited for previous similar cavity computations. It has thus been demonstrated that a computation with simplified geometry and physics can capture the essential cavity resonance phenomena accurately enough to provide useful information much more rapidly than has previously been possible.

This ability to perform many computations rapidly has been employed to help clarify the relationship between the upstream boundary layer and the aero-acoustic properties of the cavity. Not unexpectedly, integral measures of boundary-layer thickness (the displacement and momentum thicknesses) gave a better indication of the effect a given profile would have on the cavity PSD than the nominal thickness δ . Robustness of the dominant peak to changes in the upstream profile was observed (when grid resolution is sufficient), as was the fact that the amplitudes of the higher-frequency peaks decrease with increasing profile thickness in the high-speed resonant state. The reverse was observed in the low-speed resonant state.

The nature of the low- and high-speed resonance states was investigated by examining the evolution of the PSD as the Mach number is increased. Force spectra showed that the low-speed resonance is a simple longitudinal disturbance, but the high-speed resonance is a more complex mode with significant transverse motion. The high-speed resonance condition is also associated with significant streamwise vorticity in the shear layer above the cavity opening, in contrast to the more two-dimensional shear layer seen in the low-speed case. Pressure isosurfaces in the cavity showed the low-speed resonance is characterized by a one-dimensional longitudinal acoustic mode, while the high-speed resonance is associated with a more complex two-dimensional mode.

While it is clear that particular care must be taken in developing correct inflow boundary layer specifications, the present results show that the approach outlined in this paper for fast, accurate, computations of cavity acoustic problems is a feasible means for making runs at many more flight conditions than would otherwise be possible. When used knowledgeably and in conjunction with experimental testing and full-aircraft unsteady computations for validation, this approach can provide a fuller and more detailed understanding of acoustic and aerodynamic conditions throughout the flight envelope.

Acknowledgments

Comments by T. Bui, J.A. White and J.H. Morrison on this paper are gratefully acknowledged, as are helpful discussions with T. Bui throughout the period during which this work was carried out. The assistance of the SOFIA project, particularly aerodynamic lead S. Cumming, in providing information and data is also gratefully acknowledged, as is a useful discussion with S. Schmid. Computer time was provided by the NASA Advanced Supercomputing Division.

References

- ¹Rose Engineering and Research, Inc., "SOFIA — Conceptual Door Design and Aero-Optics Wind Tunnel Test: SOFIA IV: Final Report," Technical Report, 1996.
- ²Rose Engineering and Research, Inc., "SOFIA V Design Validation Test Final Report," Technical Report, 1998.
- ³Raytheon Aircraft Integration Systems, "Aeroacoustics Wind Tunnel Test Report for the Stratospheric Observatory for Infrared Astronomy (SOFIA) Program," Technical Report, 2000.
- ⁴Strelets, M., "Detached Eddy Simulation of Massively Separated Flows," AIAA 2001-0879.
- ⁵Menter, F. R., "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA J.* Vol. 32, No. 8, 1994, pp. 1598-1605.
- ⁶Mendonca, F., Allen, R., de Charentenay, J. and Kirkham, D., "CFD Prediction of Narrowband and Broadband Cavity Acoustics at $M=0.85$," AIAA 2003-3303.
- ⁷Basu, D., Hamed, A. and Das, K., "DES and Hybrid RANS/LES models for unsteady separated turbulent flow prediction," AIAA 2005-503.
- ⁸Tramel, R., Rock, S. and Ellis, J., "Comparison of Large Cavity Aeroacoustic Computations with Flight Test Results," AIAA 2005-2800.
- ⁹Schmid, S., Lutz, T. and Kramer, E., "Numerical Simulation of the Flow Field Around the Stratospheric Observatory For Infrared Astronomy," in *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, C. Tropea, S. Jakirlic, H. J. Heinemann, R. Henke and H. Honlinger, Eds., Vol. 96, Springer, Berlin, 2007.

- ¹⁰Schmid, S., Lutz, T., Kramer, E. and Kuhn, T., "Passive Control of the Flow of the Stratospheric Observatory for Infrared Astronomy," AIAA 2008-6717.
- ¹¹Schmid, S., Lutz, T. and Kramer, E., "Simulation of the Flow Around the Stratospheric Observatory For Infrared Astronomy (SOFIA) Using URANS and DES," in *Transactions of the Third Joint HLRB and KONWIHR Status and Result Workshop*, Dec. 3–4, 2007, Leibniz Supercomputing Center, Garching/Munich, Germany, S. Wagner, M. Steinmetz, A. Bode, M. Brehm, Eds. 2009.
- ¹²Schmid, S., Lutz, T. and Kramer, E., "Simulation of the Unsteady Cavity Flow of the Stratospheric Observatory For Infrared Astronomy," in *Proceedings of the IUTAM Symposium "Unsteady Separated Flows and their Control"*, Braza, M. and Hourigan, K., Eds., 2009.
- ¹³Litton, D. K., Edwards, J. R. and White, J. A., "Algorithmic Enhancements to the VULCAN Navier-Stokes Solver," AIAA 2003-3979.
- ¹⁴Roe, P.L., "Characteristics-Based Schemes for the Euler Equations," *Annual Review of Fluid Mechanics*, Vol. 18, 1986, pp. 337–65.
- ¹⁵Edwards, J. R., "A Low-Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations," *Computers and Fluids* Vol. 26, No. 6, 1997, pp. 635–659.
- ¹⁶Venkatakrishnan, V., "On the Accuracy of Limiters and Convergence to Steady-State Solutions," AIAA 93-0880, 1993.