On Problems Associated with Grid Convergence of Functionals

Manuel D. Salas* and Harold L. Atkins
NASA Langley Research Center, Mail Stop 499, Hampton, VA 23681-2199, United States

Submitted to the Journal of Computers and Fluids

* Tel.: +1 757 864 4611; fax: + 1 757 864 8469
E-mail address:m.d.salas@nasa.gov

On Problems Associated with Grid Convergence of Functionals

Manuel D. Salas* and Harold L. Atkins
NASA Langley Research Center, Mail Stop 499, Hampton, VA 23681-2199, United States

Abstract

The current use of functionals to evaluate order-of-convergence of a numerical scheme can lead to incorrect values. The problem comes about because of interplay between the errors from the evaluation of the functional, e.g., quadrature error, and from the numerical scheme discretization. Alternative procedures for deducing the order property of a scheme are presented. The problems are studied within the context of the inviscid supersonic flow over a blunt body; however, the problems and solutions presented are not unique to this example.

Keywords: Code verification, grid convergence, supersonic blunt-body, drag functional, error norm.

1. Introduction

Computational Aerodynamicists conduct most of their grid convergence studies by studying the behavior of solution functionals, e.g., drag, lift and moment coefficients, as the computational grids are refined. A basic assumption underlying the use of functionals is "that the *order* of the method applies globally as well as locally"[1]. Functionals are used for several reasons: first, their accurate evaluation is of intrinsic value; and second, they provide a means of determining convergence properties of a numerical scheme without looking directly at hundreds of thousands of field point values. Ideally, an error measure should be used to examine order properties of grid convergence studies; however, exact solutions are usually not available for flows of practical interest. Therefore, estimating convergence properties using functionals is frequently the only course of action available.

However, there are some subtle problems associated with the use of functionals for grid convergence studies, and if these problems are not recognized and resolved, the results that follow from the use of functionals can be very misleading. The alternative to using functionals, evaluating convergence from field point values, is equally marred with problems that are actually harder to resolve. It is the purpose of this paper to expose these problems, and where possible, suggest solutions.

There are many aspects of a numerical order-properties analysis that must be done correctly in order for the analysis to be reliable. Paramount among these are: that grid refinements must be uniform, preferably with grids sequences that are hierarchical; and the iterative methods used to solve the discrete equations on a given grid must be sufficiently converged, preferably with residuals reduced several orders of magnitude

* Tel.: +1 757 864 4611; fax: + 1 757 864 8469
E-mail address:m.d.salas@nasa.gov

below the solution error. Of course, this is complicated by the fact that the errors are not known a priori.

In this work, we focus strictly on problems associated with the use of functionals in grid refinement studies. We assume all other conditions required for a reliable order-property analysis are satisfied. In § 2, we present well known methods for computing the order properties from numerical results. The third section applies one of these methods to the results of a numerical simulation to demonstrate the problems that can occur when applying a standard method to the drag functional. Section 4 considers how the drag functional should behave relative to the behavior of the surface pressure. Sections 5 and 6 investigate one of the problems associated with the evaluation of the order-of-convergence from a functional. However, a second problem in the order property evaluation prevents a complete solution. Section 7 discusses the implementation of Richardson's extrapolation. In § 8, we identify the source of the difficulty that is inherent to functionals by applying the standard method to realistic models of the error. Section 9 introduces two techniques to reliably predict order properties from functionals, or closely related metrics that are applicable when the exact solution is not known. Conclusions and recommendations are made in § 10. The mathematician Paul Halmos when asked how he went about doing mathematics replied, "First, think of a question. Second, I look at examples, and then third, guess the facts"[2]. We followed Halmos' advice in this exposition.

2. Order-of-convergence of field point values

In Ref. [3] it was shown that in order to study the grid convergence of a multi-dimensional problem with a single grid size measure, h, the grid aspect ratio has to remain constant over all k-grids. In two dimensions, this condition is:

$$\chi = \frac{h_{y,k}}{h_{x,k}} = \text{constant}, \tag{1}$$

where h_x and h_y are the grid spacing in the x and y directions, respectively. If this necessary condition is satisfied, then convergence can be studied with a single measure, say $h_k = \sqrt{h_{x,k} \times h_{y,k}}$ and we say that the grids belong to the same family. We note that, as in [3], we assume that the physical domain is mapped to a uniformly spaced computational domain, and all mesh spacings refer to those in the computational domain.

Let h_k be a measure of the spacing on the k^{th} grid satisfying (1). Let U_e be the exact solution of some function of interest, and $u_{c,k}$ be its computed value on the k^{th} grid at some mesh point n,m with coordinates x,y. Let the grids be ordered such that $1 > h_1 > h_2 \ldots > h_{k-1} > h_k$. Furthermore, the grids are nested in the sense that if S_k is the set of points of the k^{th} grid, then $S_k \subset S_{k+1}$. In the asymptotic convergence range, a numerical algorithm of order p behaves like:

$$u_{c,k}(x,y) = U_e(x,y) + \alpha(x,y) h_k^{p(x,y)} + O\left(h_k^{\,p+1}\right), \quad k = 1,2,3,\ldots. \tag{2}$$

Although it is common to say that a numerical algorithm has an order-of-convergence p, the reality is that the order-of-convergence changes from point to point. This fact is recognized in the notation of (2), and it will become important in § 4. After dropping the higher-order terms in (2), at a given x, y there are three unknowns: U_e, α, and p. If the exact solution is known, then p can be readily determined from the solutions on two grids as follows:

$$z_e \equiv \frac{u_{c,k} - U_e}{u_{c,k+1} - U_e} = \left(\frac{h_k}{h_{k+1}} \right)^p \equiv \sigma^p . \tag{3}$$

A similar relation can be derived when the exact solution is not known, by using the solutions on a sequence of three grids.

$$z_k \equiv \frac{u_{c,k} - u_{c,k+1}}{u_{c,k+1} - u_{c,k+2}} = \frac{\left[h_k^p - h_{k+1}^p \right]}{\left[h_{k+1}^p - h_{k+2}^p \right]} = \frac{\left[1 - \left(h_{k+1} / h_k \right)^p \right]}{\left[1 - \left(h_{k+2} / h_{k+1} \right)^p \right]} \left(\frac{h_k}{h_{k+1}} \right)^p . \tag{4}$$

If

$$\frac{h_{k+1}}{h_{k+2}} = \frac{h_k}{h_{k+1}} = \sigma , \tag{5}$$

then the right hand side of Eq. (4) becomes identical to that of Eq. (3). In both cases, the order-of-convergence is given by

$$p = \log_\sigma (z) , \tag{6}$$

where $z = z_e$, or z_k. In what follows we will assume that condition (5) is satisfied[1].

Another common practice, used when the exact solution is not known, is to apply Eq. (3) using the solution on a very fine grid as a surrogate for the exact solution. However, this is an approximation that should be used with caution. Formally substituting this approximation into Eq. (2) and dropping the higher-order terms gives

[1] Note that condition (5) is different from (1).

$$z_{k,k\_ref} \equiv \frac{u_{c,k} - u_{c,k\_ref}}{u_{c,k+1} - u_{c,k\_ref}} = \frac{\left[h_k^p - h_{k\_ref}^p\right]}{\left[h_{k+1}^p - h_{k\_ref}^p\right]} = \frac{\left[1 - \left(h_{k\_ref}/h_k\right)^p\right]}{\left[1 - \left(h_{k\_ref}/h_{k+1}\right)^p\right]}\left(\frac{h_k}{h_{k+1}}\right)^p, \tag{7}$$

where k_ref denotes the very fine reference grid. If both $\frac{h_{k\_ref}}{h_{k+1}}$ and $\frac{h_{k\_ref}}{h_{k+2}} << 1$, then the right hand side of Eq. (7) approximately equals σ^p. There is not a single closed form solution that applies for all k, k_ref combinations, but we can examine some specific cases. Evaluating Eq. (7) for $k\_ref = k+2$ gives

$$z_{k,k\_ref} = \frac{\left[1 - \left(h_{k\_ref}/h_k\right)^p\right]}{\left[1 - \left(h_{k\_ref}/h_{k+1}\right)^p\right]}\left(\frac{h_k}{h_{k+1}}\right)^p = \frac{\left(1 - 1/\sigma^{2p}\right)}{\left(1 - 1/\sigma^p\right)}\sigma^p = \left(\sigma^p + 1\right). \tag{8}$$

Thus, if z_{k,k_ref} is naively substituted into Eq. (6) we get $\hat{p} = \log_\sigma(\sigma^p + 1)$, and the order-of-convergence is over predicted by $O\left(1/\sigma^p \log(\sigma)\right)$. Table 1 gives $\hat{p} = \log_\sigma(z_{k,k\_ref}(p))$ from Eq. (7) for $k\_ref = k+2$ and $k+3$ for a range of p. From this we see that Eq. (7) is a good approximation for $p > 2$ when $k\_ref = k+2$, and for $p > 1$ when $k\_ref = k+3$.

3. The blunt-body problem

We illustrate the problems that are associated with the use of functionals for the study of grid convergence rates with numerical results from the computation of a blunt-body in an inviscid supersonic stream. However, we emphasize that the problems we discuss are not unique to the blunt-body problem, indeed they are not unique to fluid mechanics, and may occur in any grid convergence study involving functionals. The blunt-body in a supersonic stream is a well understood problem. The problem is non trivial without being overly complicated, and has interesting flow physics features. Its solution by finite-difference methods dates back to the mid-sixties. A detailed review of the rich history of this problem as well as the physical properties of the flow can be found in Ref. [4]. The particular case we study is the Mach 6 flow of an inviscid gas over a circular cylinder. In our numerical implementation the problem is solved as a time dependent problem with the bow shock wave fitted as a boundary of the flow. By fitting the shock, the numerical scheme acts only on a smooth flow region, as illustrated by the pressure contours shown in Fig. 1. Thus the computation is limited to the layer bounded by the bow shock, the circular cylinder, the symmetry line ($\theta = 0^\circ$), and a supersonic out flow boundary imposed at some $\theta = \theta_{\max}$, see Fig. 1. The physical plane is transformed to a computational plane by the transformation,

$$Z = \frac{r - b(\theta)}{s(\theta,t) - b(\theta)},\ Y = \theta / \theta_{\max},\ T = t. \tag{9}$$

Here, r,θ and t are the radial, circumferential and time coordinates, respectively. In general, the body shape is defined by $b(\theta)$, for the case under study $b(\theta) = 1$. The shock wave shape is defined by $s(\theta,t)$ and is computed as part of the solution. In the computational plane, mesh points are uniformly distributed between the body and the shock and between the symmetry line and the outflow line. Let N be the number of mesh point intervals between the body and the shock, and let M be the number of mesh point intervals between the symmetry line and the outflow line. The details of the numerical scheme used to solve the Euler equations are described in Ref. [5], and previous results from this method can be found in Ref. [6]. In the present implementation, the Lax-Wendroff scheme described in Ref. [5] is replaced by the predictor-corrector scheme of MacCormack [7], formally a second order scheme. The code is written for the MATLAB® environment, which by default uses double precision, i.e., 64 bits on a 32 bits CPU. All the results presented were obtained on a laptop computer with a 3.2 GHz Pentium 4 CPU and 896 MB of RAM. Running times vary with grid size. For the grids used here, typical running times ranged from a few seconds to a few minutes.

Table 2 shows results obtained with a series of grids. Columns 5 and 6 display the inviscid drag coefficient computed with the trapezoidal rule (TR) and with Simpson's rule (SR). The drag coefficient order-of-convergence p computed using equations (4) and (6) is shown in the last two columns. The order-of convergence for $k = 3$ for TR and SR is showing a large discrepancy and both results are significantly greater than the formal order of the scheme. For $k = 4$, the drag coefficient is not monotone and the order-of-convergence evaluation fails. (Recall that the order-of-convergence for grid k depends on the solutions from grids k, k+1, and k+2).

In order to establish that there is reason to suspect these results, we consider the behavior of the error norm in total temperature. For this problem in the steady state, the total enthalpy, and hence the total temperature, is constant. The L_2 and L_∞ norms of the total temperature error are shown in Fig. 2 for grids $k = 3,4,5$ and 6. The first norm is an indication of convergence in the mean while the second indicates absolute convergence. The order-of-convergence based on the L_2 and L_∞ norms is 2.03 and 1.84, respectively. These are in fairly good agreement with the formal order of the scheme. Why then is the order-of-convergence of the drag functional misbehaving?

4. Order-of-convergence of functionals

To answer the last question, we first consider the following question: If the surface pressure converges with order $p(\theta)$, what should the expected order-of-convergence of the drag-functional be? To this end, let the computed surface pressure, P_c, normalized by the free stream pressure, P_∞, be given by

$$\frac{P_c}{P_\infty} = \frac{P_e}{P_\infty} + \alpha(\theta) h^{p(\theta)}, \tag{10}$$

where P_e is the exact surface pressure. Note that by design the computed pressure is in the asymptotic range and has an order-of-convergence $p(\theta)$. The sectional drag coefficient is defined by

$$C_d = \frac{1}{\frac{1}{2}\gamma M_\infty^2 A_{ref}} \int_0^{\theta_{\max}} \left(\frac{P_c}{P_\infty} - 1 \right) b(\theta) \cos(\theta) d\theta, \tag{11}$$

where γ is the ratio of specific heats, M_∞ is the free stream Mach number, $b(\theta) = 1$, and A_{ref} is a reference surface area, here taken as the projected plan-form area. Substituting (10) into (11) we find

$$C_d = \frac{1}{\frac{1}{2}\gamma M_\infty^2 A_{ref}} \left[\int_0^{\theta_{\max}} \left(\frac{P_e}{P_\infty} - 1 \right) \cos(\theta) d\theta + \int_0^{\theta_{\max}} \alpha(\theta) h^{p(\theta)} \cos(\theta) d\theta \right].$$

If $\alpha(\theta) h^{p(\theta)}$ is continuous on $\left[0, \theta_{\max}\right]$, then the second mean value theorem for integrals[8] tells us that there is a number $\xi \in \left[0, \theta_{\max}\right]$ such that

$$\int_0^{\theta_{\max}} \alpha(\theta) h^{p(\theta)} \cos(\theta) d\theta = \alpha(\xi) h^{p(\xi)} \int_0^{\theta_{\max}} \cos(\theta) d\theta.$$

Therefore,

$$C_d = C_{d,e} + \Gamma h^{p(\xi)}, \tag{12}$$

where

$$\Gamma = \frac{\alpha(\xi) \sin(\theta_{\max})}{\frac{1}{2}\gamma M_\infty^2 A_{ref}} = constant, \tag{13}$$

and $C_{d,e}$ is the exact value of the drag coefficient; therefore, the drag-functional converges with an order $p(\xi)$ representing a cosine-weighted average value of the surface pressure order-of-convergence on the interval $\left[0, \theta_{\max}\right]$. Note that if the surface pressure order-of-convergence is a constant, then the drag-functional converges at the same rate as the pressure. This result holds for all integral functionals in which the integrand depends linearly on the solution.

5. The problem with quadrature

If we assume that the pressure order-of-convergence behaves like the total temperature order-of-convergence, then the result just obtained is not consistent with the results of Table 2. Thus, we return to the question we asked at the end of § 3. To answer this question, we must look in some detail at the numerical integration of (11). The integral (11) is approximated by a quadrature taken over M equally spaced intervals on the surface of the cylinder,

$$\int_0^{\theta_{\max}} f d\theta \approx \sum_1^{M+1} a_i f_i .$$

The quadrature is an approximation to the exact integral with a leading error of order h^q,

$$\sum_1^{M+1} a_i f_i = \int_0^{\theta \max} f d\theta + O(h^q) + H.O.T.$$

If the quadrature is based on the extended TR, then

$$\sum_{TR} af = \frac{h}{2}\left(f_1 + 2f_2 \ldots + 2f_M + f_{M+1}\right) = \int_0^{\theta_{\max}} f d\theta + \frac{\theta_{\max}}{12} h^2 f'' + H.O.T.,$$

and if the quadrature is based on the extended SR, then

$$\sum_{SR} af = \frac{h}{3}\left(f_1 + 4f_2 + 2f_3 \ldots + 2f_{M-1} + 4f_M + f_{M+1}\right) = \int_0^{\theta_{\max}} f d\theta + \frac{\theta_{\max}}{180} h^4 f^{iv} + H.O.T.,$$

see Ref. [8] for more details on these and other quadrature rules. Therefore, for the extended TR we have from (12) that

$$C_d = C_{d,e} + \beta h^2 + \Gamma h^{\bar{p}} + O(h^{2+\bar{p}}), \tag{14}$$

where $\bar{p} = p(\xi)$ and $\beta = \theta_{\max} f''/12$. Introducing (14) in Eq. (4) we find

$$z = \frac{C_{d,e} + \beta h^2 + \Gamma h^{\bar{p}} - \left[C_{d,e} + \beta\left(\frac{h}{2}\right)^2 + \Gamma\left(\frac{h}{2}\right)^{\bar{p}} \right]}{C_{d,e} + \beta\left(\frac{h}{2}\right)^2 + \Gamma\left(\frac{h}{2}\right)^{\bar{p}} - \left[C_{d,e} + \beta\left(\frac{h}{4}\right)^2 + \Gamma\left(\frac{h}{4}\right)^{\bar{p}} \right]}, \tag{15}$$

where h is the coarse grid spacing. The drag-functional order-of-convergence is given by

$$\log_2(z) = \bar{p} + \log_2\left\{\frac{2^{\bar{p}-2}\beta\left[1-2^2\right]+h^{\bar{p}-2}\Gamma\left[1-2^{\bar{p}}\right]}{2^{2(\bar{p}-2)}\beta\left[1-2^2\right]+h^{\bar{p}-2}\Gamma\left[1-2^{\bar{p}}\right]}\right\}. \tag{16}$$

Equation (16) shows that there is interplay between the order-of-convergence of the drag-functional and the order-of-error of the TR. The left hand side of Eq. (16) provides a good estimate of the order-of-convergence of the drag-functional only if the log term is small, i.e. the log term is the error. This term vanishes only if $\bar{p}=2$ or $\beta=0$. For $\bar{p}\neq 2$ if the ratio $\beta/\Gamma<0$, the estimated order-of-convergence of the drag-functional can be singular. In the limit $h\to 0$ we have

$$\lim_{h\to 0}\log_2(z) = \begin{cases} 2, \ if\ \bar{p}\geq 2, \\ \bar{p}, \ if\ \bar{p}<2. \end{cases}$$

The behavior of (16) is depicted in Fig. 3 for four different values of β/Γ as a function of $\bar{p}$ for $h=.01$. For β/Γ positive, the estimated drag-functional order-of-convergence is greater than its true value as long as $\bar{p}<2$. For $\bar{p}>2$, the estimated drag-functional order-of-convergence levels off at a value of two. Thus, the TR cannot be used to predict order-of-convergence greater than two. Any result that shows order-of-convergence greater than two based on the TR (or greater than four based on the SR) is in error. The latter occurs when β/Γ is negative resulting in singular behavior of the log term, as shown on the right hand panel of Fig. 3, and very low or very high values of the predicted order-of-convergence. This is the cause for the puzzling results of Table 2 for both the TR and SR. This problem has gone unnoticed in the literature. For example, in Ref. [9] a resistance coefficient, drag-like functional, is computed using the TR leading to order-of-convergence of 4.4 and 9.5 for a formally second order accurate scheme. The authors fail to provide any explanation for this behavior.

The extension of (16) to other quadrature rules is

$$\log_2(z) = \bar{p} + \log_2\left\{\frac{2^{\bar{p}-q}\beta_q\left[1-2^q\right]+h^{\bar{p}-q}\Gamma\left[1-2^{\bar{p}}\right]}{2^{2(\bar{p}-q)}\beta_q\left[1-2^q\right]+h^{\bar{p}-q}\Gamma\left[1-2^{\bar{p}}\right]}\right\},\ q=\begin{cases} 2\ trapezoidal\ rule, \\ 4\ Simpson's\ rule, \\ 6\ Bode's\ rule, \end{cases} \tag{17}$$

and

$$\lim_{h\to 0}\log_2(z) = \begin{cases} q, \ if\ \bar{p}\geq q, \\ \bar{p}, \ if\ \bar{p}<q. \end{cases}$$

6. How to eliminate the quadrature problem

The quadrature problem and its solution can be found by studying Eq. (15). Consider the numerator of (15). The numerator is the drag coefficient of the medium grid minus

the drag coefficient of the coarse grid. The medium grid has a quadrature error of order $(h/2)^2$, while the coarse grid has a quadrature error of order h^2. These errors do not cancel out and their interplay with the algorithmic error causes the problems illustrated on Fig. 3. The solution is to implement the quadrature in such a way that the quadrature errors of the medium and coarse grids cancel. To do this, we evaluate the medium grid quadrature using an h interval, i.e., we use only every other point of the medium grid. The same idea is applied to the denominator. For the denominator, the quadrature for the fine and medium grids can be evaluated either using h or $h/2$. It is only important that both be evaluated using the same spacing. When this is done, the log term in (16) vanishes and we are left with:

$$\log_2(z) = \bar{p}\,.$$

Now there is no formal dependence on h or on β/Γ. With this modification, the TR can be used to determine the functional order-of-convergence of schemes with $\bar{p} > 2$.

Applying this modification to the drag calculation of the blunt-body problem, we obtain the results shown in Table 3. Note that now the TR and SR give consistent values for $\bar{p}$. What is interesting about this method is that by implementing a less accurate quadrature we obtain a more consistent evaluation of the order-of-convergence. We have eliminated the quadrature error from the order-of-convergence evaluation. However, the new results indicate very fast convergence. Are these results correct? We answer this question in § 8 when we look in detail at what happens when the solution is near but not in the asymptotic range.

7. Richardson extrapolation

Since we have introduced a quadrature of the drag based on the coarse grid points, a question remains: How should we do a Richardson extrapolation of the drag? That is, should the Richardson extrapolation of the drag be based on a quadrature using all available grid points or only the coarse grid points? The Richardson extrapolation is given by

$$C_{d,RE} = \frac{C_{d,k+1,j+1} 2^{\bar{p}} - C_{d,k,j}}{2^{\bar{p}} - 1}, \tag{18}$$

where

$$C_{d,k,j} = C_{d,e} + \beta h_j^q + \Gamma h_k^{\bar{p}}\,. \tag{19}$$

The second term on the right hand side is the quadrature error, and the last term is the algorithmic error. Introducing (18) into (19) we get

$$C_{d,RE} = C_{d,e} + \frac{\beta\left(2^{\bar{p}} h_{j+1}^{q} - h_j^{q}\right) + \Gamma\left(2^{\bar{p}} h_{k+1}^{\bar{p}} - h_k^{\bar{p}}\right)}{2^{\bar{p}} - 1}. \tag{20}$$

Since $h_{k+1}^{\bar{p}} = \left(h_k / 2\right)^{\bar{p}}$, Eq. (20) reduces to

$$C_{d,RE} = C_{d,e} + \frac{\beta\left(2^{\bar{p}} h_{j+1}^{q} - h_j^{q}\right)}{2^{\bar{p}} - 1}. \tag{21}$$

If we use the new quadrature method, then $h_{j+1} = h_j$ and Eq. (20) reduces to

$$C_{d,RE} = C_{d,e} + \beta h_j^{q}. \tag{22}$$

If we use the usual quadrature method, then $h_{j+1} = h_j / 2 = h_k / 2$ and Eq. (21) reduces to

$$C_{d,RE} = C_{d,e} + \frac{2^{\bar{p}-q} - 1}{2^{\bar{p}} - 1} \beta h_k^{q}. \tag{23}$$

As long as $q > 0$ the factor $\left|\left(2^{\bar{p}-q} - 1\right) / \left(2^{\bar{p}} - 1\right)\right| < 1$, therefore, since $h_j \geq h_k$, Eq. (23) has a smaller error than Eq. (22). Hence, for Richardson extrapolation of the drag we should use the most accurate drag values available, i.e. those from Table 2.

8. Identifying the source of the anomalous behavior.

This section will examine an additional cause of the anomalies that have been observed in the previous sections. We do this by assuming several reasonable models of the error, applying the method given in § 2, and observing the result.

It is important to keep in mind that the actual error in any numerical solution on any given grid contains a full hierarchy of errors that are unknown, but are generally assumed to be of the form

$$u_{c,k} = U_e + \sum_{n=p}^{\infty} \alpha_n h_k^{n}, \tag{24}$$

where p is the unknown order of the numerical method. A primary goal of the verification process is to either 1) verify that a method has a particular design order, or 2) to determine the order of a method when it is applied outside of its ideal design space (non-smooth grids, discontinuous solutions, etc). The standard method for deducing order properties from grid convergence, described earlier, is obtained by fitting a single error mode of the form $u_{e,k} = U_e + \alpha h_k^{p}$ to the actual error. For sufficiently small h, the actual error is dominated by the lowest order term, and the single mode model provides a good

fit and an accurate prediction of the order-of-convergence. However, for larger h above this asymptotic region, multiple error modes are competing in the actual error, and their projection onto a single mode can be erroneous. This can be clearly illustrated by assuming several slightly more realistic models for the error and observing the result of applying equations (4) and (6).

The first model of the error, beyond the standard of Eq. (2), is simply one that contains two modes

$$u_{c,k} = U_e + \alpha_1 h_k^{p_1} + \alpha_2 h_k^{p_2}. \tag{25}$$

Substituting into Eq. (4), we obtain an equation equivalent to Eq. (16) and conclude that the result depends only on p_1, p_2 and the ratio $\alpha_2/\alpha_1 \equiv \alpha$. Plotting the predicted order of convergence as a function of p_1 with p_2 fixed for $\alpha < 0$ or > 0 produces the curves similar to those shown previously in Figures 3 (a) and (b), respectively. Here, it is clear that the singular behavior occurs because the error measure briefly goes through zero as the competing error terms cancel each other.

However, we are more interested in the behavior of the standard method as the grid is refined. Figures 4 (a) and (b) illustrate the error and the order predicted by equations (4) and (6) for the case where $p_1 = 1$, $p_2 = 2$, and $\alpha = \pm 1$. (Note that larger or smaller values of α do not alter the general behavior of the curves.) The error is smooth when α is greater than zero, and the predicted order-of-convergence transitions smoothly from $p = 2$ to $p = 1$ as h decreases. However, when α is less than zero, the denominator of Eq. (4) passes through zero and the predicted order properties range from $\pm\infty$, taking every value but those between one and two. Furthermore, there is a region where the numerator and denominator of Eq. (4) are of opposite sign. We note that $\log(z) = \log(|z|) + i\pi$ when $z < 0$. The plots show $\log(|z|)$ instead of p, and the regions of $z < 0$ are denoted with a dash-dot line segment.

Since functionals are not proper error norms, they are likely to be the result of many competing and canceling terms. From the above illustration, it is clear that such functionals will not converge in a monotone manner, and this will lead to the zeros in the numerator of z and regions where z is negative, as observed above. To further support this notion, we examine a spatial error model that will allow the order to be predicted from either a functional of the solution or from a proper error norm of the solution. The spatial error model is motivated by a Fourier stability analysis, which in essence, provides an exact eigensolution to a discrete problem on a simple domain. Assume for the moment a simple advection problem in which the exact solution is given by $U_e(x,t) = G(x-t)$, and a numerical solution with both amplitude and phase error terms that is modeled by

$$u_{c,k} = (1 + \alpha_1 h^{p_1}) U_e(x, t(1 + \alpha_2 h^{p_2})). \tag{26}$$

To enable a computation, let $G(x) = \cos(\pi x)$; and to avoid fortuitous cancellation that can occur on periodic domains, let the physical domain be the fractional period $0 < x < 7/4$. In this exercise, we will define the functional, denoted by I_k, to be the integral of the solution approximated by the trapezoidal rule. We will also define and examine several error metrics. The first is the true error defined as a norm of $u_{c,k} - U_e(x, t)$, and denoted by $\varepsilon_{k,e}$. The second error measure, denoted $\varepsilon_{k,k+1}$, is a norm of $u_{c,k} - u_{c,k+1}$, which is simply the difference between solutions on sequential grids. The last error measure, denoted ε_{k,k_ref}, is a norm of $u_{c,k} - u_{c,k\_ref}$, which is the difference between the current solution and that on a very fine reference grid. Computations performed using L_1, L_2 and L_∞ norms gave similar results. Therefore, the following figures and tables give only the results using the L_2 norm, and with error model parameters $p_1 = p_2 - 1 = 2$ and $\alpha_1 = 2\alpha_2 = 1$. Table 4 shows errors and convergences rates for the integral functional and the 3 different error norms. The convergence rate of the functional, I_k, is erratic. The entry in "()" indicates that z is negative for that case and the value given is $\log(|z|)$. However, the convergence of the L_2 norm of the actual error is between 2 and 3 on the coarse grid and approaches 2, as the grid is refined. The L_2 norm of the local relative error, $\varepsilon_{k,k+1}$, is similar to the real error. The L_2 norm of the error with respect to a fine grid reference solution, ε_{k,k_ref}, initially trends like the exact error, but asymptotes to ~2.3 on the finest grid, as expected. It is interesting to note that since the norm of the quantities $\varepsilon_{k,e}$, $\varepsilon_{k,k+1}$, and ε_{k,k_ref} are evaluated on grid k, the issues associated with quadrature identified earlier are naturally eliminated.

Figures 5 (a), (b) and (c) give the solution, the absolute value of the relative error $\left|u_{c,k} - u_{c,k+1}\right|$, and local values of z computed from the relative error, respectively. The solutions on the three finest grids are indistinguishable from each other. The relative error appears well behaved. The local minimums in the absolute value of the relative error are where the relative error crosses through zero and changes sign. The shift in the location of the zero will have little effect on any norm of this function. However, any integrated functional will be strongly influenced by the cancellation that occurs between the positive and negative contributions that exist on either side of each zero point. The shifts are quite large even though, in this case, the phase error is the higher order term in the error model. The local value of z is simply the ratio of two adjacent relative error curves (with the sign restored). The zeros of the relative error cause z to approach $\pm\infty$ and the shifts in the zero points cause regions of negative z, just as was previously seen in the two-mode model.

9. Accurate alternative procedures for deducing order property of functionals

The previous discussion suggests two approaches to reliably and accurately predict the order-of-convergence of a functional. The first approach recognizes that the fundamental problem is that functionals may have cancellation of various order error terms embedded within them, and constructs a positive norm of a related quantity that

minimizes this cancellation. The second approach accepts the presence of the embedded cancellation, and fits the functional to a higher-order multi-mode error model of the functional.

9.1 Norm based on relative error

Assuming the functional is an integrated quantity, the first approach is to base the prediction of the order-of-convergence on a norm of the relative error of the integrand of the functional. That is, if the functional F is defined as

$$F \equiv \int_{\Omega} f d\Omega, \tag{27}$$

then the order-of-convergence is computed from a norm of $f_k - f_{k+1}$. Applying this technique to the simulation results described in § 3 produces the L_2 error norms and order-of-convergence rates given in Table 5. Here, the convergence rates are between 2 and 3 as we would expect; however, the rate is not smoothly (and monotonically) transitioning from 3 to 2 as was observed in the model case. The cause of this behavior is discussed below. Figure 6a shows the spatial distribution of the relative error metric. As in the earlier case with the spatial model for the error, it is clear that shifts in the zero point will cause large cancellations within the integrated functional, but will have little effect on the norm of this metric. Another feature revealed in Fig. 6b is that the solution in the down stream region is converging at a higher rate than is the upstream region. This indicates that the error in this region is dominated by higher-order truncation error terms. The norm of the relative error results in an average over the domain, and this produces the non-monotone transition in the convergence rate.

This last result reveals an important distinction between the spatial error model of § 8 and typical simulation results. The model has a single wavelength that is uniformly resolved; thus, the convergence transitions smoothly from $p+1$ to p. However, in real simulation results, the flow is not uniform and the grid distribution is not ideal; thus some regions may be well resolved and others are not and higher-order error terms may locally dominate the convergence. In this case, the norm of the relative error will produce a representative average convergence rate, but may not converge monotonically.

9.2 Increasing the order of the error model

The second method accepts that the functional has embedded cancellations due to higher order effects, and resolves the issue by increasing the order of the error model. Using a sequence of four grids, and assuming $p_1 = p_2 - 1 \equiv p$, it is possible to directly fit the two-mode model to the local solution or to a numerical functional.

$$u_{c,k} = U_e + \alpha_1 h_k^p + \alpha_2 h_k^{p+1} \tag{28}$$

After some manipulation, we find

$$(2z-1)(z-1)(2\Delta u_{k-1,k}z^2 - 3\Delta u_{k-2,k-1}z - \Delta u_{k-2,k-3}) = 0, \tag{29}$$

where $\Delta u_{i,j} = u_{c,i} - u_{c,j}$.

The root of interest is given by:

$$z = \frac{3\Delta u_{k-2,k-1} - \sqrt{9\Delta u_{k-2,k-1}^2 + 8\Delta u_{k-1,k}\Delta u_{k-2,k-3}}}{4\Delta u_{k-1,k}} \tag{30}$$

Figure 7 shows the order-of-convergence of the surface pressure from the blunt-body solutions on grids $k = 4,5,6$ using the standard error representation, i.e., Eq. (2). The singular behavior occurs at the body sonic point ($\theta \approx 46^\circ$). This is to be contrasted with Fig. 8 where the order-of-convergence is evaluated using the higher order analysis described here with the surface pressure from the blunt-body solutions on grids $k = 3,4,5,6$. The new result is near 2 over most of the domain and the singular behavior is eliminated. Similarly, if we compute the drag order-of-convergence based on the higher order analysis, we obtain the results shown in Table 6 which are consistent with the formal analysis of § 4.

We can also solve for the other unknowns, U_e, α_1 and α_2:

$$\alpha_2 = \frac{8(\Delta u_{k+1,k} - z\,\Delta u_{k+3,k+4})z}{3(2z-1)}, \tag{31}$$

$$\alpha_1 = \frac{\Delta u_{k+1,k} + (2z-1)\alpha_2}{2(z-1)}, \tag{32}$$

and $\hat{U}_e = u_{c,k} - \alpha_1 h_k^p - \alpha_2 h_k^{p+1}$. The "hat" denotes that this is not the exact solution but a prediction of it. This prediction of the exact solution serves as a high-order version of the standard Richardson extrapolation. Of particular interest is the ratio of the $p+1$ and p terms:

$$\frac{\alpha_2 h_k^{p+1}}{\alpha_1 h_k^p} = \alpha_2 h / \alpha_1.$$

This ratio gives an indication of whether or not a region is dominated by the lower order terms, and thus, is in the asymptotic range. Figure 9 gives this ratio for grids $k = 3,4,5,6$. Here we see that the downstream region is dominated by the higher order terms on all but the finest grid. This agrees with the results in the previous section in which the local

convergence of the relative error measure also converged at a higher rate in the downstream region.

Applying the higher-order fit given by equations (30) and (6) directly to the drag functionals given in Table 2 produces the order-of-convergence shown in Table 7. The higher order model predicts that the drag functional convergence rate approaches 1.7 as the mesh is refined.

10. Conclusions and recommendations

With the increased reliance in both science and engineering on the numerical solution of partial differential equations, the subject of code verification has become increasingly significant and prominent. An important element of code verification is the study of grid convergence and the determination of order-of-convergence. Most studies today of this subject have been at best superficial and in many cases painfully inadequate. A brief survey of the many papers presented in the AIAA drag prediction workshop series [10] should suffice in establishing this observation. This paper is an attempt to reverse this trend by first highlighting a series of problems that exist in the standard order-of-convergence analysis, particularly as it relates to the evaluation of functionals, and second by providing a number of solutions and workarounds to these problems. These problems are certainly important, however, we believe that the most important message from this work is that order-of-convergence studies are not trivial exercises and their proper execution requires a high degree of control over grid properties and the capability of systematically performing many levels of grid refinement. Highly complex problems, such as those used in the drag prediction workshop series, are just not good candidates for these studies.

It is important to distinguish between a code verification effort and an effort to determine if a particular solution to a specific problem is sufficiently accurate for some intended use. The two tasks are very different. A rigorous grid convergence and order-of-convergence study can aid in determining if an algorithm has been implemented correctly. However, such a rigorous study requires grids of the same family and grid refinements that are uniform, preferably with grids sequences that are nested, as defined in § 2. For steady-state problems, iterative or time convergence should be obtained with residuals reduced several orders of magnitude below the solution spatial error. In the second task of evaluating solution accuracy, limited time and resources often lead to compromising one or more attributes of a rigorous study. While non-uniform mesh refinement may lead to some improvement in the solution, especially when performed by an expert, order-of-convergence properties computed from non-uniform refinements or ill-converged solutions sets are meaningless. The common practice of performing such order-of-convergence predictions should be avoided.

Whenever possible, error norms should be used to establish the order-of-convergence. If functionals are used, first an analysis should be performed to establish the formal dependence of the functional on the order-of-convergence of its functions, and if the functional is an integral, then care must be exercised to avoid the interplay of quadrature

accuracy errors and order-of-convergence errors. Integral functionals, such as lift or drag, are subject to cancellation effects that can greatly delay the onset of the asymptotic convergence regime. These cancellation effects are strongly influenced by the higher-order contributions to the error. The order-of-convergence computed from a norm of the relative error of the solution or the functional integrand is effective in estimating an average order-of-convergence, is naturally immune to quadrature errors, and can provide some insight into local convergence behavior of a solution. The higher order analysis developed in § 9.2 is the best way to evaluate if the grids used are within the asymptotic range and to establish if more levels of grid refinement are needed to reach the asymptotic range. It should be part of any rigorous grid convergence study.

Acknowledgment

Four decades ago, Gino Moretti shared his blunt-body code with the first author, along with his love of fluid mechanics, and his many insights into numerical analysis. The code has survived rigorous verification tests and the rigor of the passage of time. For this and more, the first author is forever grateful.

References

[1]Roach, PJ. Verification and validation in computational science and engineering. New Mexico: Hermosa Publishers; 1998, p. 111.
[2]Ewing, JH and Gehring, FW. Paul Halmos: Celebrating 50 years of mathematics. Springer-Verlag, New York, 1991.
[3]Salas, MD. Some observations on grid convergence. Computers & Fluids, **35**, 2006, 688-692.
[4]Rusanov, VV. A blunt body in a supersonic stream. Ann. Rev. Fluid Mech., 1976, 377-404.
[5]Moretti, G. and Abbett, M. A time-dependent computational method for blunt body flows. AIAA J., **4**, 1966, 2136-2141.
[6]Moretti, G. Inviscid blunt body shock layers: two-dimensional symmetric and axisymmetric flows. Polytechnic Inst. of Bklyn.; PIBAL Report No. 68-15,1968.
[7]MacCormack, RW. The effect of viscosity in hypervelocity impact cratering. AIAA Paper No. 69-354, 1969.
[8]Hildebrand, FB. Introduction to numerical analysis. 2nd Ed. Dover Publications, Inc.; 1987.
[9]Wilson, RV, Stern, F, Coleman, HW, Paterson, EG. Comprehensive approach to verification and validation of CFD simulations – Part 2: application for RANS simulation of a cargo/container ship. ASME J. Fluids Eng., **123**, 2001, 803-810.
[10]Hemsch, MJ, Morrison, JH. Statistical analysis of CFD solutions from 2nd drag prediction workshop. AIAA Paper No. 2004-556, 2004.

k_ref \ p	1	2	3	4
k+2	1.58	2.32	3.17	4.09
k+3	1.22	2.07	3.04	4.00

Table 1. Predicted order-of-convergence, $\hat{p}$, when a fine grid solution is used as a surrogate for the exact solution.

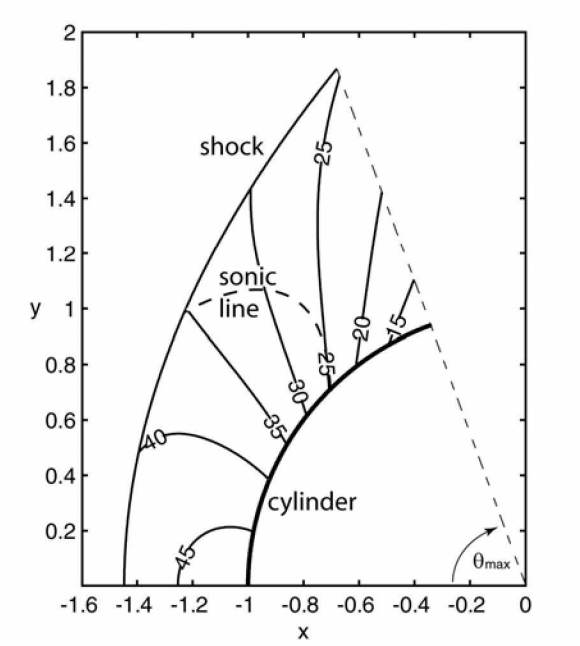

Figure 1. Details of the computed pressure field for a Mach 6 inviscid flow past a circular cylinder, $\theta_{\max} = 70^{\circ}$.

k	N	M	h	C_d TR	C_d SR	p TR	p SR
1	6	10	0.1290994	1.8755919	1.8767669	1.92	1.90
2	12	20	0.0645497	1.8706109	1.8709412	2.73	2.58
3	24	40	0.0322748	1.8692925	1.8693766	4.03	3.25
4	48	80	0.0161374	1.8690942	1.8691147	—	—
5	96	160	0.0080687	1.8690821	1.8690872		
6	192	320	0.0040343	1.8690859	1.8690872		

Table 2. Mach=6 cases investigated. C_d is the drag coefficient and p is its order-of-convergence. For the three finest grids the trapezoidal rule (TR) computed drag is not monotone and for both rules the computed drag exhibits super-convergence.

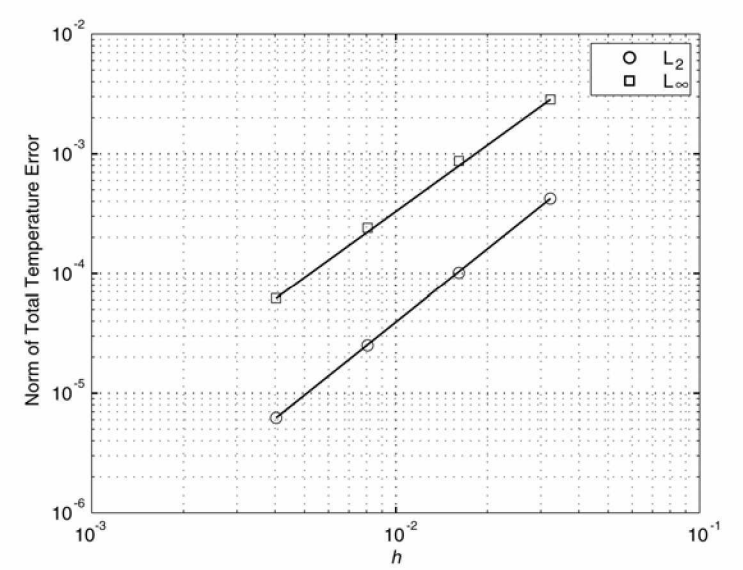

Figure 2. L_2 and L_∞ of total temperature error for entire shock layer, based on results from grids $k = 3,4,5,6$. In each case the line connects the data points for $k = 3$ and $k = 6$. The slope of L_2 is 2.03, that of L_∞ is 1.84.

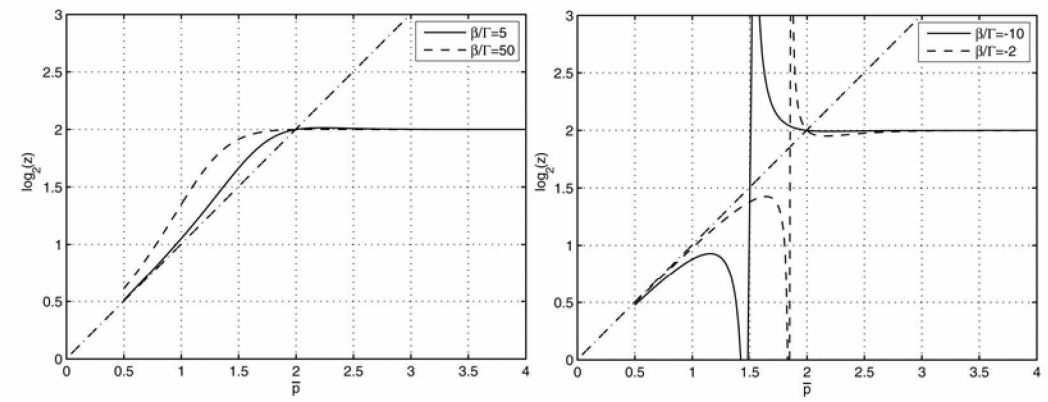

Figure 3. The figure shows the estimated order-of-convergence of the drag-functional as a function of the true order-of-convergence of the drag-functional. The deviation from the diagonal line indicates the error introduced by the TR quadrature. All the results shown are for $h = .01$.

k	C_d TR	C_d SR	$\overline{p}$ TR	$\overline{p}$ SR
1	1.8755919	1.8767669		
2	1.8696197	1.8709503		
3	1.8680572	1.8693560	1.93	1.87
2	1.8706109	1.8709412		
3	1.8690402	1.8693679		
4	1.8687875	1.8691135	2.63	2.63
3	1.8692925	1.8693766		
4	1.8690328	1.8691145		
5	1.8690056	1.8690872	3.26	3.26
4	1.8690942	1.8691147		
5	1.8690669	1.8690873		
6	1.8690668	1.8690872	8.15	7.75

Table 3. Drag coefficient and its order-of-convergence evaluated using the coarse-grid spacing of each grid subset in the quadrature rules.

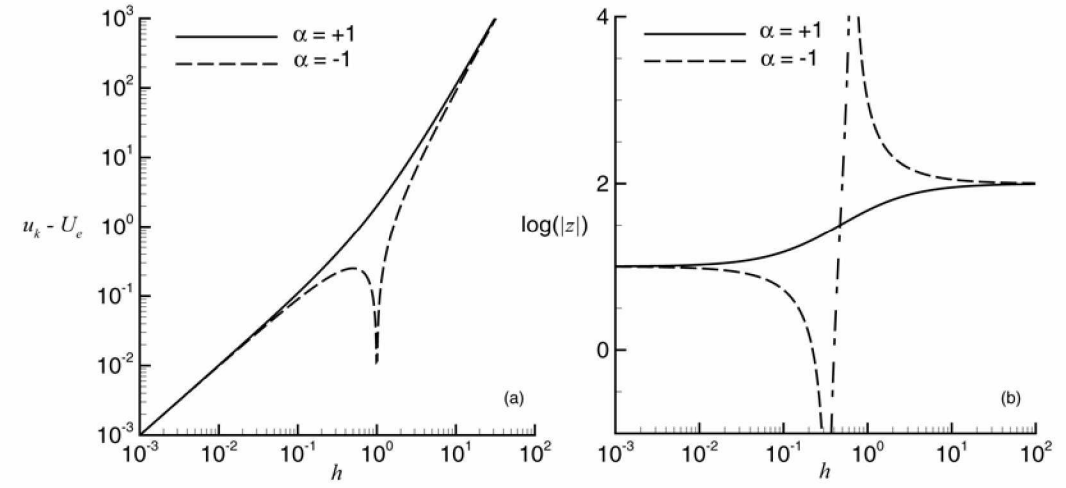

Figure 4. Error and $\log(|z|)$ for two mode model. The dash-dot line segment denotes the region where $z<0$ for the $\alpha=-1$ case.

	Metric				Convergence of metric			
h/h_0	I_k	$\varepsilon_{k,e}$	$\varepsilon_{k,k+1}$	ε_{k,k_ref}	I_k	$\varepsilon_{k,e}$	$\varepsilon_{k,k+1}$	ε_{k,k_ref}
1.0	-0.18641	1.56764	1.47571	1.56784	(4.73)	2.63	2.81	2.63
0.5	0.25319	0.24594	0.20292	0.24546	0.99	2.41	2.50	2.42
0.25	0.23664	0.04770	0.03701	0.04708	1.81	2.12	2.16	2.19
0.125	0.22830	0.01098	0.00833	0.01033	1.92	2.04	2.05	2.36
0.0625	0.22592	0.00267	0.00201	0.00201	—	—	—	—
0.03125	0.22529	—	—	0.0	—	—	—	—

Table 4. Error metrics and convergence rates for spatial error model. The convergence rate of the functional, I_k, is erratic; however, the convergence of the L_2 norm of the actual error is between 2 and 3 on coarse grid and approaches 2 as the grid is refined. The L_2 norm of the local relative error, $\varepsilon_{k,k+1}$, is similar to the real error. The L_2 norm of the error with respect to a fine grid reference solution, ε_{k,k_ref}, initially trends like the exact error, but asymptotes to 2.3 on the finest grid, as expected. The entry in "()" indicates that z is negative for that case, and the value given is $\log(|z|)$.

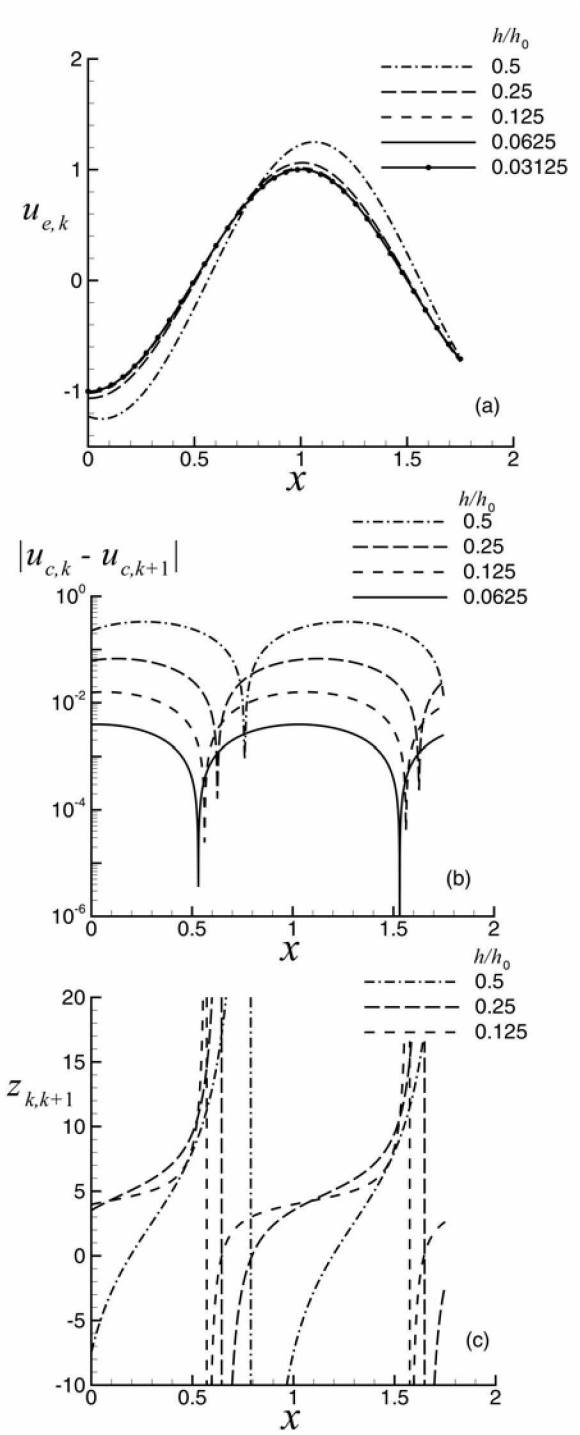

Figure 5. Spatial distributions produced by the spatial error model on a sequence of grids for (a) the solution $u_{c,k}$, (b) the absolute value of the relative error $\left|u_{c,k}-u_{c,k+1}\right|$ and (c) the local values of z computed from the relative error. The solutions on the three finest grids are indistinguishable from each other. The local minimums in the relative error are where it crosses through zero and changes sign. Integrated quantities will be strongly influenced by the cancellation that occurs between the positive and negative contributions that exist on either side of each zero point. The local value of z is simply

the ratio of two adjacent relative error curves (with the sign restored). The zeros of the relative error cause z to approach $\pm\infty$ and the shifts in the zero points cause regions of negative z.

h/h_0	Error metric $\varepsilon(P)_{k,k+1}$	$\varepsilon(P)_{k,k\_ref}$	Convergence of error metric	
1.0	30.49477	38.661806	2.13	2.20
0.5	6.957833	8.4258712	2.43	2.47
0.25	1.290086	1.5198184	2.62	2.65
0.125	0.209531	0.2418774	2.54	2.75
0.0625	0.035942	0.0359420	—	—

Table 5. Norms of error metrics of pressure and their convergence rates for the numerical simulation results from § 3.

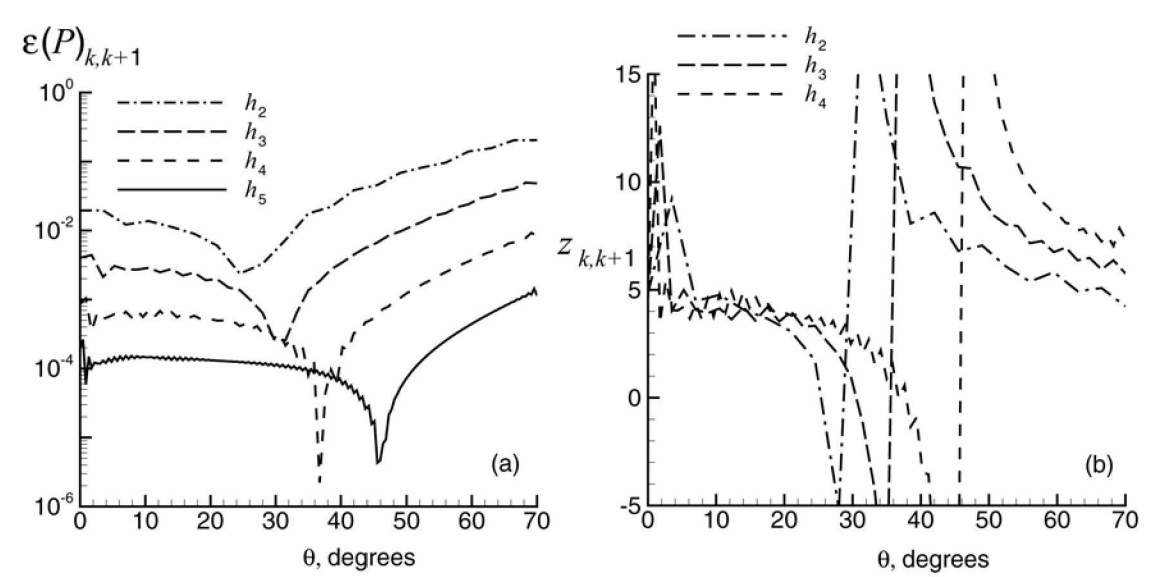

Figure 6. Spatial distribution of (a) the relative error, and (b) the local value of z, from the numerical simulation results of § 3. The rightward shift in the position of the zero point as the grid is refined would result in large cancellation effects in the integration of the (signed) functional, but would have little effect on any norm of the quantity.

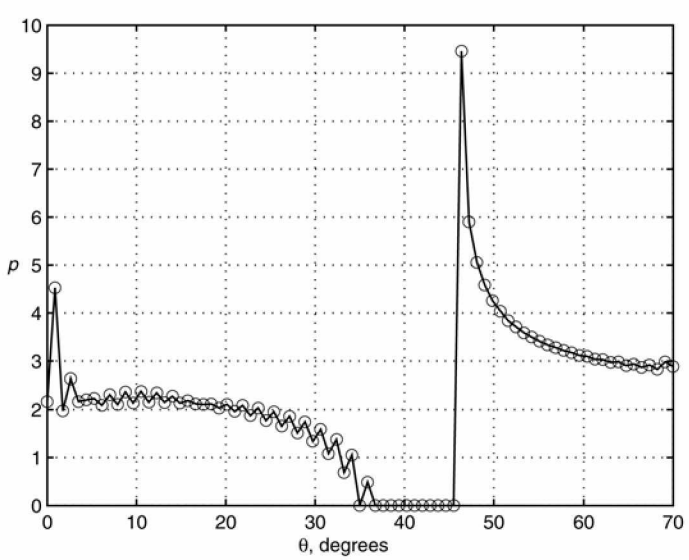

Figure 7. Surface pressure order-of-convergence based on standard analysis using blunt-body results on grids $k = 4,5,6$. Singular behavior near $\theta \approx 46^{\circ}$ corresponds to body sonic point.

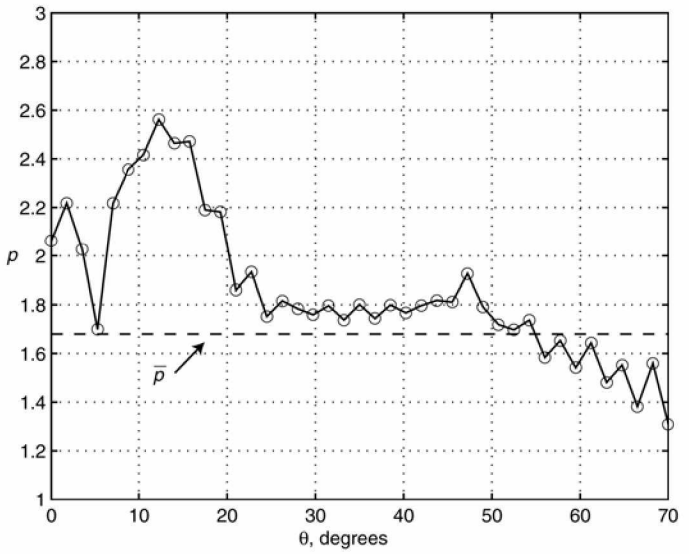

Figure 8. Surface pressure order-of-convergence based on higher order analysis using blunt-body results on grids $k = 3,4,5,6$. Dashed line indicates drag order-of-convergence from Table 6.

k	C_d TR	C_d SR	$\bar{p}$ TR	$\bar{p}$ SR
3	1.8692925	1.8693766	1.674	1.681
4	1.8690328	1.8691145		
5	1.8690056	1.8690872		
6	1.8690056	1.8690871		

Table 6. Drag order-of-convergence using higher order method and coarse grid interval, $k = 3$, for quadrature rules.

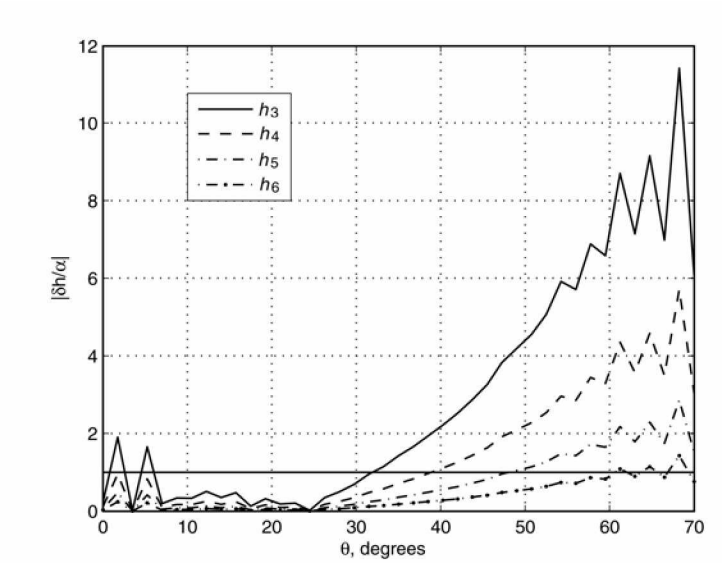

Figure 9. Ratio $\alpha_2 h/\alpha_1$ to for grids $k = 3,4,5,6$. In asymptotic range this ratio should be less than one.

k	C_d TR	C_d SR	$\bar{p}$ TR	$\bar{p}$ SR
1	1.87559192503	1.87676688882	0.56	0.57
2	1.87061085848	1.87094124452	1.30	1.26
3	1.86929251463	1.86937661231	1.70	1.67
4	1.86909418267	1.86911465349		
5	1.86908208985	1.86908715853		
6	1.86908589935	1.86908716957		

Table 7. Order-of-convergence predicted by the 4-grid high order fit applied to the drag functional data given in Table 2.