

Chapter Five: Convergence

Introduction

Before results of any experiments can be meaningfully discussed or compared, measures must be taken to establish and verify the quality and repeatability of the experiments, and numerical experiments are no exception. To that end this chapter features an investigation of convergence. Two kinds of convergence are considered, convergence to steady state and grid convergence. Both can be difficult to achieve in complex flows.

Convergence to Steady State

The present investigation, like most similar ones, employs the time-dependent Navier-Stokes equations to solve for steady-state flow solutions. When using this technique, an initial “guess” solution (such as freestream flow) is supplied and the time derivative of the solution vector, $\partial\bar{q}/\partial t$, is treated as a relaxation or error term. The solution is then iterated until the error reaches zero, and the determination of that zero-error (“steady”) state becomes one of the most critical parts of the solution procedure.

Perhaps the most straight-forward measure of convergence to steady-state is the residual itself, as defined in Equation 4.1. Since the residual vector is directly proportional to the time rate of change of the primitive variable vector, it would seem logical that steady-state convergence could not be reached while the residual (or a norm thereof) remains nonzero. This approach is limited by a number of factors. First is the finite resolution of the computer in use, which makes an exact representation of any number (such as zero) unlikely. Second is the use of nondimensionalization, as such simple techniques as dividing the residual by its initial value can change dramatically the apparent size and significance of a later value. For these reasons and reasons like them many researchers consider the change in the residual more useful in determining convergence than its actual value. A decrease in the size of the residual indicates a

proportional reduction in the time rate of change in the solution itself, so a decrease of, say, two orders of magnitude in the residual means that the steady-state error is one hundred times smaller than it was at the beginning, and so forth. Clearly there might be a point beyond which further changes in the solution are small enough to be insignificant. Researchers have different opinions about the change in the residual necessary to reach such a point, but a common value might be three to five or even eight orders of magnitude decrease in the residual, depending on the problem (flow physics and solution scheme) in question.

Unfortunately, such a scheme for judging convergence is fraught with uncertainty. Even in "simple" problems (meaning laminar, largely inviscid, nearly uniform flows calculated on orthogonal grids), engineering judgement is necessary to determine what decrease in the residual is sufficient evidence that further iteration would not significantly change the solution, and convergence criteria for highly viscous, turbulent flows are more uncertain still. Additional factors introduced by turbulence modeling, flux limiters, flux vector splits, etc., may prevent the residual or its norm from ever decreasing by the orders of magnitude considered necessary for convergence in simpler flows. Under circumstances such as these, the norm of the residual or even the residual itself might cease to change appreciably, either oscillating about a fixed point or asymptotically approaching a certain value. This lack of clear decrease in the residual does not prove that the solution itself is no longer changing, however, just as a change in the residual does not always indicate that significant flow evolution is occurring. In these cases additional convergence criteria are useful.

Another obvious rule for stopping a calculation is that a calculation should never be considered converged to steady state while the solution vector itself is changing. In theory, one could store the solution vector (the updated vector of primitive or conserved variables at all points in the flow) periodically and

subtract successive solutions until the difference reaches machine zero for all variables at all points in the flowfield. With this approach it is the sheer volume of information that limits practicality. Multispecies flows with differential-equation turbulence models frequently calculate upwards of a half dozen, a dozen, or even more flow variables at each cell in the computational domain, and highly viscous, complex flowfields often require on the order of one to ten million cells for accurate calculation. The researcher is therefore presented with literally millions of pieces of data for analysis, especially when one expands the list of variables for comparison to include commonly used functions of the primitive variables, such as skin friction, vorticity, etc. It is unreasonable to expect that all variables at all points in the flow will be equally important as convergence indicators, or that they will reach steady-state after the same number of iterations. The latter is especially true when local timestepping is employed or functions of interest require derivatives of other variables. If all variables were considered equally important in determining convergence, the researcher might be forced to extend the iteration process, increasing greatly the expense of the calculation in order to improve the computed value of a few, isolated data points which might be unimportant to the specific purpose of the calculation.

A sensible modification of the method described above is to monitor the evolution of only a limited number of flow variables and features, and to stop the calculation when all the flow features of interest cease changing. Experience has shown this rule to be a useful one, within its limitations. From one point of view, if the purpose of the calculation is to obtain information about a few quickly converged flow features, one can save much time by stopping the calculation before more slowly evolved, irrelevant features have converged. From another point of view, no set of convergence criteria, no matter how elaborate or conservative, can properly pronounce steady-state convergence if some flow feature of interest is still changing. But, researchers seldom know in advance to

exactly what uses their calculations might be put, and care must be taken to allow for unexpected applications of the data. As a result, the above rule is necessary but not sufficient, and other convergence criteria are also important.

The list of possible and useful convergence criteria is long. Experienced researchers often have their own list of favorite convergence indicators, consisting primarily of flow features and variables of frequent interest to their particular application or those that are known to evolve slowly. Another common indicator of the attainment of steady-state is the conservation of mass: steady-state convergence is said to be impossible until inflow and outflow mass flow rates are constant and match to within some set limit (for example, a few tenths of a percent). As a convergence indicator the conservation of mass is a simplification of the variable-monitoring criteria discussed above, as it considers only one variable, mass, in integrated form instead of pointwise, and that at only two locations in the flowfield, inflow and outflow boundaries. Yet another approach is to assume a connection between convergence and the movement of particles and information through the domain. Interpreting the CFL number as the number of grid-cells a typical particle crosses during one iteration (though it may not be this exactly, depending on the characteristic length and speed chosen for the CFL number calculation), the number of iterations needed for a particle to completely pass through the computational domain can be calculated. Practitioners of this technique then multiply that number of iterations by some factor, typically between three and ten, and expect convergence when the iteration count indicates reaches that level. Neither this nor any of the other criteria, however, is by itself sufficient, and a good deal of experience and judgment are required to use them meaningfully.

A number of criteria were set to define convergence to steady state in the current investigation. Among them were convergence of centerline (symmetry line) surface pressure coefficient and skin-friction coefficient plots, and general convergence of contours of mass

fraction, Mach number, and static pressure. These contours were examined on various planes within and downstream of the primary mixing regions and at the exit planes of the zones. Particular attention was paid to helium mass fraction and total pressure because it was these data, respectively, which were to be compared with the experiments. Care was also taken to ensure convergence of small, viscous-dominated cells near walls. (Remember that with local timestepping the time increment applied to a given cell is proportional to its characteristic length, so that smaller cells evolve more slowly in time.) In this investigation, the helium mass fraction contours were usually slower and more difficult to converge than any other monitored variable, especially near the walls in the downstream regions. Hence, they proved to be the practical test of convergence. Mass conservation was monitored and proved to be a useful tool for error analysis, though as a measure of convergence it was no more useful than some other techniques. The iteration-count technique discussed above was also considered, though the use of freestream velocities in calculation of the CFL number (a conservative approach to improve stability) reduced its usefulness.

As discussed above, mesh sequencing was employed as a means of convergence acceleration. The converged solution from the coarse mesh was interpolated for use as an initial solution for the medium mesh, and likewise with the converged medium-mesh solution and the initial fine-mesh solution. In such cases the initial value of the residual for the medium and fine-mesh solutions is smaller than it might be otherwise, because of the more accurate initial solution, and it decreases by a lesser amount. Whereas the residuals for coarse-mesh solutions in globally iterated zones, which were initialized with a freestream solution, sometimes dropped as much as five orders of magnitude without leveling off, the normalized residuals for the medium and fine meshes usually began approximately two orders of magnitude lower than those of the coarse meshes and decreased by only one and a half to three additional orders of magnitude before asymptotically approaching or

fluctuating about some fixed value. Particularly for the fine meshes, the solutions often continued to change after (sometimes long after) the residual leveled off.

A similar technique was used for parallel calculations with different turbulence models. In the interest of conservation of computer time, the solution from the Chien $K-\epsilon$ calculation was used as an initial solution for the RNG calculations of the same problem. When the Chien- $K-\epsilon$ solutions were applied to an RNG equation set, the residuals showed an increase from their previous levels, as would be expected. For the coarse meshes the increase was small, but for the fine meshes it was sometimes as much as two orders of magnitude or more. The RNG residuals then began to converge smoothly, and eventually returned to level approximately equal to those of the $K-\epsilon$ calculation. But in these cases, as in the case of mesh sequencing, the residuals in globally iterated zones were of limited usefulness.

In space-marched zones, solutions were usually initialized with data from the upstream plane. Convergence was easier in these situations and the residuals provided slightly more information. Five orders of convergence were usually possible, at the end of which the solution was spot-checked for convergence. The five orders were almost always found to be sufficient.

Grid Convergence

The second kind of convergence investigated was grid convergence. Grid convergence, also known as grid independence, is investigated by comparing results from all levels of the meshes used in mesh sequencing. If grid independence has been reached, the solutions from the two finest meshes should be identical and the researcher can be confident that a finer mesh would not improve the solution. Unfortunately, proving grid convergence requires completing two sets of calculations, one far larger than actually necessary to produce a good solution for the problem of interest. (A typical size for the second calculation would be eight times larger

than the first, resulting from doubling the number of cells in each direction.) Even in cases where computer memory and time are sufficient to produce the first good-quality solution, very rarely do researchers have sufficient resources to prove the quality of their solutions by completing a second, larger calculation. As a result, grid convergence remains unproven in many calculations of the present type.

The present problems suffered from the same limitations as many others of their type. Though the fine-mesh solutions were expected to be reasonably well-resolved, the limitations imposed by computer memory and processor time eliminated the possibility of completing yet a larger solution for each case. Instead, the coarser, smaller, medium-mesh solutions were used for grid convergence comparisons with the fine-mesh solutions. Technically speaking, this grid convergence study was a study of the quality of the medium-mesh solutions, not the fine-mesh solutions, and it should have come as no surprise that grid convergence was not proven for the medium meshes. By virtue of their smaller size, these medium meshes were not expected to produce results identical to those produced by the finer meshes, and they did not. But this failure does not comment directly on the quality of the fine-mesh solutions themselves. That information must be inferred indirectly.

Flat Plate

As mentioned in Chapter Four, two different meshes were tested for the flat plate investigation. Because this was a two-dimensional calculation, extremely fine meshes were possible and the coarse and fine grids had dimensions 101×101 and 201×201 , respectively. Figure 5.1 shows the velocity profile at the outflow plane, a distance of 1.00 m from the leading edge of the plate, for the two meshes. Figure 5.2 presents the static temperature profile at the same location, and Figure 5.3 shows static density. Clearly there is very little difference, if any, between the two solutions for any of the three variables. Based on this data, the coarse mesh, with 101 grid

points in each of the normal and streamwise directions, appears to be as close to complete grid convergence as is practicable and necessary. The fine mesh, with twice as many computational cells in each direction, should certainly be sufficient for the intended turbulence model comparison.

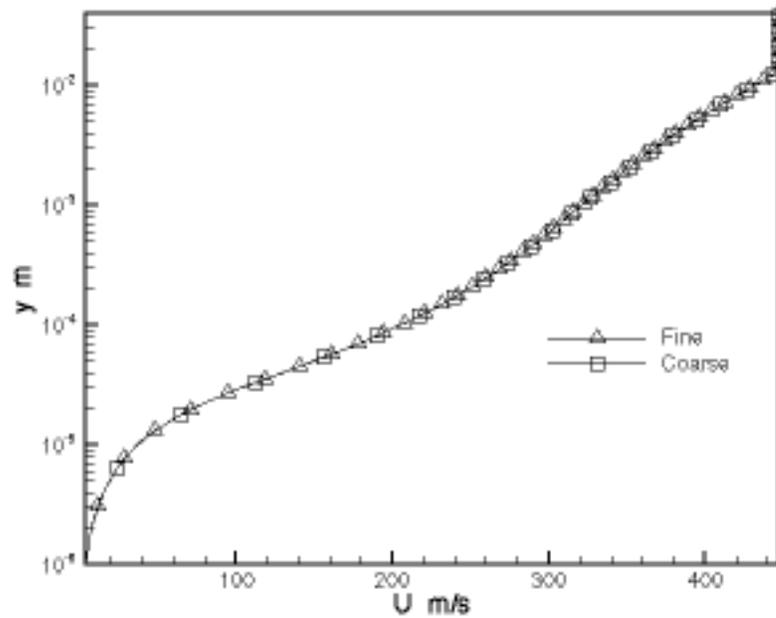


Figure 5.1. Streamwise Velocity versus Wall-Normal Distance for Coarse and Fine Meshes for Flat Plate Test Case at Plane $x = 1.00$ m. Every Fourth Grid Point is Marked.

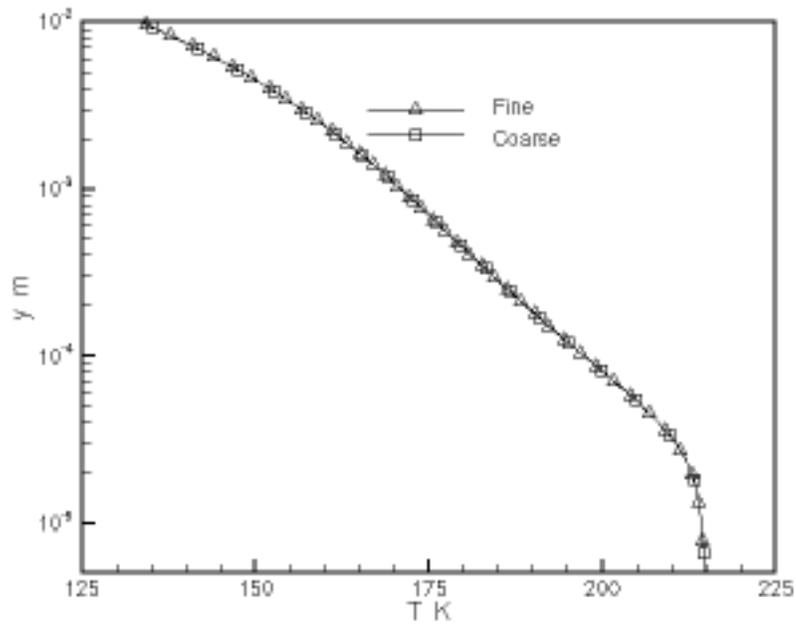


Figure 5.2. Static Temperature versus Wall-Normal Distance for Coarse and Fine Meshes for Flat Plate Test Case at Plane $x = 1.00$ m. Every Fourth Grid Point is Marked.

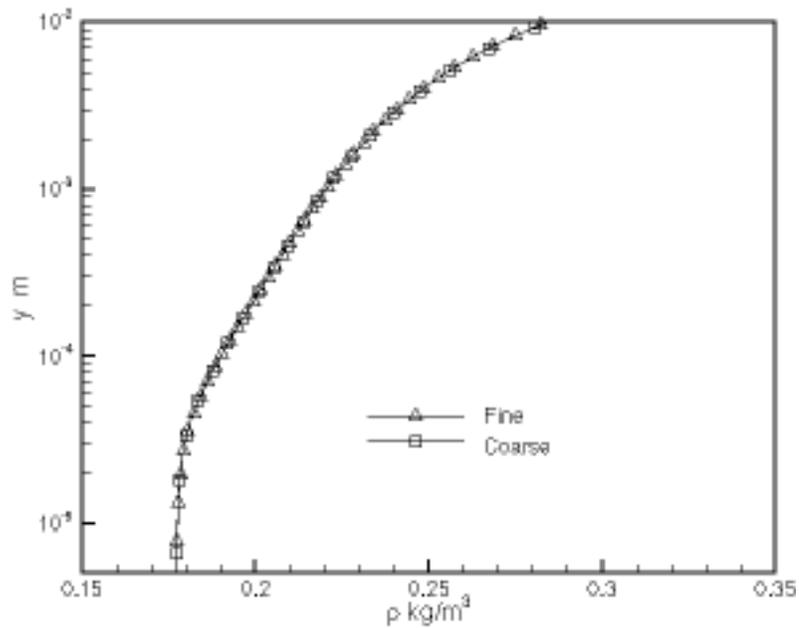


Figure 5.3. Static Density versus Wall-Normal Distance for Coarse and Fine Meshes for Flat Plate Test Case at Plane $x = 1.00$ m. Every Fourth Grid Point is Marked.

Wall Injector Array

In the interest of simplicity and conservation of resources, only one solution set, the Chien $K-\epsilon$ case, was used to judge grid convergence for the wall injector array. Solutions with the other turbulence models were generated only on the fine grids, as will be discussed below. This approach assumes that grid convergence characteristics will be at least marginally independent of turbulence model, an unproven but reasonable hypothesis.

Because helium mass fractions were the only data available for this test case, helium mass fractions were chosen as the primary variable for use in grid independence comparisons. Table 5.1 compares peak mass fractions of helium in the core of the injectant plume for each mesh at three axial locations, two in the nearfield (six and nine effective diameters downstream of the center of the first row of injectors) plus the outflow plane, 50.39 effective diameters downstream from the same point. All data is taken from Chien $K-\epsilon$ calculations. The values presented are for core-flow only; where a higher mass fraction occurred in the boundary layer or near-wall region, it was discarded.

Table 5.1 Peak Helium Mass Fractions at Three Axial Locations for All Mesh Sequence Levels for Nine-Hole Injector Array

Location	Coarse Mesh	Medium Mesh	Fine Mesh
$x/d_{\text{eff}}=6.00$.9424	.9568	.9686
$x/d_{\text{eff}}=9.00$.7121	.7083	.7803
$x/d_{\text{eff}}=50.39$.2148	.2949	.2513

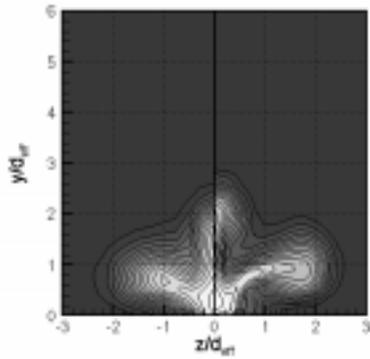
Clearly Table 5.1 confirms the hypothesis that grid independence would not be proven with grids feasible for this investigation. Beyond this, however, little is immediately clear. Most problematic is the unexpected tendency of the medium-mesh solution to represent an extremum, a maximum or minimum value,

instead of some value intermediate to the others. This pattern disturbs the convergence trends and makes judging grid convergence more difficult. Possible explanations for and consequences of this phenomenon will be presented shortly.

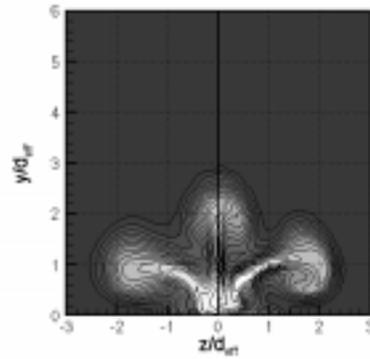
One additional inference from Table 5.1 is that a small number of point values cannot adequately summarize data from an entire flowfield. To assist in the assessment of grid convergence, Figures 5.4 through 5.6 present contour plots comparing helium mass fractions for the coarse, medium, and fine-mesh solutions at the axial locations referenced above. Figures 5.7 through 5.9 show contours of static pressure for the same grid levels and locations.

A number of conclusions can be drawn from Figures 5.4 through 5.9. The first, which is supported by the data in Table 5.1, is that grid convergence is more closely approached in the nearfield region of the flow. The mass-fraction contours corresponding to the medium and fine meshes in Figures 5.4 and 5.5, less than ten effective diameters downstream of the first row of injectors, bear a strong resemblance to one another; the corresponding contours in Figure 5.6, much further downstream, do not. Similarly, the first row of values in Table 5.1, the peak helium mass fractions at $x/d_{\text{eff}} = 6.00$, show strong similarity and clear trends, something data taken further downstream does not show. The pressure contours in Figures 5.7 through 5.9 illustrate the same feature, with nearfield data showing definite trends toward grid convergence and the farfield data showing essentially none.

This increase in grid convergence in the nearfield region is not unexpected. Because of the presence of the injectors, the grid was clustered more strongly in the nearfield region, giving rise to a smaller average grid dimension. Additionally, the flow in the nearfield region is expected to be dominated by the effects of inviscid mixing, which tends to be less sensitive to grid spacing than small-scale viscous and turbulence effects. Finally, the downstream portions of the flow solution incorporate the cumulative effects of all

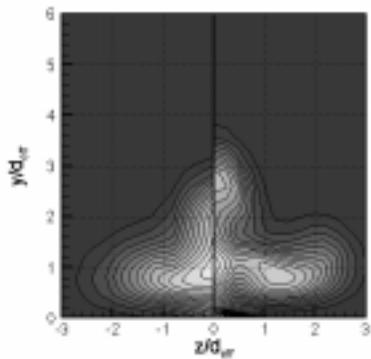


(a) Coarse (left) and Medium Meshes

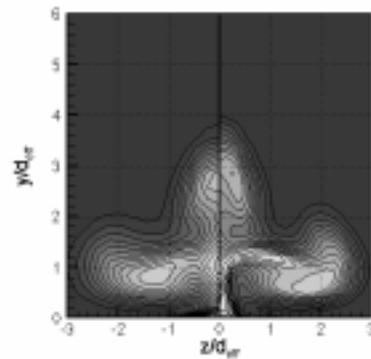


(b) Medium (left) and Fine Meshes

Figure 5.4. Helium Mass Fraction Contour Plots at $x/d_{eff} = 6.00$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.



(a) Coarse (left) and Medium Meshes



(b) Medium (left) and Fine Meshes

Figure 5.5. Helium Mass Fraction Contour Plots at $x/d_{eff} = 9.00$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.

upstream grid-resolution errors, magnifying the appearance of the error.

The second conclusion to be drawn from the figures is that the coarse-mesh solution is fundamentally different from the fine-mesh solution and even from the medium-mesh solution. While the near-

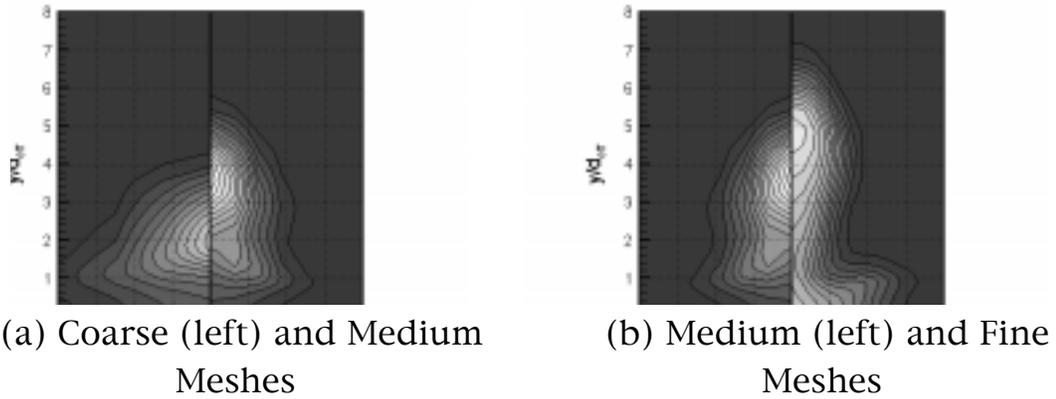


Figure 5.6. Helium Mass Fraction Contour Plots at $x/d_{\text{eff}} = 50.39$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.

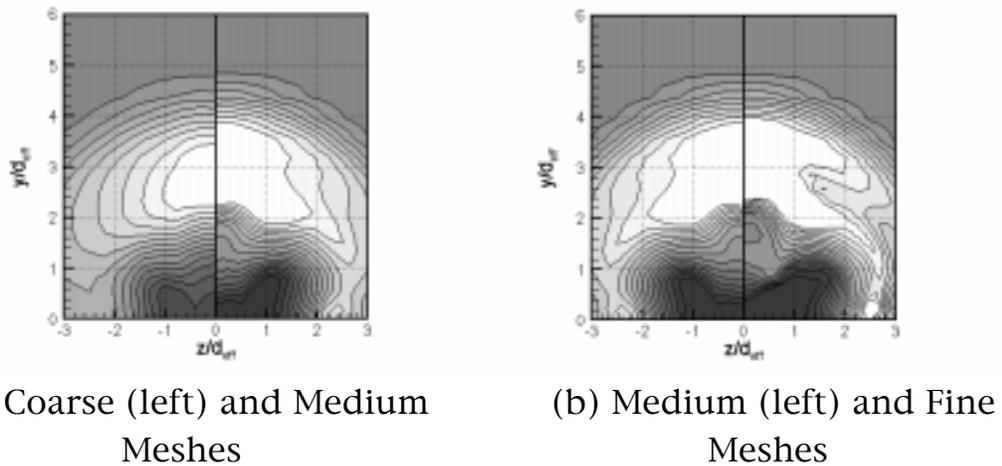
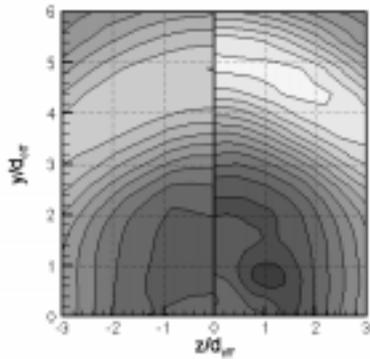
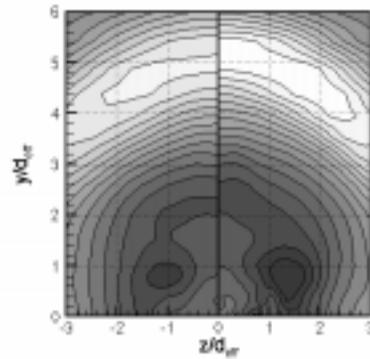


Figure 5.7. Static Pressure Contour Plots at $x/d_{\text{eff}} = 6.00$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.

field data corresponding to the coarse grid does somewhat resemble those of the finer meshes, especially in helium mass-fraction plots, the downstream data clearly does not. One can only conclude that the coarse-mesh solution is sufficiently under-resolved to fail to predict one or more important flow features and/or mixing mechanisms, particularly in the farfield region. With this in mind, it is likely that the inconsistent trends noticed in Table 5.1 are more

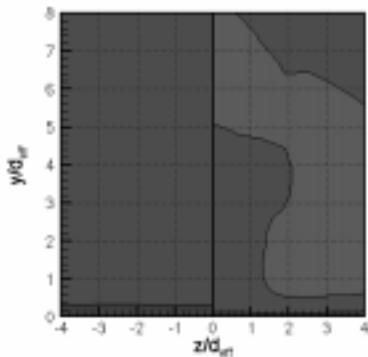


(a) Coarse (left) and Medium Meshes

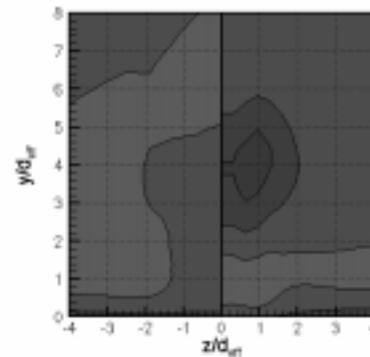


(b) Medium (left) and Fine Meshes

Figure 5.8. Static Pressure Contour Plots at $x/d_{eff} = 9.00$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.



(a) Coarse (left) and Medium Meshes



(b) Medium (left) and Fine Meshes

Figure 5.9. Static Pressure Contour Plots at $x/d_{eff} = 50.39$ for All Mesh Sequence Levels for the Nine-Hole Injector Array.

the fault of an overly crude coarse-mesh solution than a medium or fine-mesh solution that is flawed in some fundamental manner.

To provide additional insight into the flow in the nearfield region, Figures 5.10 and 5.11 plot the coefficient of pressure on the lower wall along the centerline in the vicinity of the injectors. In the

figures the origin of the x axis is again at the center of the first row of injectors. Figure 5.10 presents a large-scale view of the data, showing clearly the elevations of pressure associated with each centerline injector. Note that at the lower wall, which is the injector exit plane, each injector plume has a different static pressure, as the influence of the surrounding flowfield penetrates into the injector ports themselves. The peak pressures predicted by the medium-mesh calculation agree much more closely with those of the fine-mesh calculation than with those of the coarse mesh, suggesting reasonable grid convergence in this area for this problem. Figure 5.11 allows a closer examination of the low-pressure wake regions that follow each high-pressure injector plume. These do not indicate grid convergence to the extent that the high-pressure plume regions do, but even in these regions the medium-mesh solution more closely resembles the fine-mesh solution than it does the coarse-mesh solution. In the second and third wake regions, the fine-mesh solution shows sharp curvature the medium-mesh solution does not, perhaps indicating a failure of the medium-mesh solution to resolve some flow feature. Thus, Figures 5.10 and 5.11 present a mixed, if somewhat favorable, message on grid convergence.

Data presented has not fully confirmed grid convergence for the nine-hole injector array. Given the relatively small size of the grid (the third-row injectors were represented by only four cells each in the spanwise and streamwise directions), the complexity of the flowfield, and the inability to examine grid convergence of the fine mesh, proof of grid convergence was not expected for this case. Comparisons between the medium and fine-mesh solutions in the nearfield region have indicated that the medium mesh is approaching grid convergence in that region, and this evidence further suggests that the fine-mesh solution should be somewhat reliable (from the standpoint of grid convergence) in this region, even though it is not expected to be perfectly grid-converged. Further downstream, even the medium and fine-mesh solutions diverge, suggesting the need for greater resolution and indicating

that caution should be employed in the interpretation of farfield data.

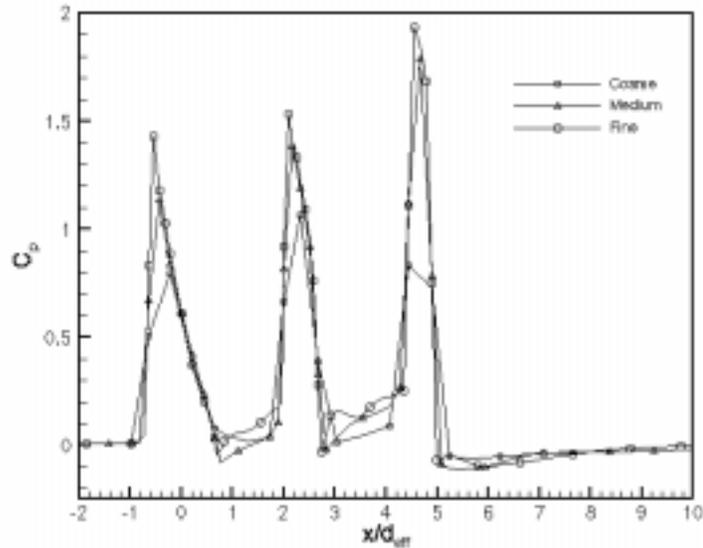


Figure 5.10. Centerline Pressure Coefficient versus Axial Distance in the Vicinity of the Injectors for the Nine-Hole Injector Array.

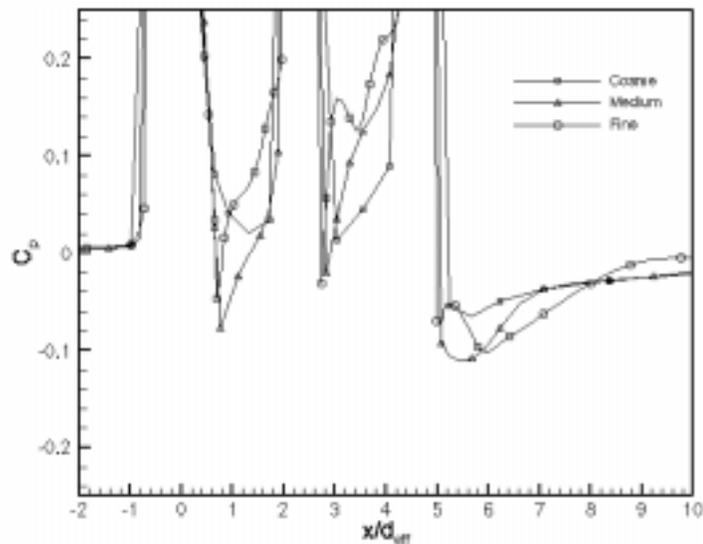


Figure 5.11. Centerline Pressure Coefficient versus Axial Distance in the Vicinity of the Injectors for the Nine-Hole Injector Array. Detail of Wake Regions.

Ramp Injector

Grid convergence studies for the ramp injector, like those for the nine-hole injector array, involved only calculations performed using the Chien $K-\epsilon$ turbulence model. Once again this approach assumes that grid convergence is relatively independent of the turbulence model used in the calculation, an assumption considered reasonable given the similarity of the turbulence models in question. One difference between the grid convergence study of the ramp injector and that of the nine-hole array is the mesh sequencing: recall that the ramp injector calculations employed mesh sequencing in all three directions, while the nine-hole injector array calculations did not sequence normal to the lower wall. As a result, the size difference between the coarse and fine grids for the ramp case is larger than that for the array case, and the coarse and fine-mesh solutions are expected to differ more strongly.

Taking the data intended for comparison with experiment to be among the most important for grid convergence studies, helium mass fraction and pressure were the primary variables for this portion of the analysis. Table 5.2 presents the maximum helium mass fraction at three axial locations, $x/h=0.5$, $x/h=2.5$, and $x/h=20.0$. Axial coordinates in the ramp investigation in present and future chapters are nondimensionalized by ramp height and referenced to a streamwise origin at the ramp face. All data is taken from Chien $K-\epsilon$ calculations. The values presented are for core-flow only; for this case high concentrations of helium were not present in the boundary layer.

The nearfield ($x/h=0.500$) data in Table 5.2 show the same anomaly present in the nine-hole injector array data in Table 5.1, in that the medium-mesh value is not bounded by the other two. In this case, however, the maximum difference between values is less than two percent, and the unboundedness of the middle value may not be significant. Moreover, the trend is not seen at the downstream locations. The data at $x/h=2.50$ exactly follows expected

Table 5.2 Peak Helium Mass Fractions at Three Axial Locations for All Mesh Sequence Levels for Four-Hole Ramp Injector

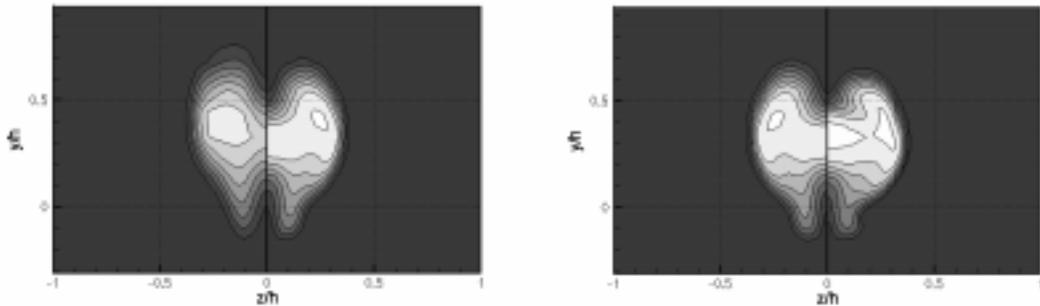
Location	Coarse Mesh	Medium Mesh	Fine Mesh
x/h=.500	.08314	.08457	.08326
x/h=2.50	.07623	.08087	.08327
x/h=20.0	.01271	.01310	.01359

trends, with the medium-mesh value bounded by the coarse and fine-mesh values but in closer agreement with the fine. At $x/h=20.0$ the trend is almost as positive. Again, the medium-mesh value is bounded between the other two, though it is slightly closer to the coarse-mesh value than the fine. While the lack of agreement with the fine-mesh value is disappointing, the difference between the fine-mesh value and the coarse is less than six and a half percent, a reasonably narrow range given that the coarse mesh has sixty-four times fewer grid points than the fine. Recognizing that a few point values cannot determine convergence and acknowledging the anomalous behavior of the nearfield data, there is nonetheless a clear trend toward convergence. This trend will be further explored with additional data below.

Figures 5.12 and 5.13 expand upon the data in Table 5.2 by presenting cross-flow plots of helium mass fraction at two of the axial locations, $x/h=2.50$ and $x/h=20.0$, respectively. Figures 5.14 and 5.15 present similar plots of static pressure at the same locations. In each figure coarse and medium-mesh solutions are presented in part (a) and medium and fine-mesh solutions are presented in part (b). Contour levels at the downstream location, $x/h=20.0$, have different values than those in the nearfield. Otherwise, the levels are consistent throughout this analysis.

The nearfield helium mass-fraction data presented in Figure 5.12 clearly shows a great deal of similarity between the solutions generated on all three meshes. The sizes, shapes, and locations of the

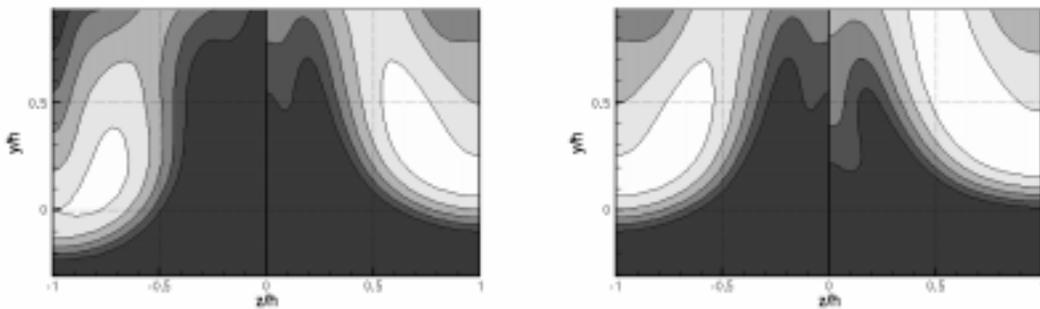
outer (lowest-valued) contours are remarkably consistent. There is a slight tendency toward overdiffusion on the coarser meshes, as evidenced by the slightly smaller sizes and closer spacing of the outer contours of each solution with respect to those coarser. This tendency is consistent with that of the innermost (highest-valued) contours, which show an increase in helium mass fraction as the grid



(a) Coarse (left) and Medium Meshes

(b) Medium (left) and Fine Meshes

Figure 5.12. Helium Mass Fraction Contour Plots at $x/h = 2.50$ for All Mesh Sequence Levels for the Ramp Injector.



(a) Coarse (left) and Medium Meshes

(b) Medium (left) and Fine Meshes

Figure 5.13. Helium Mass Fraction Contour Plots at $x/h = 20.0$ for All Mesh Sequence Levels for the Ramp Injector.

is refined. While all three solutions show generally similar patterns, the fine-mesh solution captures a centerline concentration peak the medium mesh plot does not reveal. That is not to say the medium-mesh solution completely misses the centerline peak; mostly likely, a lower-valued peak exists but was not captured by the (somewhat arbitrarily) chosen contour values. The other features of the inner and outer contours of the medium-mesh solution are remarkably

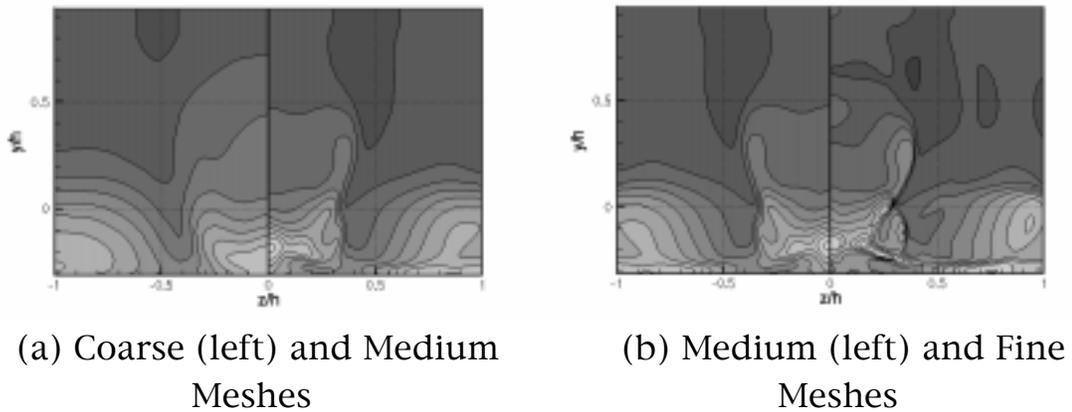


Figure 5.14. Static Pressure Contour Plots at $x/h = 2.50$ for All Mesh Sequence Levels for the Ramp Injector.

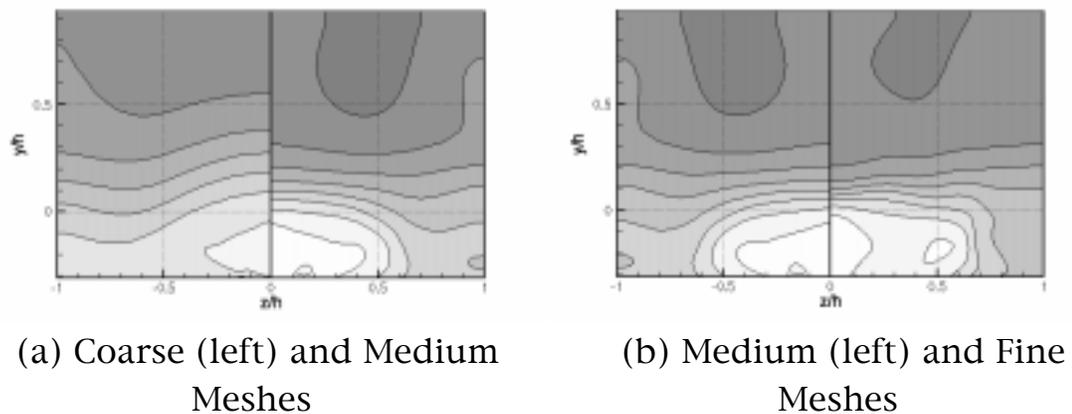


Figure 5.15. Static Pressure Contour Plots at $x/h = 20.0$ for All Mesh Sequence Levels for the Ramp Injector.

similar to those of the fine-mesh solution. Differences between the coarse and medium-mesh solutions are much larger, with the coarse-mesh solution completely missing the innermost helium contour. While absolute grid convergence is not indicated for the medium mesh, there is a strong trend toward relative grid convergence. For many applications the level of grid convergence indicated by Figure 5.12 for the medium mesh would certainly be acceptable. It is therefore likely that the fine mesh is at least approaching a level of grid convergence appropriate for the present investigation, at least as far as nearfield helium data is concerned.

Farfield helium data presented in Figure 5.13 does not show quite such strong trends toward grid convergence, though it does show a general tendency. While all three meshes capture the overall shape of the helium plume, the coarser meshes overdiffuse the helium with respect to the fine-mesh solution, producing larger, lower-valued bands of helium with smaller peaks near the centerline. Differences between the coarse and medium-mesh solutions are on the whole considerably larger than those between the medium and fine-mesh solutions, suggesting a measure of grid convergence has been reached for the medium mesh. As with the nearfield helium data, this data supports the hypothesis that the fine-mesh solution is at least approaching a level of grid convergence appropriate for the present investigation.

The nearfield pressure data presented in Figure 5.14 show a tendency toward convergence, though not convergence itself. The medium and even coarse-mesh solutions do seem to capture the major features of the pressure field at this location, though they fail to resolve the small details revealed by the fine-mesh solution. In particular, the fine-mesh solution contains much sharper gradients and more and stronger extrema, with higher peaks and lower minimum values than those of the medium-mesh solution. Yet the qualitative trends in the pressure fields of the medium and fine-mesh solutions are quite similar. Certainly there is a tendency toward grid convergence, though convergence itself has not been

reached. For certain qualitative purposes, this tendency toward convergence may be sufficient. Once again this conclusion does not impact directly upon the grid convergence of the fine mesh; one can infer, however, that the fine mesh should be reasonably well resolved for many purposes.

The pressure field at the downstream location, $x/h=20.0$, shown in Figure 5.15 is much less eventful than that of the nearfield, and there appear to be fewer strong gradients to resolve. The fine-mesh solution does show sharper gradients than the others do, with a smaller region of high pressure and a more uniform pressure field away from that region. Nonetheless the differences between the medium and fine-mesh solutions are smaller than those between the coarse and medium-mesh solutions, suggesting a measure of grid convergence for the medium mesh. Given the denser mesh of the fine-grid solution, it is likely that the degree of grid convergence exhibited by the fine-mesh solution, could it be quantified, would be acceptable for the present purpose.

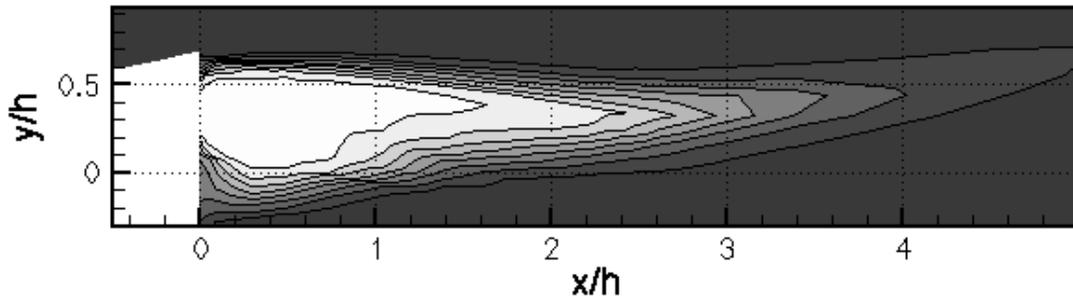
Figures 5.16 through 5.18 continue the analysis by presenting helium mass fraction and static pressure on a vertical plane at the centerline of the ramp (the symmetry plane). Once again the streamwise origin is located at the face of the ramp. In these figures the ramp itself appears as a white trapezoid or triangle at the lower edge of the figure. Figure 5.16 shows helium mass fraction downstream of the ramp. Figure 5.17 shows static pressure contours above the ramp itself, and Figure 5.18 shows static pressure contours downstream of the ramp.

The centerline helium data presented in Figure 5.16 is consistent with that of Figure 5.12. While absolute grid convergence has not been achieved for the coarse or medium meshes, the qualitative agreement between all solutions is strong. As was evident in Figure 5.12, there is a tendency for the coarser solutions to diffuse helium too quickly, resulting in shorter, narrower plumes and an artificially low value of helium mass fraction at any given

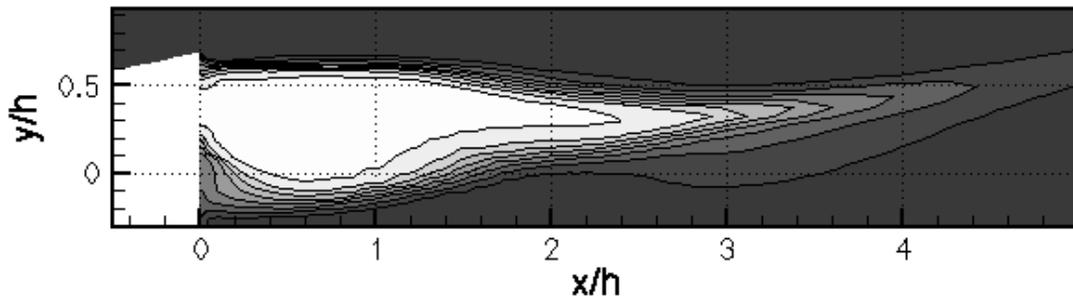
location. Nonetheless the qualitative agreement between the solutions strongly suggests a measure of grid convergence.

Figure 5.17 highlights the pressure rise associated with the oblique shock that forms at the upstream end of the ramp. The fine-mesh solution does show a crisper shock than either of the others, though the differences between the medium and fine-mesh solutions are smaller than those between the coarse and medium-mesh solutions. Downstream of the ramp, in Figure 5.18, similar trends are evident, as the coarse-mesh solution tends to smooth out extrema the finer meshes capture. The fine-mesh solution itself shows sharper contrasts in pressure than the medium-mesh solution, with higher-valued local maxima and larger regions of low pressure, but the shapes of the contours and their placements are quite similar. By contrast, the coarse-mesh solution is much more diffuse. While the coarse-mesh solution cannot be considered grid converged, the medium-mesh solution might be sufficiently converged for some qualitative purposes, and there is reason for some optimism about the grid convergence of the fine-mesh solution.

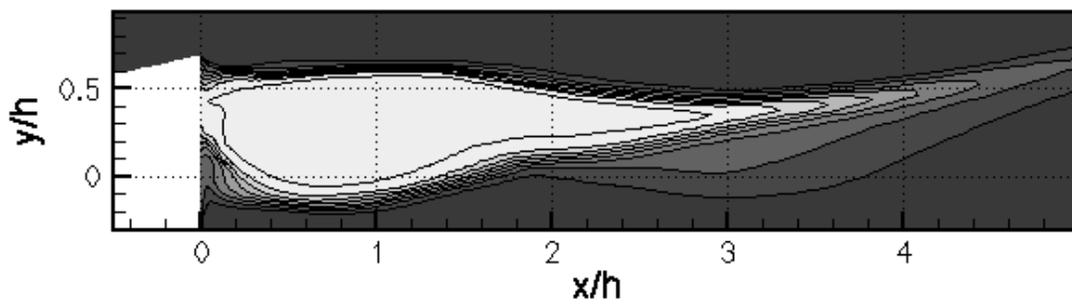
Grid convergence is difficult to achieve in complex, three-dimensional flowfields, and in the case of the four-hole ramp injector the medium mesh was not expected to be perfectly grid-independent. Nonetheless a remarkable tendency toward grid convergence has been observed for the medium mesh. Though the grid independence of the fine mesh cannot be investigated directly, the degree of grid independence observed in the medium-mesh data suggests that the fine-mesh solution is likely to be adequately grid-converged for the present purpose.



(a) Coarse Mesh

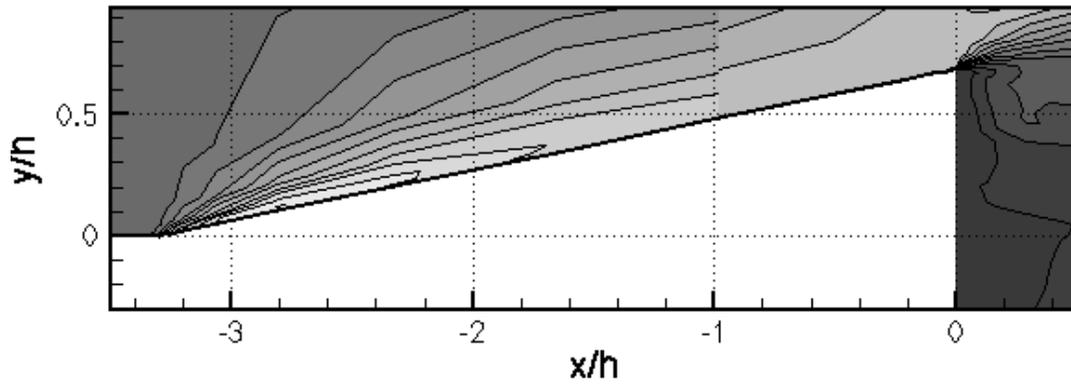


(b) Medium Mesh

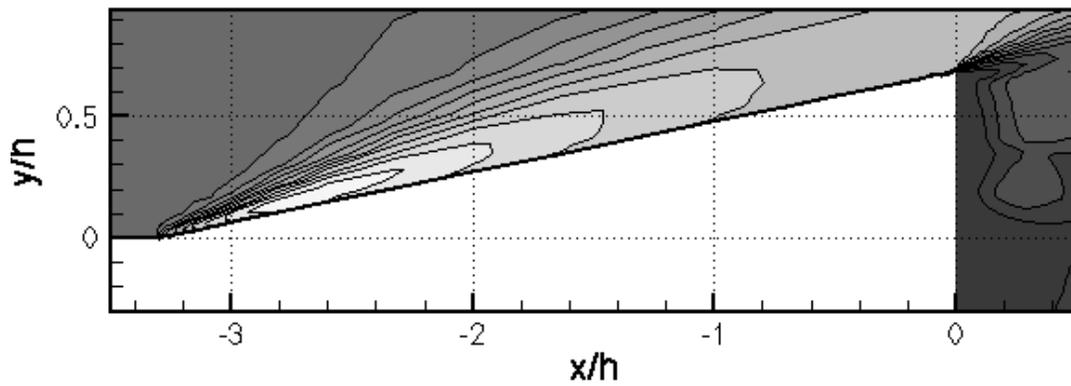


(c) Fine Mesh

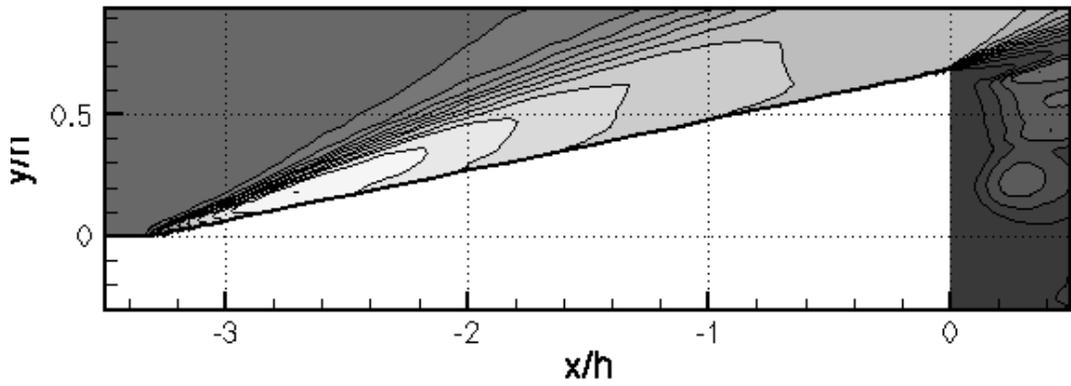
Figure 5.16. Helium Mass Fraction Contours on the Centerplane of the Ramp Injector for All Mesh Sequence Levels.



(a) Coarse Mesh

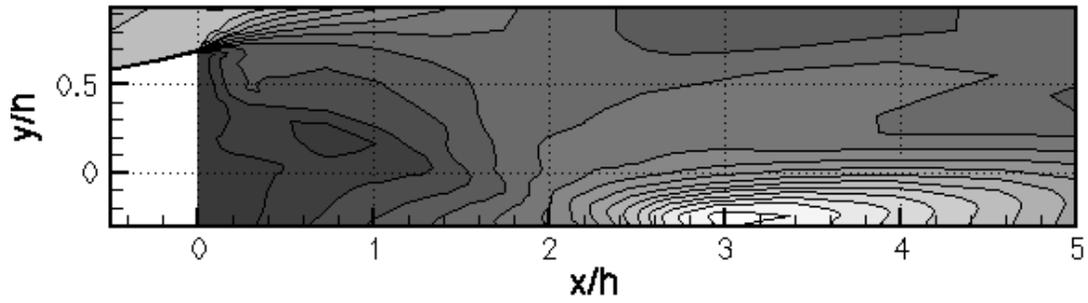


(b) Medium Mesh

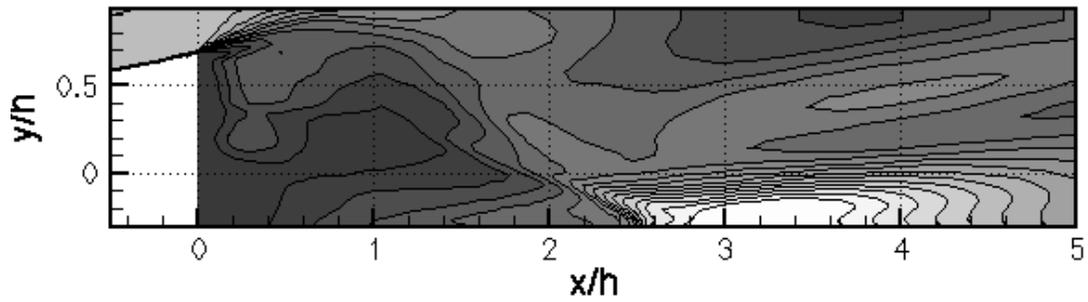


(c) Fine Mesh

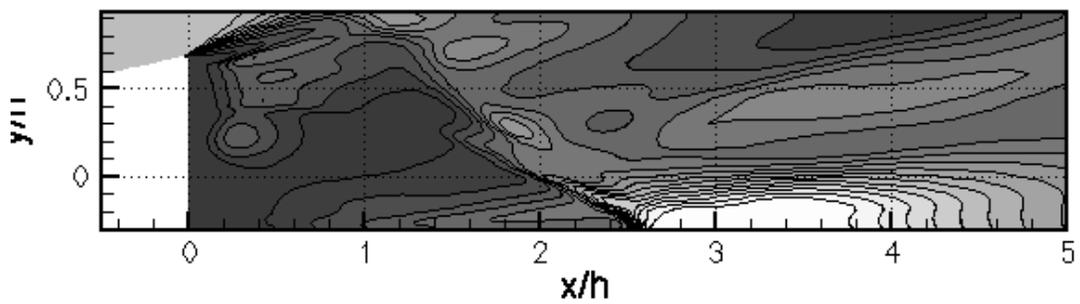
Figure 5.17. Static Pressure Contours Above the Ramp Injector for All Mesh Sequence Levels.



(a) Coarse Mesh



(b) Medium Mesh



(c) Fine Mesh

Figure 5.18. Static Pressure Contours Downstream of the Ramp Injector for All Mesh Sequence Levels.