JOURNAL OF AIRCRAFT
Vol. 46, No. 3, May–June 2009

Grid Quality and Resolution Issues from the Drag Prediction Workshop Series

Dimitri J. Mavriplis*
University of Wyoming, Laramie, Wyoming 82071
John C. Vassberg†
The Boeing Company, Huntington Beach, California 92647
Edward N. Tinoco†
The Boeing Company, Seattle Washington 98124
Mori Mani†
The Boeing Company, St. Louis, Missouri 63301
Olaf P. Brodersen‡ and Bernhard Eisfeld‡
DLR , German Aerospace Center, 38108 Braunschweig, Germany
Richard A. Wahls§ and Joseph H. Morrison¶
NASA Langley Research Center, Hampton, Virginia 23681
Tom Zickuhr** and David Levy**
Cessna Aircraft Company, Wichita, Kansas 67218
and
Mitsuhiro Murayama††
Japan Aerospace Exploration Agency, Chofu, Tokyo 182-8522, Japan

DOI: 10.2514/1.39201

The drag prediction workshop series, held over the last six years and sponsored by the AIAA Applied Aerodynamics Committee, has been extremely useful in providing an assessment of the state of the art in computationally based aerodynamic drag prediction. An emerging consensus from the three-workshop series has been the identification of spatial discretization errors as a dominant error source in absolute as well as incremental drag prediction. This paper provides an overview of the collective experience from the workshop series regarding the effect of grid-related issues on overall drag prediction accuracy. Examples based on workshop results are used to illustrate the effect of grid resolution and grid quality on drag prediction, and grid convergence behavior is examined in detail. For fully attached flows, various accurate and successful workshop results are demonstrated, and anomalous behavior is identified for a number of cases involving substantial regions of separated flow. Based on collective workshop experiences, recommendations for improvements in mesh generation technology that have the potential to impact the state of the art of aerodynamic drag prediction are given.

I. Introduction

OVER the last six years, the AIAA Applied Aerodynamics Committee has sponsored three Drag Prediction Workshops (DPW), with the aim of assessing the state of the art of current computational fluid dynamics solvers at predicting absolute and incremental drag changes on generic transonic transport aircraft configurations. The first workshop, DPW I, held in Anaheim, California, in June 2001, employed the DLR-F4 wing-body geometry [1,2] due to 1) its simplicity and relevance to modern transonic transport aircraft design studies and 2) the availability of publicly released experimental data. The second workshop, DPW II [3], held in Orlando, Florida, in June 2003, employed the slightly different DLR-F6 geometry, which included a wing-body configuration and a wing-body-nacelle-pylon configuration, both of which have been tested experimentally. DPW II thus offered the possibility to study increments in drag between two similar configurations, as would be typically done in a design study. Results from both the first and second workshops showed considerable scatter in terms of drag prediction and agreement with experimental data. This scatter was largely attributed to the fact that the test cases contained substantial regions of turbulent separated flow, which are known to be very sensitive to turbulence model fidelity, numerical dissipation, and discretization errors.

For the third and most recent workshop, DPW III [4], held in San Francisco, California, in June 2006, a wing-body fairing was designed [5] and added to the DLR-F6 wing-body configuration to suppress flow separation in the wing-body juncture region. The baseline DLR-F6 wing-body configuration was also retained as a test case, providing a bridge between the two most recent workshops and enabling the comparison of incremental values between clean and separated-flow cases. Furthermore, to reduce the extent of the trailing-edge separation observed on the baseline DLR-F6 configuration studied in DPW II, the Reynolds number for these cases was raised from 3×10^6 to 5×10^6 in DPW III. In addition to the two

Presented as Paper 930 at the 46th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 7–10 January 2008; received 18 June 2008; revision received 29 January 2009; accepted for publication 31 January 2009. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the $10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0021-8669/09 and $10.00 in correspondence with the CCC.

*Professor, Department of Mechanical Engineering. Associate Fellow AIAA.
†Boeing Technical Fellow. Associate Fellow AIAA.
‡Research Engineer, Institute of Aerodynamics and Flow Technology.
§Assistant Head, Configuration Aerodynamics Branch. Associate Fellow AIAA.
¶Assistant Head, Computational Aerosciences Branch. Senior Member AIAA.
**Senior Specialist Engineer. Senior Member AIAA.
††Researcher, Civil Transport Team, Aviation Program Group. Member AIAA.

DLR-F6 wing-body configurations, a set of two wing-alone configurations were included as a second test case for DPW III. The purpose of this second test case was to provide a simpler wing-alone geometry, which would enable the use of higher grid resolution, thus enabling a more complete grid resolution study, and to provide a well-behaved test case with no anticipated regions of separated flow. Although one of the objectives of DPW III was to test the hypothesis that the large scatter in the results of the previous workshops was due to the presence of separated flow, considerable scatter was also observed for the DPW III results, even for the cases with negligible flow separation such as the wing-body-fairing cases and the simple wing cases. Furthermore, in many instances, this scatter was not observed to decrease with increasing grid resolution.

Over the course of these three workshops [6–8] and numerous follow-up studies undertaken over the years [9–27], a consensus has been emerging suggesting that the dominant sources of error in aerodynamic drag prediction are related to spatial discretization error, which in turn stems largely from grid resolution and grid-quality issues. The gradual evolution of this consensus can be seen in the focus of the three DPW workshops. In DPW I, the workshop results were all presented for a standard grid size, which at the time was in the neighborhood of 3 million points [6]. The considerable scatter of the collective results of the workshop, along with subsequent independent grid resolution studies [10–13], prompted the organizers to mandate a grid resolution study as part of DPW II. Although the results of the second workshop were analyzed more rigorously with regard to spatial error or grid convergence, various anomalies were identified, such as inconsistent grid convergence behavior and collective scatter, which remained independent of the level of grid resolution. The fact that many of these cases contained substantial regions of separated flow further clouded the issue of which error source remained most important (modeling or discretization), an issue that the third workshop attempted to address through the addition of the fairing, the inclusion of simpler wing alone geometries, and the use of a more extensive grid refinement exercise with much finer grids than previously employed. The fact that the bulk of the DPW III results have been presented in the context of a grid convergence study is a testament to the importance attributed to determining the effect of discretization error on solution accuracy.

The identification of spatial discretization effects as a dominant error source has important implications for advancing the state of the art of aerodynamic simulations, because improvements in other areas such as turbulence and transition modeling will either be inconsequential or impossible to assess in the presence of dominant spatial discretization error. Thus, understanding and quantifying grid quality and resolution effects on solution accuracy remains one of the most important requirements for improving aerodynamic predictive ability. The collective DPW results have established some surprising facts concerning the sensitivity of aerodynamic drag prediction with respect to computational meshes. Over the course of the workshop series, the level of mesh resolution employed for calculations has steadily increased, from grids of several million points in DPW I [6] up to grids of over 40 million points in DPW III [8], with some individual studies employing grids up to 100 million points [20,22,25]. This steady increase in available grid resolution has been enabled largely due to advances in computational hardware and improved solution algorithms. However, the collective workshop results have not demonstrated substantially reduced variability over the course of the workshop series, in spite of over an order-of-magnitude increase in grid resolution from the first to the third workshop. To be fair, much of the observed variability in the collective results stems from the large numbers of different solvers, turbulence models, grid types, and particular grid constructions used throughout the workshop series by the participants. Individual participant workshop results have demonstrated monotone grid convergence behavior, particularly for attached-flow cases. However, in many instances, the demonstration of apparent grid convergence is contradicted by results obtained with the same code on grids of different construction. These results raise important questions with respect to the effect of the computational mesh on aerodynamic predictive ability, such as how can one minimize grid induced simulation errors, what constitutes a good mesh for aerodynamic simulations, and how should one go about assessing grid convergence?

We attempt to address some of these questions in this paper for transonic cruise drag prediction. In the following section, we begin with a description of the DPW geometries and test cases, followed by an overview of the gridding guidelines established and used for the provided workshop grids. In the subsequent section, we discuss some of the important issues related to mesh generation for aerodynamic simulations, such as the establishment of adequate grid resolution, the ability to assess grid convergence, and the determination of grid-quality effects on solution accuracy. Finally, we suggest a path forward for the community for addressing some of these open questions.

II. Geometry and Test Case Description

The baseline wing-body configuration used for both DPW II and DPW III, referred to as the DLR-F6, corresponds to a transport aircraft wind-tunnel model designed to cruise at transonic speeds. The basic planform characteristics of this model are depicted in Fig. 1, including the position of the wing-mounted nacelles, which were part of the DPW II study but omitted for DPW III. The wing has a leading-edge sweep of 27.1 deg, a quarter-chord sweep of 25 deg, and a dihedral angle of 4.787 deg. The DLR-F6 model has an acute angle of roughly 60 deg between the fuselage and wing upper surface at the trailing edge, which leads to flow separation at the wing-body juncture. A wing-body fairing was developed by Vassberg et al. [5], with the goal of completely eliminating the separation bubble in this region at the design conditions. The DLR-F6 wing-body with additional fairing used in the DPW III study is referred to as the DLR-F6 with FX2B, or simply FX2B, configuration. Figure 2 shows the computed flow on the DLR-F6 and FX2B configurations at the design condition [Mach $= 0.75$, $C_L = 0.5$, and $Re = 5 \times 10^6$ (based on the mean aerodynamic chord)] and illustrates the separated flow for the original configuration as well as the fully attached flow for the new FX2B configuration. DPW III was designed as a blind test because no experimental data on the FX2B configuration were available at the time of the workshop, although wind-tunnel data have recently been obtained on both configurations in the NASA Langley Research Center (LaRC) National Transonic Facility (NTF) and will be used for comparison with numerical predictions in subsequent studies. For DPW III, participants were required to provide computed results for both the baseline DLR-F6 geometry and the DLR-F6 with FX2B fairing configuration at the following flow conditions:

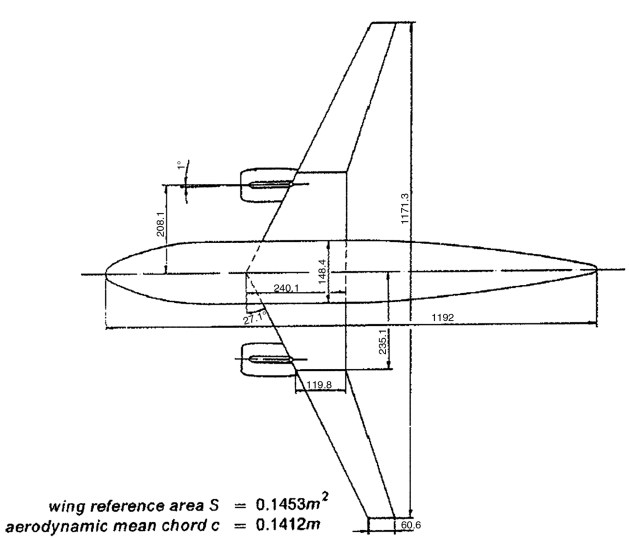

Fig. 1 DLR-F6 wing-body planform definition including standard nacelle configuration used in DPW II.

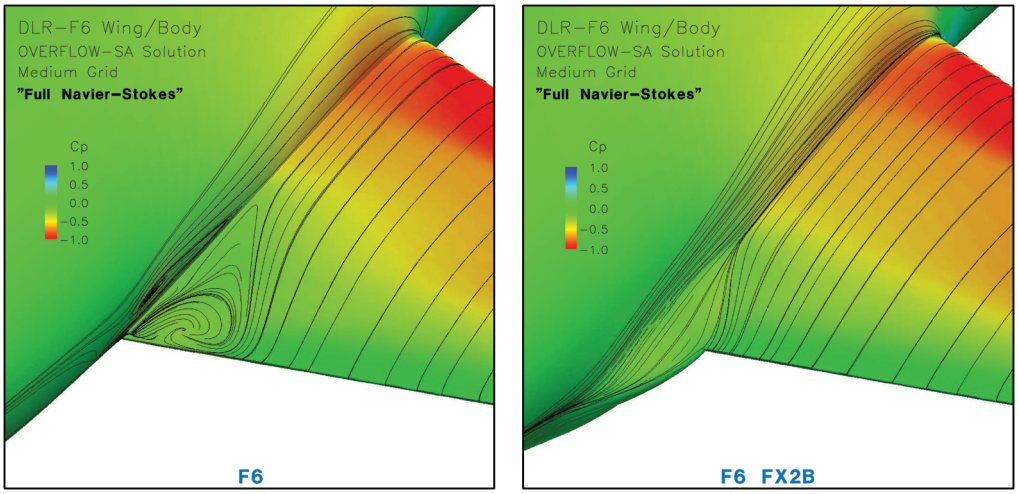

a) b)

Fig. 2 Illustration of a) separated flow in wing-root junction region for DLR-F6 configuration and b) attached flow for DLR-F6 with FX2B fairing configuration.

1) For a fixed-C_L single-point grid sensitivity study on three grids, Mach $= 0.75$, $C_L = 0.5$, and $Re = 5 \times 10^6$.

2) For drag polar on a medium grid, Mach $= 0.75$, $Re = 5 \times 10^6$, and incidences of -3.0, -2.0, -1.0, -0.5, 0.0, 0.5, 1.0, and 2.5 deg.

As mentioned previously, DPW III also included a simpler set of wing-alone geometries for enabling more extensive grid refinement studies in the presence of simpler geometries with fully attached flow. Two closely related wing configurations were designed by DPW committee members, labeled wing 1 and wing 2 (or W1 and W2). Both have a common simple trapezoidal planform shape with no leading-edge or trailing-edge breaks and incorporate a leading-edge sweep of 17.2 deg and quarter-chord sweep of 15.0 deg. The airfoil selected for wing 1 is representative of a supercritical section found on most modern transonic aircraft, producing a nearly flat rooftop pressure distribution with a shock located at approximately 57% chord at the design conditions of Mach $= 0.76$ and $C_L = 0.50$. The second configuration, wing 2, was created by optimizing the wing 1 configuration for drag reduction at the design point using only camber and twist changes. Additional information regarding the design and definition of these wings is provided in [4] in the workshop presentations on geometry review. For DPW III, participants were required to compute the following cases:

1) For a fixed-incidence single-point grid sensitivity study on four grids, Mach $= 0.76$, incidence is 0.5 deg, and $Re = 5 \times 10^6$.

2) For drag polar on a medium grid, Mach $= 0.76$, $Re = 5 \times 10^6$, and incidences of -1.0, 0.0, 0.5, 1.0, 1.5, 2.0, and 2.5, and 3.0 deg.

These wing-alone configurations were devised purely for computational studies, and no experimental data exist for these cases.

Although the workshop results called for both fixed-point grid refinement studies and the calculation of drag polars on a fixed grid, we will concentrate on the results of the fixed-point grid refinement studies for the purposes of this paper, with the intent of focusing on grid-related errors for drag prediction. Note that the grid refinement studies are specified at fixed C_L for the DLR-F6 and FX2B configurations, and these are specified at a fixed angle of attack for the wing-alone cases.

III. DPW Gridding Guidelines

The wide range of scales present in aerodynamic problems requires that even the simplest simulations employ highly variable mesh resolution, tailored to the problem at hand. In the absence of fully automated adaptive meshing procedures, this means that mesh generation for computational aerodynamics will remain somewhat of an art based on practitioner experience and intuition, at least for the near term. In an attempt to ensure some level of uniformity over the various grids used for the DPW calculations, the committee developed a set of gridding guidelines that codify much of the collective experience of the committee members in aerodynamic grid generation practices. These guidelines have remained relatively unchanged over the course of the three workshops, and thus only those published for DPW III are discussed herein. For the latter workshops, the additional focus on grid convergence issues also resulted in the issuance of guidelines for generating consistent families of uniformly refined meshes for grid convergence studies. The DPW III gridding guidelines are summarized next, where C, M, F, and X refer to coarse, medium, fine, and extra-fine mesh levels, respectively.

1) For the boundary-layer region, $Y^+ \leq [1, \frac{2}{3}, \frac{4}{9}, \frac{8}{27}]$ $[C, M, F, X]$, $\Delta_1 \sim 0.0006$ mm (approximate dimensional spacing for $Y^+ = 1$), $\Delta_2 = \Delta_1$ (two cell layers of constant spacing at viscous wall), and growth rates are 1.25 (preferably 1.20).

2) The far field is $\sim 100\ C_{\text{ref}}$ lengths away from the geometry.

3) The local spacings on the medium grid are as follows: chordwise is 0.1% local chord at the wing leading edge and trailing edge, spanwise is 0.1% semispan at root and tip, and cell size on the fuselage nose and tail is 2% C_{ref}.

4) Cells across the wing trailing-edge base are [8, 12, 16, 24] $[C, M, F, X]$

5) The grid family is as follows:

a) Medium mesh is representative of the current engineering drag predictions.

b) A parametric family of uniformly refined grids in sequence is maintained.

c) The grid size is to grow $\sim 3\times$ for each level refinement (structured is 1.5 times in each I, J, and K direction).

d) Consideration is given to multigriddable dimensions on structured meshes.

e) The sample size for the DLR-F6 wing/body is [2.7, 8, 24, 72 million] $[C, M, F, X]$.

The guidelines specify an outer boundary location of approximately 100 reference chords out from the body, which has been found to minimize the effect of the outer boundary condition treatment on computed lift and drag values. Local grid spacings of 0.1% chord at critical regions such as leading edge, trailing edge, and wing tip and root areas are also specified for the medium-level grid, with the appropriate scaling of these spacings for the coarser and finer grids. For blunt-wing trailing-edge geometries, between 8 and 24 cells across the trailing-edge thickness are specified, as experience has shown the high level of sensitivity of the results to resolution along blunt trailing edges. The standard practice of resolving the boundary layer down to a minimum of $y^+ = 1$ is also specified for the coarsest grid, with appropriately smaller spacings for the finer grids. Additionally, the first two layers of cells at the wall should be of similar spacings, and the growth rate of the normal spacing through the boundary layer should not exceed 1.25. For grid convergence

Table 1 DPW III case I (DLR-F6/FX2B) submitted grids: number of grid points

		DLR-F6			FX2B		
Family	Type	Fine	Medium	Coarse	Fine	Medium	Coarse
Tinoco	MB	27,185,664	8,080,896	2,298,880	27,185,664	8,080,896	2,298,880
SAUNA	MB	9,761,201	4,731,073	2,551,989	9,761,201	4,731,073	2,551,989
Extruded	MB	28,367,120	9,343,009	2,996,626	28,367,120	9,329,185	3,028,420
Gridgen	MB	27,982,776	8,927,196	2,739,621	27,982,776	9,138,772	2,842,878
Sclafani	OS	26,892,352	7,985,236	2,387,918	26,969,192	8,020,348	2,395,170
DLR	HY	8,535,263	5,102,446	2,464,385	10,305,876	6,111,664	2,873,102
ANSYS	HY	18,120,772	8,038,922	3,059,189	20,472,520	8,272,308	3,163,605
LaRC	UN	40,014,934	14,298,135	5,354,214	41,069,036	14,598,610	5,618,073
AFLR	UN	11,374,451	3,792,485	1,492,082	11,849,212	3,178,559	1,640,590
Embraer	UN	24,030,000	8,320,000	3,550,000	24,030,000	8,320,000	3,550,000
STAR	UN	——	12,377,058	——	21,509,137	12,469,599	8,421,799
TAS	UN	17,535,215	9,431,154	5,399,929	17,219,535	9,481,477	5,422,128

Table 2 DPW III case 2 (wing-1/wing-2) submitted grids: number of grid points

	DPW-W1				DPW-W2			
Family	Extra fine	Fine	Medium	Coarse	Extra fine	Fine	Medium	Coarse
Tinoco	14,811,489	8,620,123	4,204,203	1,602,651	14,811,489	8,620,123	4,204,203	1,602,651
Sclafani	55,014,321	16,265,909	4,856,149	1,442,285	55,014,321	16,265,909	4,856,149	1,442,285
DLR	17,053,510	10,150,588	5,288,507	2,174,364	16,631,805	9,910,645	5,030,379	1,928,405
LaRC	36,956,019	11,492,625	4,495,117	1,818,508	38,462,630	11,903,329	4,658,853	1,882,672
Raytheon	12,748,678	6,138,245	2,417,082	983,633	12,419,567	5,963,713	2,325,884	947,409

studies, the specifications call for sequences of grids with uniform refinement in all areas of the domain, resulting in a growth of the number of grid points by a factor of 3 for each consecutively finer mesh, with approximate grid sizes of 2.7, 8, 24, and 72 million points for the coarse, medium, fine, and extra-fine meshes, respectively.

Tables 1 and 2 list the attributes of the various grids submitted and employed for computations for DPW III, and Figs. 3–7 illustrate the topology and resolution for several of the different baseline grids used for the workshop. These tables and figures illustrate the wide variety of grid types used for the workshop, which include overset structured meshes (OS), point-matched multiblock structured grids (MB), unstructured tetrahedral meshes (UN), and hybrid prismatic-tetrahedral meshes (HY). Only the general characteristics of these grids are shown here, and further details of the individual grid submissions can be found in companion papers [22–25,27].

Although all submitted grids were generated according to the DPW published guidelines, specific compliance with the guidelines of individual submissions could not be monitored, and it remains unclear how closely all these submissions adhere to the guidelines. Furthermore, specific characteristics of various grid generation packages may result in substantially different grid characteristics based on the same guidelines. For example, for overset structured meshes, the topology of the grid system and location and extent of the overlap regions must be considered carefully for accurate numerical drag prediction, as discussed by Vassberg et al. [28]. For unstructured meshes, much higher node counts will result for approaches that do not allow or employ spanwise stretching, especially considering the high level of trailing-edge resolution specified by the guidelines. In some cases, the trailing-edge resolution must be relaxed to maintain practical grid sizes in the absence of spanwise stretching [23]. Additionally, for unstructured tetrahedral mesh solvers, the specification of mesh resolution for cell-centered discretizations generally does not correspond to the specification for vertex-based discretizations, because the latter approach contains 5 to 6 times fewer degrees of freedom on the same mesh as the cell-centered approach. However, because vertex-based discretizations entail more complex stencils, establishing the equivalent grid resolution for similar accuracy for both approaches has remained an open question. Experience from previous DPW exercises has suggested that matching the number of triangles on the boundary surface for cell-centered approaches to the number of vertices on the surface for vertex-based discretizations (i.e., a 2:1 resolution ratio on the surface), results in vertex-based meshes containing 3 times more points than the equivalent cell-centered mesh and produces equivalent accuracy for both schemes. However, these findings are far from conclusive, and the situation becomes even more complex for hybrid meshes containing mixtures of tetrahedra, prisms, pyramids, and hexahedra.

In summary, the wide variety of grid types and technologies used for the workshop series implies that the characterization of these grids by simple node counts and guidelines is approximate at best, and substantially different characteristics can be expected between supposedly equivalent grids. On the other hand, one of the outcomes of the workshop series is a large database of publicly available grids and published computational results, which can be used to assess grid quality and perform further numerical experiments.

IV. Overview of DPW Grid-Related Results

Figure 8 depicts the collective results from DPW I, and Figs. 9 and 10 show the collective results from DPW II and DPW III, respectively. Although the DPW I results were only examined at a fixed mesh resolution, the DPW II and DPW III results were examined as a function of grid resolution, using a grid index factor in DPW III defined as the number of grid points raised to the $-2/3$ power as a measure of average cell size squared, which should produce straight-line plots for grid convergence studies based on second-order-accurate discretizations. Furthermore, in DPW III, the idealized drag coefficient [defined as $C_{D_{\text{ideal}}} = C_D - C_L^2/(\pi \text{AR})$, where AR refers to the wing planform aspect ratio] is used instead of the actual drag coefficient. By subtracting the estimated induced drag, the use of the idealized drag coefficient represents an attempt to minimize the effect of variations in lift of the submitted results. The shear amount of information present in these plots is a testament to the vast numbers of different solvers, turbulence models, grid types, and grid resolutions used by workshop participants. Because it is well known that different solvers can be more or less sensitive or tolerant to different mesh characteristics, the ideal situation for assessing grid suitability would be to run a single grid with all solvers or, conversely, all grids with a single solver. Unfortunately, the large number of different entries in the DPW data sets combined with the fact that these workshops represent goodwill efforts on the part of the

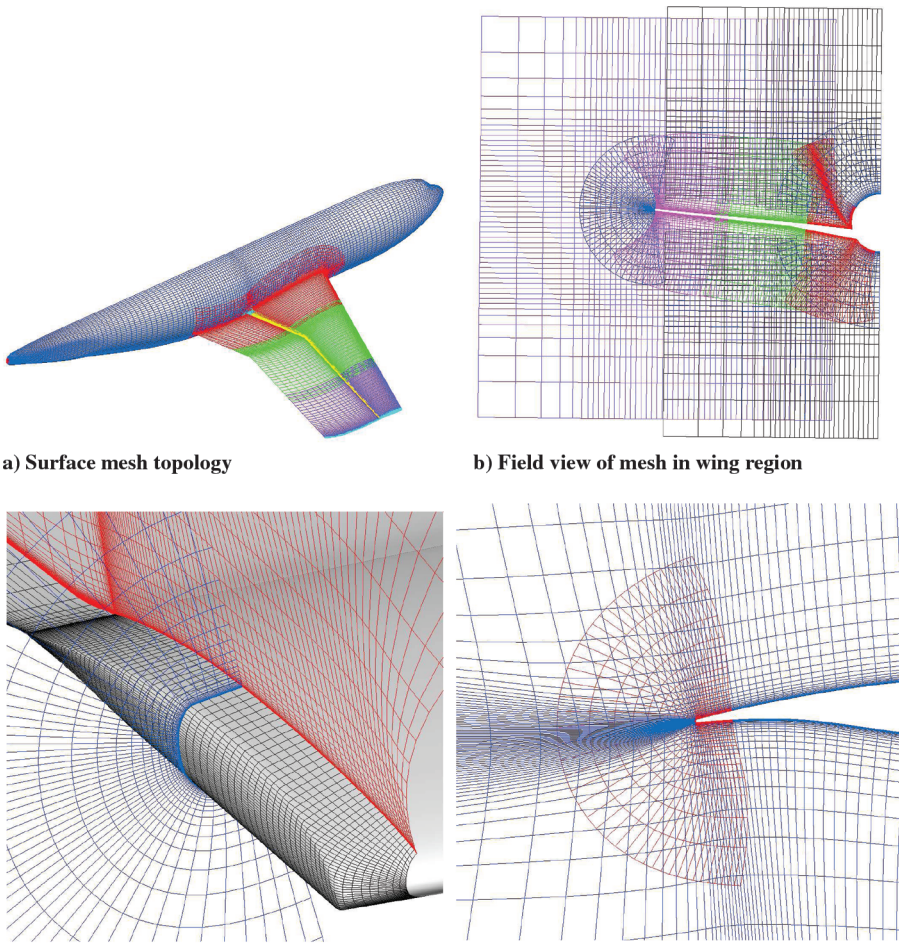

a) Surface mesh topology

b) Field view of mesh in wing region

c) Wing tip mesh topology and resolution

d) Trailing-edge mesh topology and resolution

Fig. 3 Illustration of overset mesh topology and resolution details for DLR-F6 configuration reproduced from Sclafani et al. [22].

community at large obviate the possibility of performing such types of controlled experiments. Nevertheless, various important observations can be made from the DPW workshop results regarding the influence of mesh generation technology on the accuracy of aerodynamic simulations. In the following, we separately discuss the issues of grid resolution, grid convergence, and grid quality.

A. Grid Resolution

It is generally evident from the workshop results that significantly increased grid resolution will be necessary for meeting the accuracy requirements of aerodynamic drag prediction. To some degree, a natural progression in what is considered to be practical grid resolution occurs as hardware and solver technologies advance. This is reflected in the grid sizes used in the three successive DPW workshops, which have increased by over an order of magnitude in this 5-year period, going from approximately 3 million points in 2001 (DPW I) to over 40 million points in 2006 (DPW III), with some interim and follow-up studies using grids in the vicinity of 100 million points [20,22,25]. However, given the levels of discrepancies noted in even the latest workshop results, accompanied by the grid convergence trends, it is evident that the ability to employ much finer grids, perhaps up to 10^9 grid points on a regular basis, would greatly improve the accuracy and confidence in production-level aerodynamic drag prediction.

These high grid resolution requirements stem not only from the high accuracy requirement of drag prediction, but also from the large range of scales that must be resolved in aerodynamic simulations. For example, in the vicinity of the outer boundary, the dominant scales are many times larger than the mean chord or characteristic body dimension, whereas flow details in the vicinity of blunt trailing edges, where the characteristic length is thousands of times smaller than the mean chord, can have a profound effect on lift and drag accuracy if not properly resolved. Additionally, the thin boundary layers produced at test and flight Reynolds numbers can only be resolved in a tractable way using highly stretched anisotropic meshes with aspect ratios in excess of 1000:1 or even 10,000:1. These characteristics dictate the use of highly variable resolution and anisotropic meshes for aerodynamic calculations and are the principal reason that mesh generation has remained an art largely based on practitioner experience and intuition.

The highly variable and tailored resolution required for accurate drag prediction is illustrated in Figs. 3–7 for representative sets of overset structured, point-matched block-structured, and unstructured meshes used in the workshop. For the structured-mesh sets, the figures demonstrate how the blocking or overset topology can have profound effects on the resulting local resolution. For the unstructured meshes, issues such as element type and spanwise stretching can play a significant role in the resulting resolution as well. In all cases, the figures illustrate highly variable resolution resulting from attempts to provide enhanced refinement in critical areas, such as at blunt trailing edges, and at the wing-fuselage junction, where a pocket of separated flow is anticipated for the DLR-F6 configuration.

The workshop results have demonstrated the complex nature of the interaction between grid resolution and drag prediction accuracy. In most instances, increased grid resolution was found to have only a minor effect on local quantities such as surface pressure coefficients or skin friction. This behavior is illustrated in Fig. 11, for the W1 test case, using four successively finer unstructured meshes from 1.8 to 37 million points, using the NSU3D unstructured-mesh solver [27]. In this case, the only notable differences in the surface values in going from coarse to fine meshes are in the resolution of the shock wave, which becomes sharper while rotating about its center point.

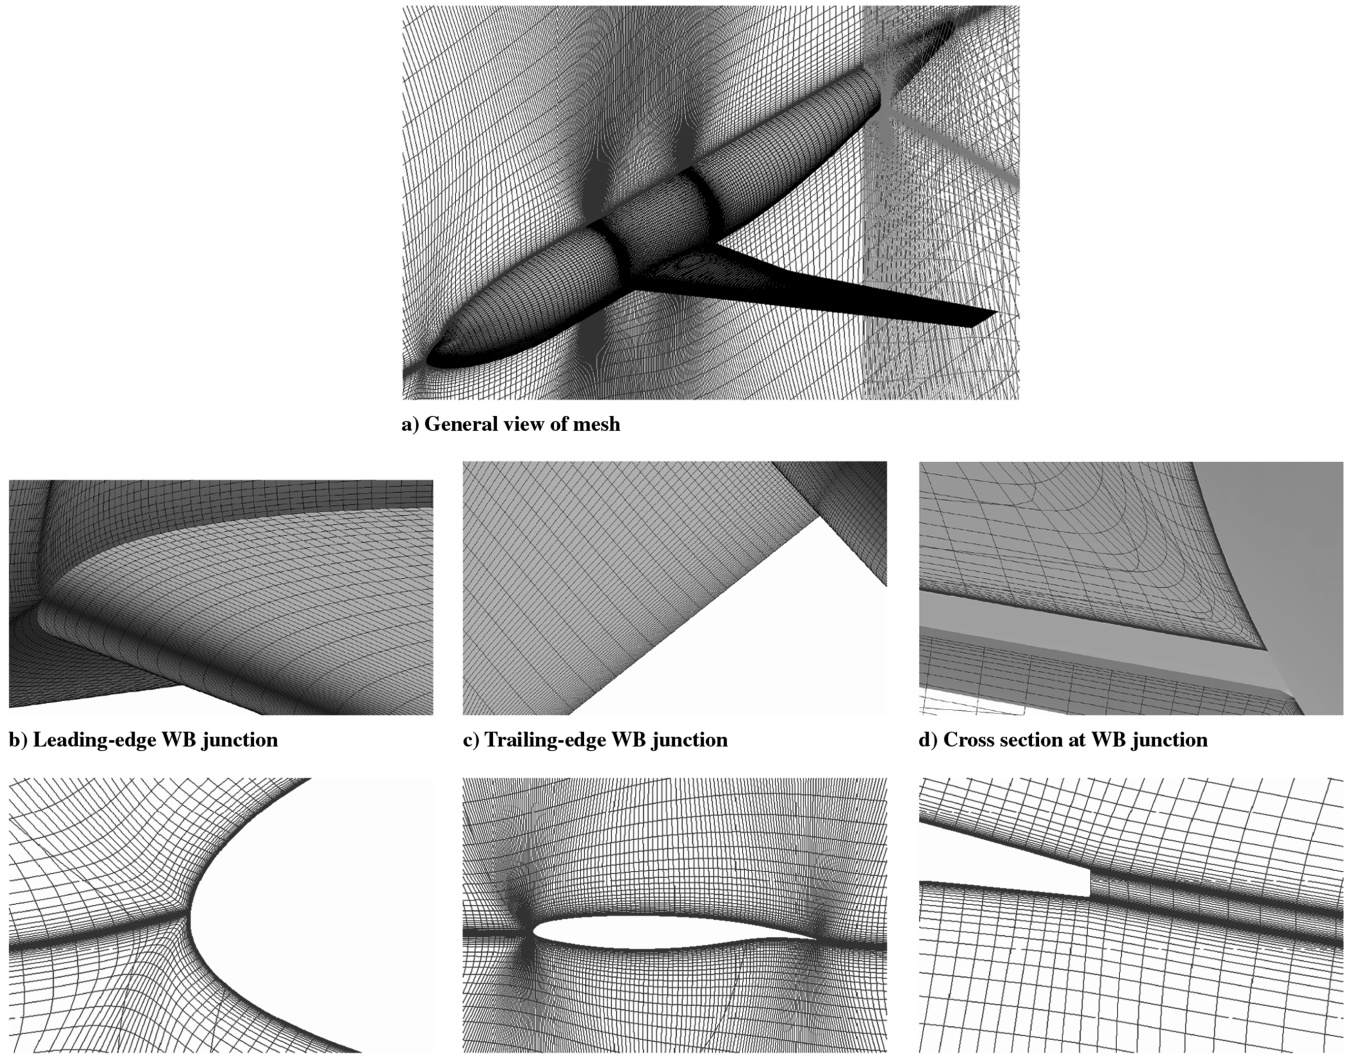

a) General view of mesh

b) Leading-edge WB junction

c) Trailing-edge WB junction

d) Cross section at WB junction

e) Wing leading edge

f) Airfoil section cut

g) Wing trailing edge

Fig. 4 Illustration of block-structured-mesh topology and resolution of Boeing generated grids from Tinoco et al. [25] for DLR-F6 configuration (WB denotes wing-body).

Although in this case the shock position remains relatively unchanged, cases in which the shock position changes slightly with grid resolution result in shocks of different strength, thus altering the computed drag values appreciably. In such cases, changes in the shock position may be due to inadequate resolution ahead of the shock, which thickens the boundary layer through the convection of upstream-generated numerical entropy. Similarly, in addition to their direct effect on friction drag, inaccuracies in computed skin friction can alter the prediction of downstream separation locations, which results in a much more profound effect on overall drag prediction. An example of this behavior was originally given by Anderson et al. [29] and reproduced in [30], in which the effects of inadequate boundary-layer resolution were found to produce premature separation on a two-dimensional high-lift configuration.

Examination of individual workshop results has shown that grid resolution requirements for drag prediction involving regions of separated flow are much more demanding and less well understood than for cases in which the flow is fully attached. For example, various participants have reported adequate grid convergence of drag values as well as relative insensitivity of these values to grid topology and grid type for the FX2B and wing-alone configurations in DPW III. On the other hand, for the DLR-F6 wing-body case involving pockets of separated flow, many participants have observed continuing sensitivity of the computed drag values with increasing grid resolution, as well as with differing grid topologies and grid types. Furthermore, no consensus on the sensitivity of drag prediction to local variations in grid resolution can be determined from the workshop results. For example, Sclafani et al. [22] observed that the size of the separation bubble at the wing-root junction of the DLR-F6 (and thus the drag value) increases continuously with increasing grid resolution for an overset mesh approach, and Tinoco et al. [25] and Murayama and Yamamoto [23] observed a relatively constant separation-bubble size with increasing grid resolution for a block-structured-mesh approach. Using an unstructured-mesh approach, Murayama and Yamamoto [23] also observed little variation in the amount of flow separation with increasing grid resolution, whereas Brodersen et al. [24] reported increased separation, and Tinoco et al. [25] observed large increases in flow separation using two different unstructured-mesh solvers. Similar inconsistencies are observed with respect to trailing-edge separation, which was found by Mavriplis [27] to increase as a function of grid resolution using an unstructured-mesh approach, but found by Murayama and Yamamoto [23] to decrease or remain relatively constant for both their unstructured and block-structured approaches, whereas Sclafani et al. [22] found no trailing-edge separation using OVERFLOW.

In the absence of a clear understanding of the necessary grid resolution requirements, experimentation with progressively finer meshes remains the most viable approach. Grid refinement studies are predicated on the assumption of asymptotic behavior of the

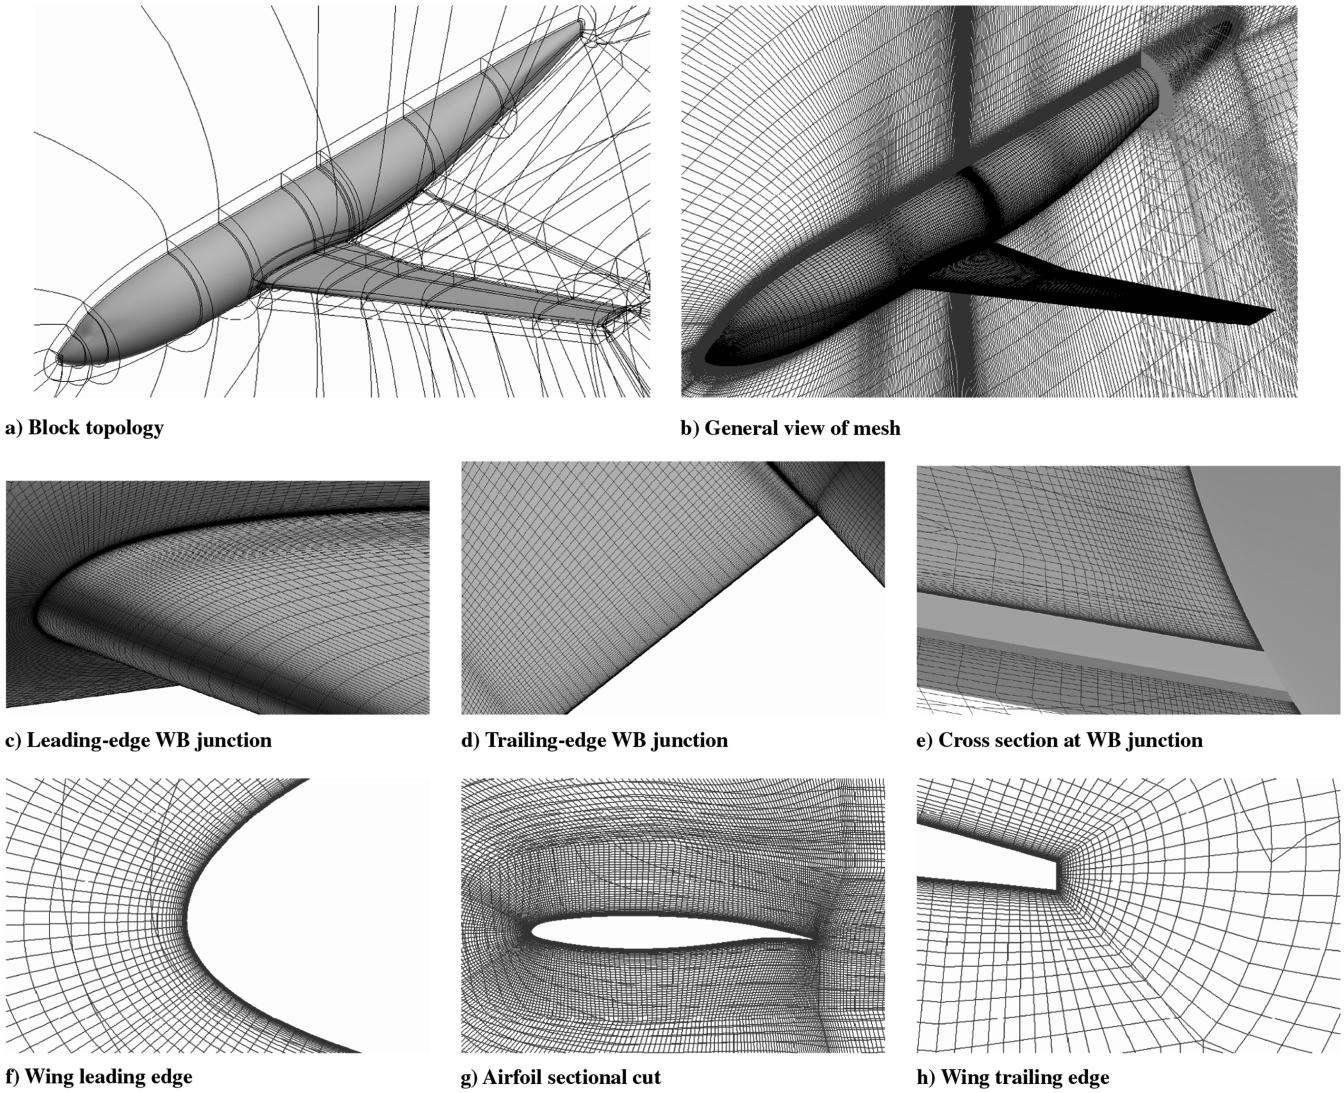

Fig. 5 Illustration of block-structured-mesh topology and resolution of JAXA generated grids from Murayama and Yamamoto [23] for DLR-F6 configuration.

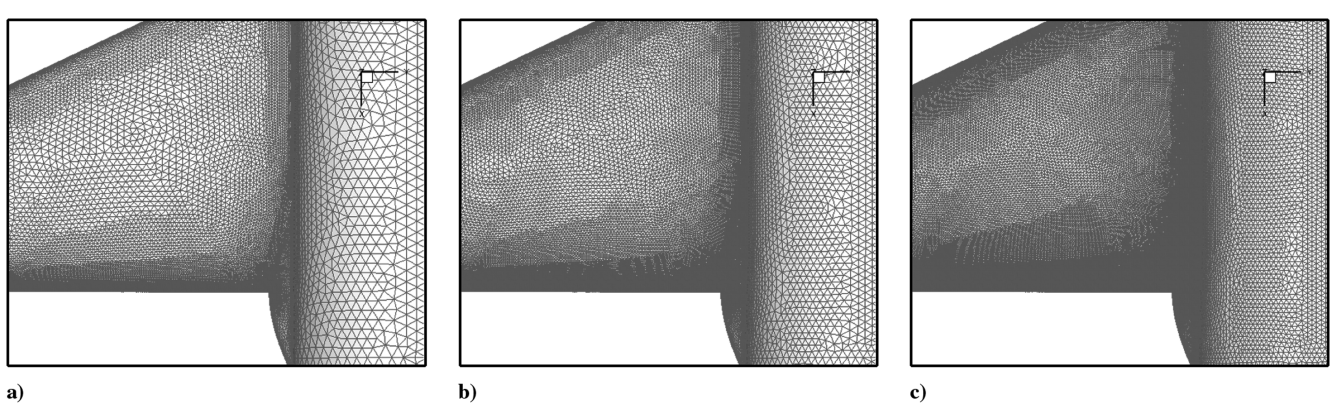

Fig. 6 Wing-body-fairing unstructured grid resolution near wing-fuselage juncture for DPW III meshes generated by LaRC and designated LaRC node-centered grids: a) coarse mesh (5.6 million points total), b) Medium mesh (14.6 million points total), c) fine mesh (41 million points total).

solution error, given some adequate level of grid resolution throughout the computational domain. Surveying the DPW III grid convergence study results, one may conclude that the coarse baseline meshes used in the workshop, containing on the order of 2 to 2.5 million points, often do not lie in the asymptotic region for absolute drag prediction, even for fully attached-flow cases (see, for example, [8,22]). Thus, the minimum adequate grid resolution for asymptotic behavior appears to be in the range of 5 to 10 million points, although this can also be expected to depend strongly on the optimality of the distribution of grid resolution throughout the domain.

One of the problems with increasing grid resolution is the rapid growth in grid sizes with successive refinements. In three dimensions, a uniform grid refinement by a factor of 2 in each direction results in an overall increase in grid size by a factor of 8, with, at best, an equivalent (but often larger) increase in required flow solution resources, for an optimally convergent and scalable flow solver.

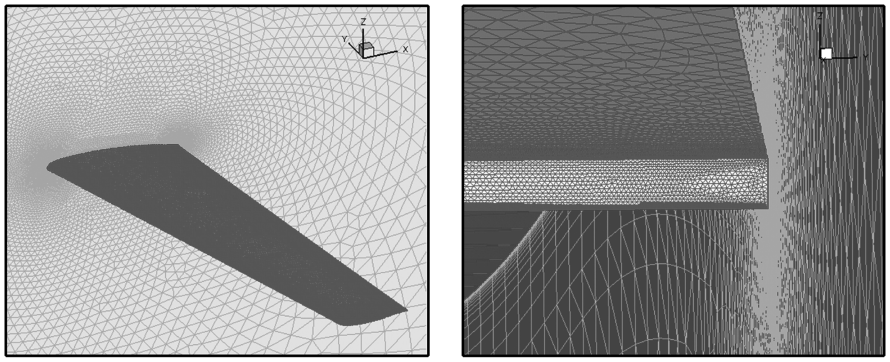

a) b)

Fig. 7 DPW III wing-1 fine unstructured mesh generated by LaRC using VGRID software. a): Fine Grid on wing-1 Configuration comprising a total of 15 million points. b): Detail of trailing-edge base resolution for fine grid on wing-1 Configuration containing approximately 16 cells across blunt trailing-edge base.

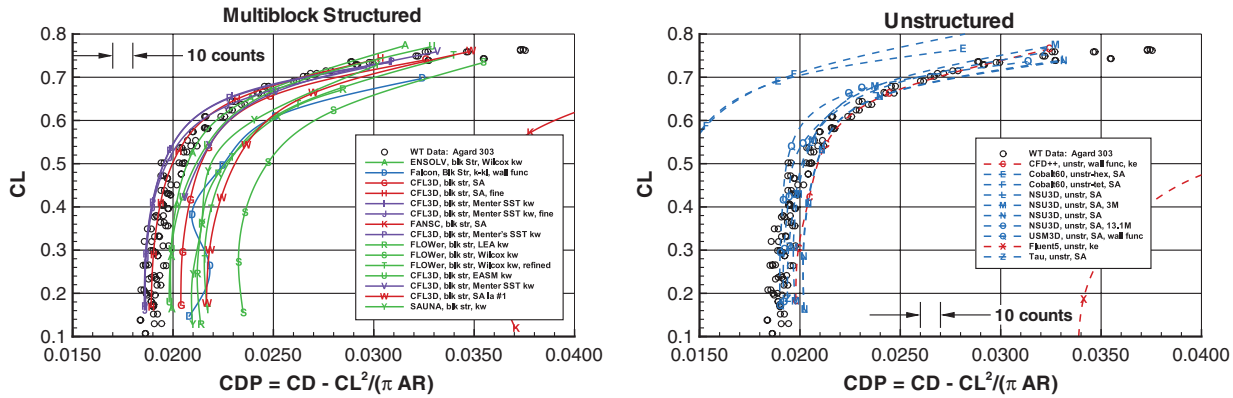

Fig. 8 Collective results from DPW I plotted as computed idealized drag coefficient versus lift coefficient for various solvers and grid types.

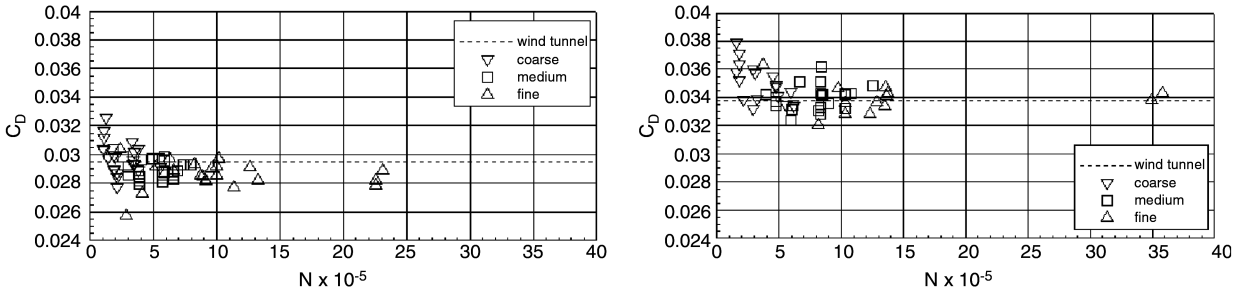

a) b)

Fig. 9 Collective results from DPW II plotted as computed drag coefficient at fixed-lift condition versus number of grid points a) wing-body configuration and b) wing-body-nacelle-pylon configuration.

Thus, a single uniform refinement of one of the fine DPW III baseline meshes of 40 million points would result in a mesh of 320 million points, which is already beyond the capabilities of most practitioners. Furthermore, few mesh generation packages are capable of generating such large-scale meshes in a robust and efficient manner.

To evaluate the suitability of a particular mesh for a given aerodynamic simulation, an a posteriori error-estimation technique is required that, when coupled with an adaptive meshing strategy, can be used to improve the mesh resolution and accuracy of the simulation. Figure 12 provides an illustration of an adapted unstructured mesh used for improving the accuracy of drag prediction in DPW III employed by Brodersen et al. [24]. In this case, the mesh was refined locally based on the gradient of total pressure in the computed solution. The principal difficulty in effective use of adaptive mesh refinement for drag prediction problems lies in the reliability of these types of gradient-based error indicators, which themselves require the solution to lie in the asymptotic error region. More sophisticated approaches such as adjoint-based functional error-estimation techniques have shown promise [31–33], but these also rely on a linearization about the current solution and will require further development to achieve the level of reliability sought for production aerodynamic calculations.

B. Grid Convergence

As mentioned previously, the DPW workshop series has gradually devoted increased attention to the issue of grid convergence. The principle concept of grid convergence states that as the mesh resolution is increased in all areas of the domain, the discretization error in the solution should vanish and the discretized solution should gradually tend toward the continuous solution. For a second-order-accurate scheme, where the discretization error varies as $\mathcal{O}(h^2)$ and

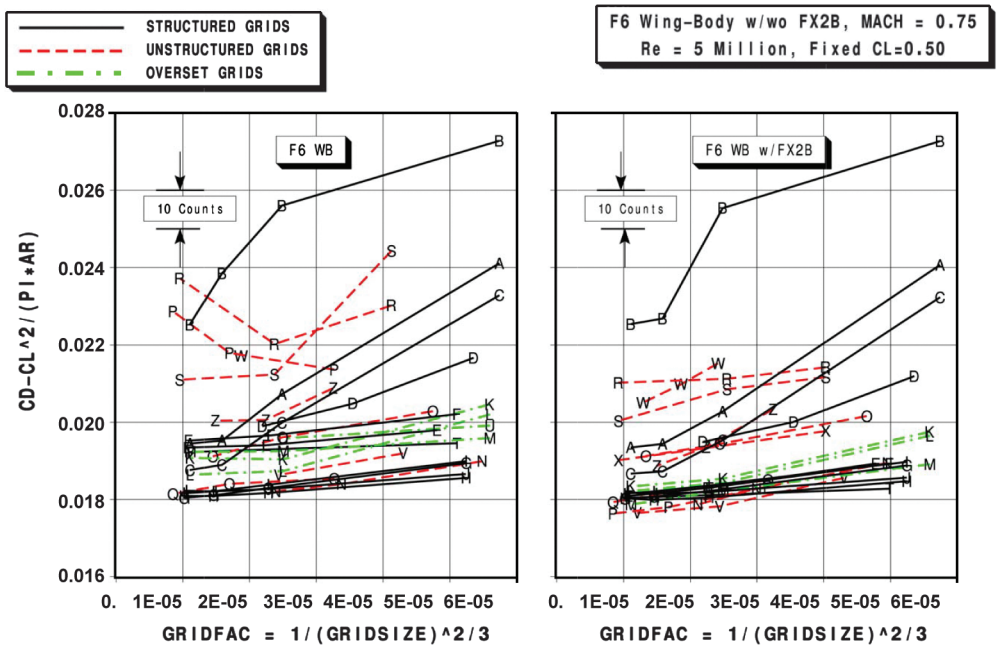

Fig. 10 Collective results from DPW III plotted as computed idealized drag coefficient at fixed-lift condition versus grid resolution index for DLR-F6 wing-body and wing-body with FX2B fairing.

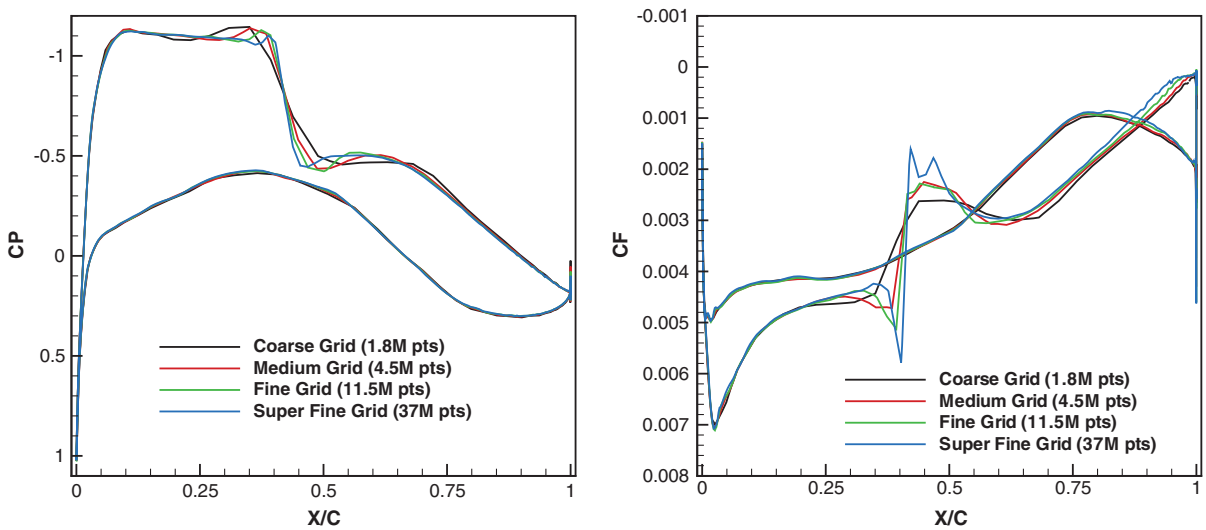

a) b)

Fig. 11 Plots of a) computed surface pressure coefficients and b) skin-friction coefficients at the 55.1% span location at Mach = 0.76, $\alpha = 0.5$ deg, and $Re = 5 \times 10^6$ on the four grids used in the grid refinement study for wing-1 with the NSU3D unstructured-mesh solver (reproduced from Mavriplis [27]).

where h represents some measure of the mesh spacing, solution values such as lift and drag plotted versus h^2 should produce a straight-line plot that extrapolates to the continuous solution values in the limit as $h \to 0$, provided that the grids are of the same family. The DPW III grid convergence studies consist of plotting solution values versus a grid index factor, defined as $N^{-\frac{2}{3}}$, where N represents the number of grid points. The grid index factor is meant to be representative of an average mesh spacing squared (i.e., h^2). However, because the grids all contain highly variable mesh spacings, two grids with vastly different local grid spacings can have similar grid index factors, and this measure only makes sense when applied to families of self-similar grids [34]. A family of self-similar grids consists of a set of coarse and fine meshes in which the resolution of the finer meshes is obtained by uniformly reducing the mesh spacing of the coarser meshes in all regions of the domain. For overset or block-structured meshes, this can be achieved by specifying increased numbers of points in the I, J, and K coordinate directions. For unstructured meshes, a global mesh spacing factor is required, which is used to rescale the specified spacings in all areas of the domain in a uniform manner. Although many unstructured-mesh generation packages such as VGRID [35] have been designed with the use of a global grid spacing factor in mind, experience has shown that the resulting meshes do not always conform to the desired uniform refinement characteristics throughout all regions of the domain, due to specifics of the mesh generation algorithm. Additionally, for both structured and unstructured meshes, the algorithmic parameters that determine the exponential growth of the normal boundary-layer resolution must be considered carefully to as closely as possible approximate the goal of uniform refinement of the anisotropic mesh in this region. The generation of consistent families of refined meshes for grid convergence studies in the presence of complex geometries is therefore a nontrivial task. The degree to which the uniform refinement objectives have been met for the families of grids used in the DPW workshops is not entirely clear,

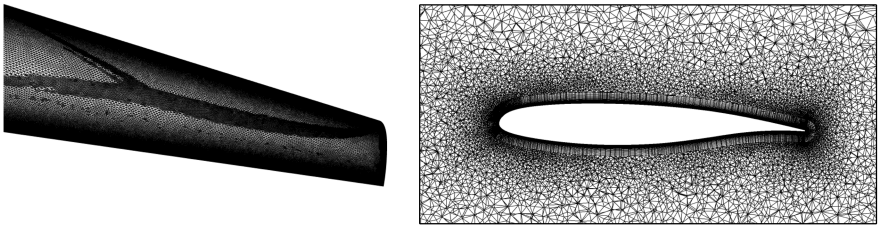

a) b)

Fig. 12 Illustration of adaptively refined mesh for DPW III wing-1 drag calculation using gradient of total pressure as a refinement criterion, and limiting the total increase in number of grid points to 70%: a) surface grid and b) field slice at $\eta = 0.17$ (reproduced from Brodersen et al. [24]).

although it can be presumed that the efforts of the workshops represent the state of the art in current grid generation technology.

Another issue that arises in the context of a grid convergence study is the ability to generate and employ sufficiently fine meshes. For grid convergence plots, a minimum of three data points (and preferably four or more points) is required to establish the degree of grid convergence of the solver. Because each successive uniform mesh refinement by a factor of 2 results in 8-times-higher grid resolution, using a four-grid sequence requires a factor-of-512 increase in the number of grid points between the coarsest and finest mesh of the sequence. The situation quickly becomes intractable, particularly given the previously stated observation that the coarsest meshes of the sequence must often be larger than several million points to fall in the asymptotic region. For this reason, the sequence of meshes in the DPW grid convergence studies were constructed using a lower refinement ratio, resulting in a threefold growth in the number of grid points between successively finer meshes (mesh spacing reduction of 1.5:1 versus 2:1). However, it is evident from this discussion that the ability to generate and employ much finer meshes remains crucial toward improving our ability for establishing grid convergence.

The specific nature of the test case must also be considered carefully in the context of a grid convergence study. The majority of the DPW test cases were computed using a fixed-C_L condition, rather than a fixed-angle-of-attack condition, because this is the industry practice for drag prediction. For example, the collective DPW III results shown in Fig. 10 depict computed (idealized) drag at a fixed $C_L = 0.5$ as a function of the grid index value. However, for a truly grid-convergent solution, all values including lift and drag should converge to their continuous values, as the mesh is refined for fixed geometric conditions (i.e., angle of attack). For these reasons, the simpler W1/W2 studies in DPW III were performed at fixed angle of attack or incidence, rather than at fixed-C_L conditions.

The DPW results have shown that grid convergence can be achieved by a number of solvers when carefully constructed families of uniformly refined grids are used, particularly for attached-flow cases. Figure 13a provides an illustration of the grid convergence achieved for the overset structured-grid code OVERFLOW on the DLR-F6 and FX2B test cases of DPW III, for the computed drag at constant lift on a sequence of grids ranging from 2.4 to 82.6 million points, reproduced from Sclafani et al. [22]. The results for the FX2B case clearly display second-order accuracy, as evidenced by the straight-line nature of the plot, and can be extrapolated reliably to an infinite resolution value by following this line to its intercept on the y axis. On the other hand, the DLR-F6 wing-body results do not produce a straight-line plot, as the drag values initially decrease and then subsequently increase with increasing grid resolution, due to the growing presence of separated-flow regions with increased grid resolution, as discussed previously. For the DPW III wing-alone test cases, reasonable grid convergence is demonstrated in Figs. 13b and 13c using OVERFLOW [22], although the drag convergence at fixed-lift conditions in Fig. 13c is seen to be superior to that at fixed-incidence conditions in Fig. 13b.

Various other workshop participant results have shown similar trends for all grid types, including overset, block-structured, and unstructured meshes. Figure 14 illustrates the grid convergence study for the FX2B configuration performed by Murayama and Yamamoto [23] using block-structured and unstructured meshes. In practice, it is strictly not valid to plot both types of grids on the same chart with a single grid index factor, because they represent two dissimilar families of grids. However, the results demonstrate that both codes achieve close to second-order accuracy on their respective family of grids. Perhaps more important, extrapolation of these plots also reveals that both codes converge to the same drag values in the limit of infinite grid resolution, particularly for the FX2B case, providing a strong verification of the consistency of these two sets of simulations on different grid types. The fact that the slope of the line for the unstructured-mesh results is steeper than the slope of the line for the structured-mesh results is indicative of higher errors associated with the unstructured-mesh results at a given level of resolution. In Figs. 14b and 14c, these drag values are further broken down into their component contributions (i.e., drag computed on the wing alone and the fuselage alone), and it is observed that the slope of the unstructured grid convergence plots is much lower for the wing alone values than for the body alone values, thus suggesting inadequate resolution of the fuselage as a principal source of error in the overall unstructured-mesh drag values, particularly for the WB case. This component drag decomposition strategy provides a rudimentary but effective approach for identifying regions of the computational domain requiring additional resolution.

The simple wing alone (W1-W2) configurations for DPW III were designed to enable an enhanced grid convergence study for cases consisting of mostly attached flow, using very fine meshes and using fixed-incidence conditions as opposed to the fixed-lift conditions employed in the previous cases.

Figure 15a displays the grid convergence achieved by the NSU3D unstructured-mesh solver [27] on this case for W1 and W2 at fixed incidence, using meshes ranging from 1.8 to 38 million points, demonstrating consistent second-order-accurate results, as evidenced by the nearly straight-line behavior of both plots. In spite of these encouraging results, it has also been shown that when the same calculations are performed on grids of different constructions, conflicting grid convergence results can be obtained. This is illustrated in Fig. 15b, in which a separate grid convergence study for the DPW III W1–W2 case was performed on a different family of uniformly refined grids, leading to a different set of nearly straight-line plots that do not coincide with the previous plots. Because these two grid convergence studies make use of different families of grids, the use of a single grid index value to characterize both families is inappropriate, and there is no reason to expect the plots to overlap using this measure. However, consistent grid convergence behavior dictates that these curves should asymptote to the same values in the limit of infinite resolution, which is clearly not the case in this experiment, as seen from Fig. 15b.

On the one hand, due to the wide variation in spatial grid resolution, it is entirely possible that some regions of the domain remain severely under-resolved, even for the finest meshes of the sequence, and that the apparent straight-line grid convergence indicative of asymptotic error behavior will be altered as these regions become more highly resolved. In this case, one or both of the convergence plots will eventually deviate from the currently observed nearly straight-line behavior as the grids are further refined, and results from both families will eventually converge together.

On the other hand, another more troubling possibility is that these two families of grids continue converging to different values as the

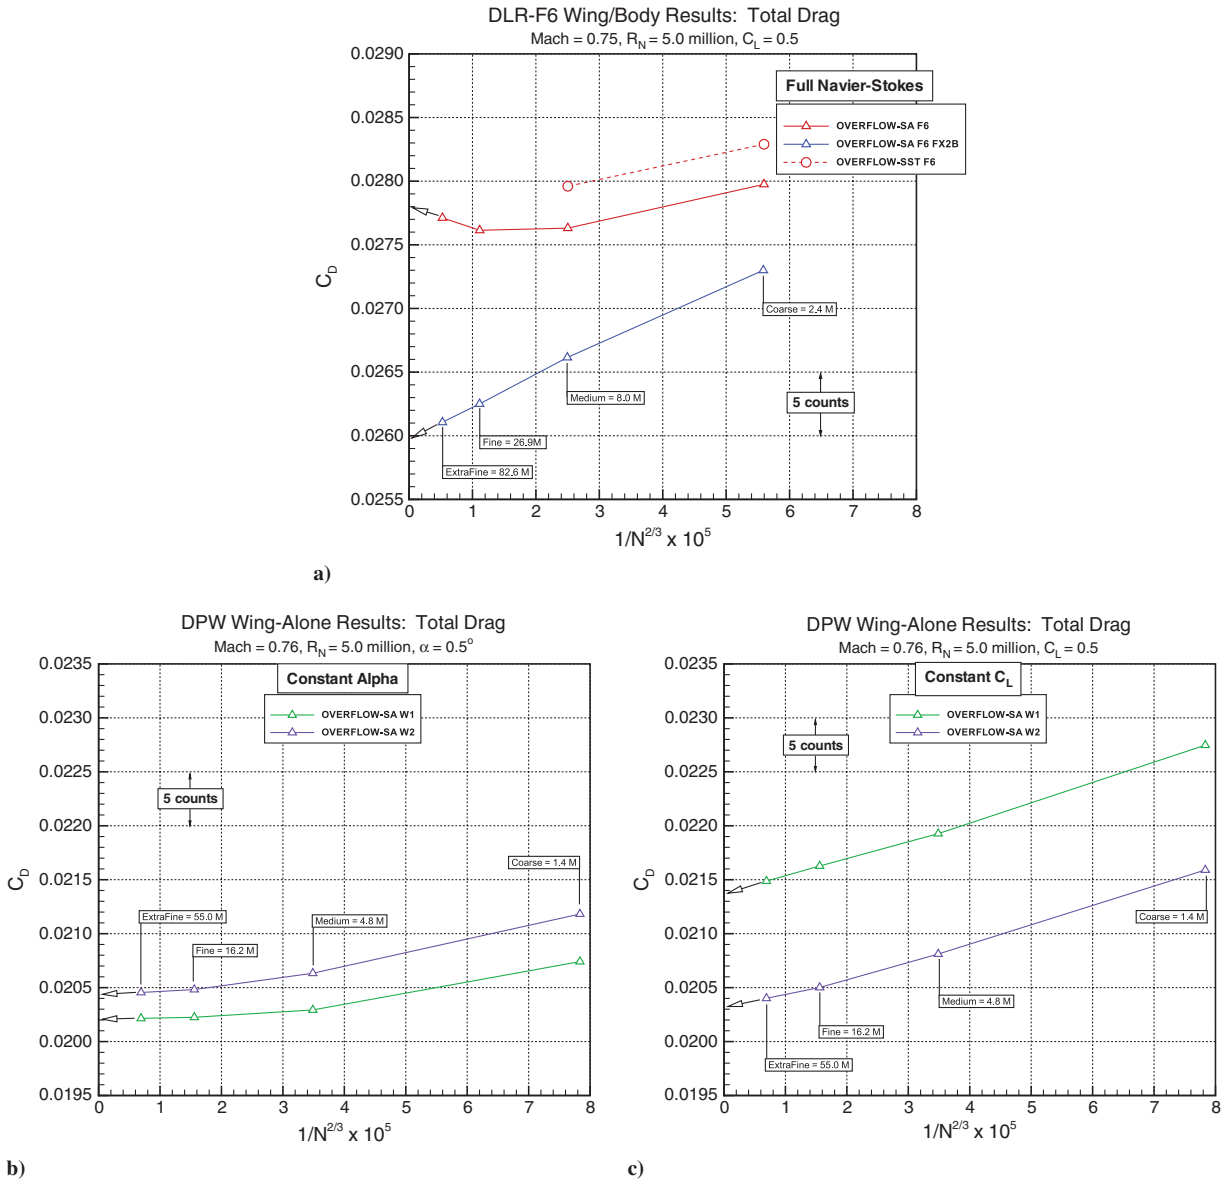

Fig. 13 Grid convergence results from DPW III using overset structured-mesh solver OVERFLOW: a) drag convergence for DLR-F6 with and without FX2B fairing at constant lift, b) drag convergence for wing-alone geometries at constant angle of attack, and c) drag convergence for wing-alone geometries at constant lift (reproduced from Sclafani et al. [22]).

grids are further refined, which would indicate that at least one of the grid configurations is inconsistent [i.e., produces errors of order $\mathcal{O}(1)$]. Obviously, to thoroughly investigate this behavior requires the ability to generate and employ much finer grids than currently available.

Other workshop results for this case, such as those previously shown in Figs. 13b and 13c, have suggested that more consistent grid convergence behavior for drag is obtained when these numerical experiments are performed at fixed-lift conditions. To explore this phenomena, the convergence of lift at fixed incidence for the DPW III wing-alone cases using the NSU3D unstructured-mesh solver is plotted in Fig. 16a, in which it is observed that the lift values computed on the two different families of grids appear to be converging to different values. In Fig. 16b, the *idealized* drag values computed at fixed incidence using NSU3D are plotted for both grid families, and these are seen to converge to similar values in the continuous limit.

These results indicate that the apparent grid convergence inconsistencies are due principally to the prediction of lift values, which in turn affect the drag values through the induced drag. However, as stated previously, consistent grid convergence implies that all computed values such as drag and lift ultimately should converge at fixed incidence.

C. Grid Quality

The preceding grid convergence studies highlight the importance of considering some measure of grid quality for specific aerodynamic simulations objectives. Here, we implicitly make a distinction between grid quality and grid resolution, noting that a grid of poor quality can be refined extensively and still lead to inaccurate solutions. Thus, in a grid convergence study, a grid-quality metric may be expected to apply equally to an entire sequence of grids from within the same family.

Although grid quality is not directly related to grid resolution, one aspect of grid quality is related to the *relative distribution* of grid resolution in space. For example, a globally fine grid with limited trailing-edge resolution would qualify as a poor-quality grid for aerodynamic drag prediction. Other characteristics that contribute to grid quality include the grid topology, element type, aspect ratio, and orthogonality for structured or semistructured (boundary-layer) grids. Grid-quality metrics may differ substantially for different

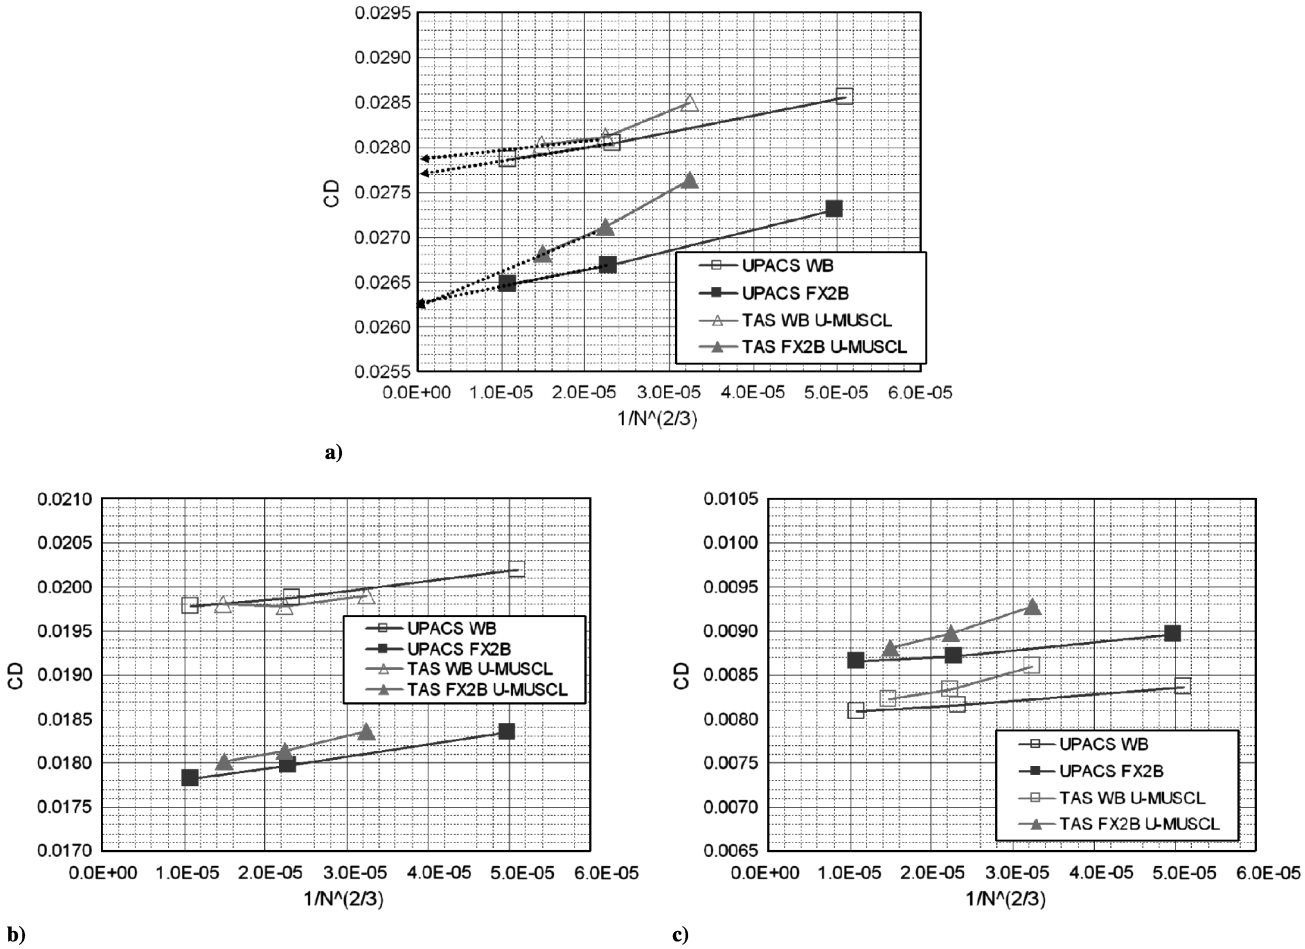

Fig. 14 Grid convergence for wing-body fairing configuration using block-structured (UPACS) and unstructured (TAS) mesh solvers showing different convergence slopes of both methods, and consistent convergence of both methods to similar infinite resolution drag values as indicated by the extrapolation arrows for standard and modified (U-MUSCL) discretizations: a) drag for complete wing-body configuration, b) drag due to wing component alone, and c) drag due to body component alone.

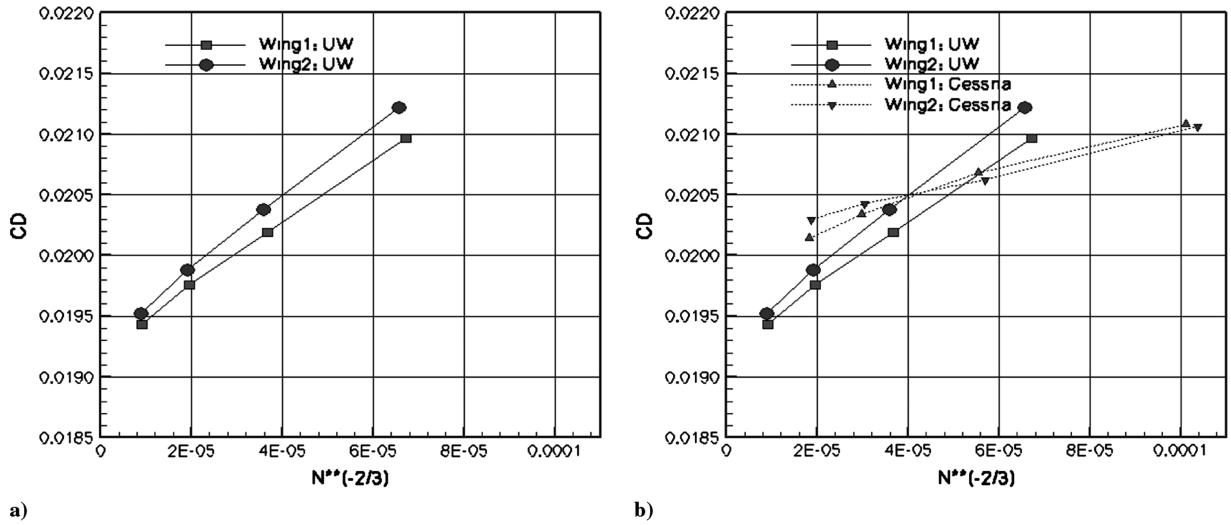

Fig. 15 Grid convergence plots for drag at fixed angle of attack for wing-1 and Wing2 configurations at Mach = 0.76, $Re = 5 \times 10^6$, $\alpha = 0.5$ deg using NSU3D unstructured-mesh solver for a) Original workshop meshes ranging from 1.8 to 37 million points, and b) comparison with second family of grids generated at Cessna Aircraft Co.

types of meshes. Experience-based guidelines for generating high-quality overset meshes for absolute drag prediction have been outlined by Vassberg et al. [28]. These include careful consideration of grid topology, avoiding singular lines in critical regions as well as locating overlap regions in noncritical areas, and matching mesh resolution in overlap regions. For unstructured meshes, solution accuracy may be affected by the choice of different element types. Although many unstructured-mesh solvers make use of fully tetrahedral elements, the most common approach is to employ prismatic elements in the highly stretched boundary-layer regions,

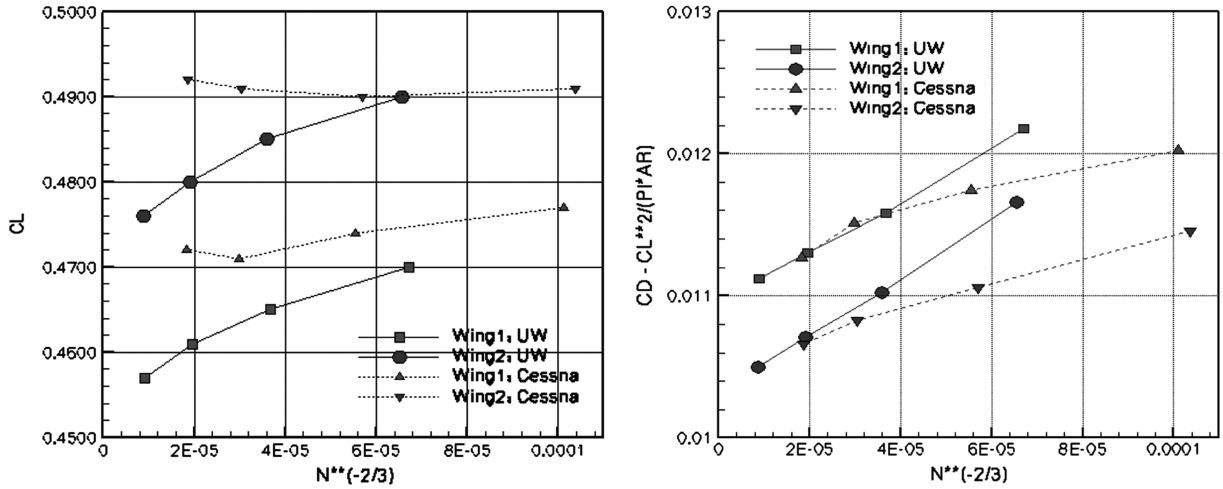

a) b)

Fig. 16 Comparison of both grid convergence at fixed angle of attack for wing-1 and Wing2 configurations at Mach = 0.76, $Re = 5 \times 10^6$, $\alpha = 0.5$ deg using NSU3D unstructured-mesh solver for a) lift coefficients and b) idealized drag coefficients.

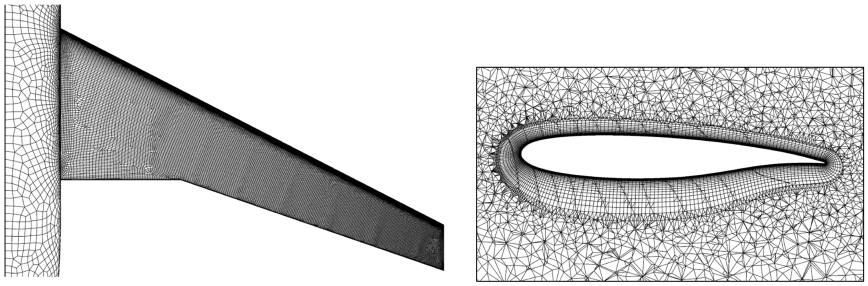

a) b)

Fig. 17 Unstructured hexahedral-tetrahedral mesh used in Brodersen et al. [24]: a) surface grid and b) field cut at 30% span location.

which results in semistructured, nearly orthogonal, meshes in these areas. A precise quantification of the accuracy effects of using fully tetrahedral meshes versus hybrid tetrahedral-prismatic meshes is still lacking, and this effect can be expected to differ for vertex-based and cell-based unstructured-mesh discretizations. Other types of elements such as unstructured sets of hexahedra may also be employed, with possibly beneficial effects for cases with multiple stretching directions, such as cases involving spanwise stretching, because the use of hexahedral elements avoids the formation of small-angle elements in such situations. Figure 17 depicts an unstructured mesh with hexahedral elements in the near-body region and tetrahedral elements in far-field regions employed by Brodersen et al. [24] in the computation of the DPW III DLR-F6 and FX2B geometries. Comparisons between the results computed on this hexahedral-tetrahedral mesh and the more traditional mixed tetrahedral/prismatic meshes using the DLR tau code have shown minimal differences in predicted drag values for the FX2B test case, although somewhat larger differences were seen for the DLR-F6 separated-flow cases. For structured-mesh simulations, different grid topologies in the wing-root junction area have been investigated by Murayama and Yamamoto [23]. The two different mesh topologies are depicted in Fig. 18, with the extruded topology shown in Fig. 18a corresponding to the traditional approach used by most workshop participants. Differences in predicted drag values using these two different structured grid topologies were found to be minimal for the attached-flow cases of the FX2B configuration, and more important differences were observed for the DLR-F6 case with separated-flow regions. The insensitivity of the attached-flow results to grid element type or topology is encouraging, and the differences observed in the separated-flow cases are not necessarily unexpected, given the

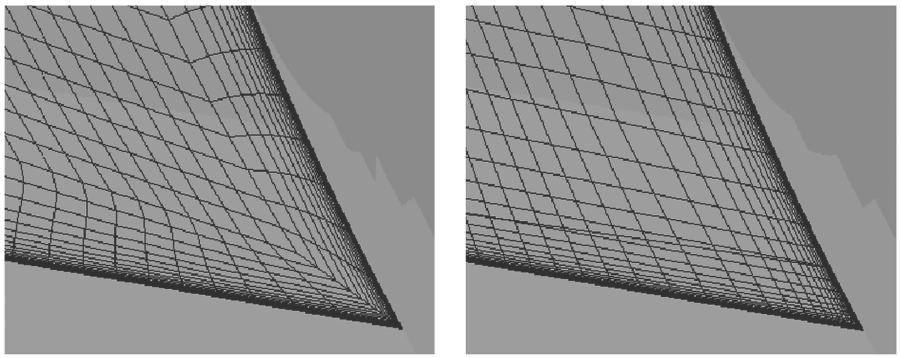

a) b)

Fig. 18 Different grid topologies in wing-root junction area used in Murayama and Yamamoto [23].

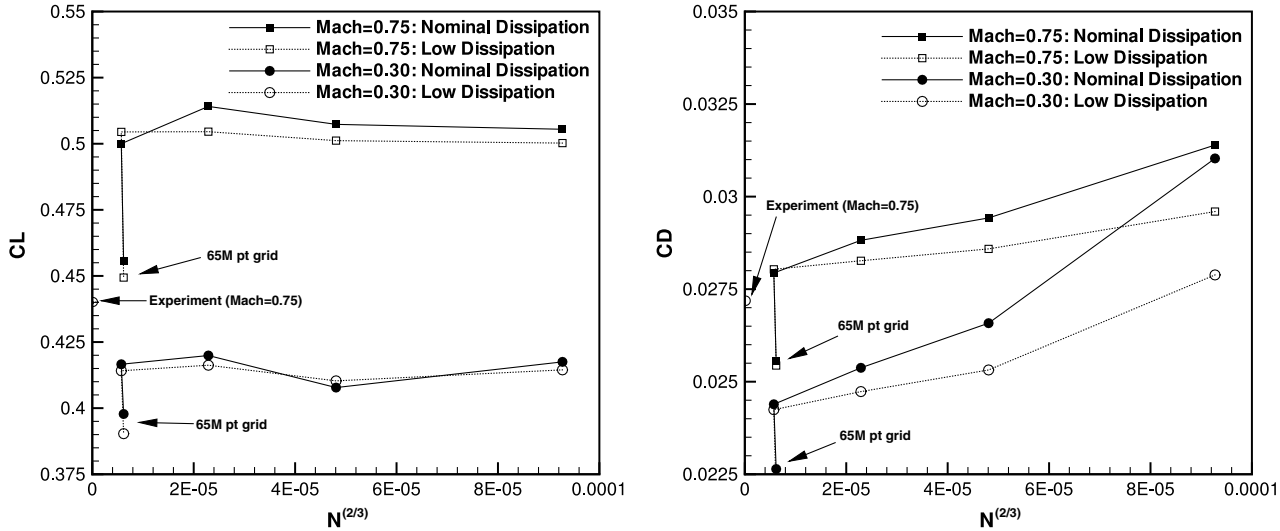

a) b)

Fig. 19 Comparison of a) computed lift and b) computed drag coefficients versus the number of grid points to the $-2/3$ power for transonic ($M = 0.75$) and subsonic ($M = 0.3$) cases at 0 deg incidence on sequence of uniformly refined unstructured meshes and alternate mesh of 65 million point for DLR-F6 configuration using the NSU3D solver (reproduced from Mavriplis [20]).

increased sensitivity of the results to grid effects in general and the possibly differing resolution in these areas for the different grid topologies and element types.

On the other hand, various situations have been documented in which different grid constructions have led to substantially different aerodynamic predictions using the same solver, even at high resolution. An example of such behavior is illustrated in Fig. 19, reproduced from Mavriplis [20], in which the computed lift and drag coefficients for the DLR-F6 wing-body configuration from DPW II on a family of uniformly refined meshes ranging up to 72 million points are compared with results obtained on an alternate mesh of 65 million points. Results on the first set of meshes display reasonably consistent grid convergence, going from 1 up to 72 million grid points, for both transonic (Mach $= 0.75$) and subsonic (Mach $= 0.3$) test cases, and values computed using two differing values of the dissipation parameter are seen to converge fairly consistently to similar infinite resolution values. However, the lift and drag values computed on the 65-million-point grid, which does not belong to the same family of grids, deviate substantially from the path of grid convergence defined by the first set of grids. A 10% drop in lift and corresponding decrease in drag are observed between the 72-million-point grid and the 65-million-point grid. Note that it is strictly not valid to plot these two grids using the same grid index, because they are not representative of a single family of grids. However, this is done in this case to illustrate the abrupt changes between the two finest grids. It remains very surprising and unexpected that there could be such a large difference between two solutions computed on such fine grids (72 million versus 65 million points), especially in light of the apparent grid convergence path determined by the original family of grids. Note that these trends (lower lift and drag) are also observed for the subsonic (Mach $= 0.3$) case and have been reproduced using the FUN3D unstructured-mesh flow solver by Lee-Rausch et al. [36], verifying that these results are not simply the product of anomalous flow solver behavior. Figure 20 illustrates details of the mesh topology in the blunt-wing trailing-edge region, in which spanwise stretching for the 72-million-point mesh is evident and isotropic resolution is observed for the 65-million-point mesh. The source of the discrepancies in the results computed on these different meshes remains unclear, although one may speculate that the use of anisotropic cells in the trailing-edge region results in the prediction of smaller trailing-edge separation regions, which in turn produces higher lift values when integrated along the span of the wing. On the other hand, the demands of isotropic trailing-edge resolution in the 65-million-point mesh may result in inadequate resolution for this mesh in other critical regions of the domain, which in turn may degrade the overall solution accuracy.

These results have prompted other follow-up studies using structured and unstructured meshes for this configuration. Sclafani et al. [22] studied the effect of additional trailing-edge refinement and lower cell aspect ratios (close to unity) using overset structured meshes, but they were found to have only minimal effect on overall drag prediction, and Murayama and Yamamoto [23] found

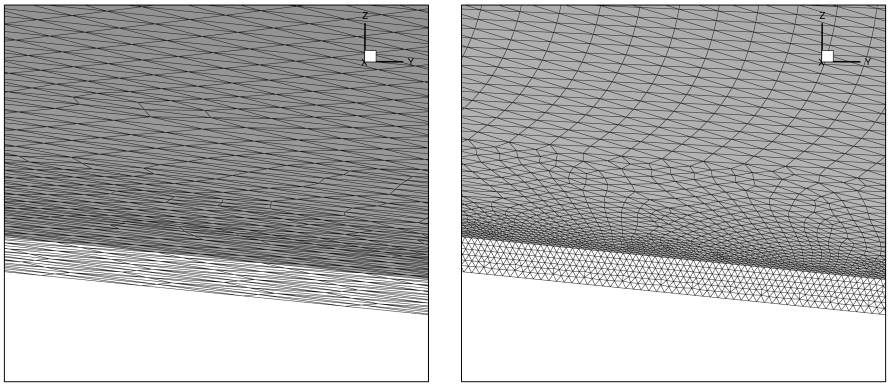

a) b)

Fig. 20 Comparison between surface resolution for a) 72-million-point grid and b) 65-million-point grid b) in trailing-edge region in the vicinity of 60% span region illustrating spanwise stretching for the 72-million-point grid.

additional trailing-edge refinement using an unstructured-mesh methodology to have limited impact on predicted drag values.

This case highlights the importance and difficulties of developing useful mesh-quality metrics. In this instance, we have two highly resolved grids that yield substantially differing results, which contrast with other grid refinement studies, but which have been verified independently by two unstructured-mesh solvers. If a correlation between solution accuracy and spanwise cell aspect ratios could be demonstrated conclusively, then a grid-quality metric could be established for differentiating between such mesh configurations (although it is entirely possible that other grid characteristics are responsible for the observed discrepancies). Many metrics have been proposed and used over the years, including orthogonality measures (for structured or semistructured grids), dihedral angles and vertex degree for unstructured grids, and min / max ratios of consecutive cell volumes, to name a few. However, the relationship between these metrics and overall solution accuracy remains tenuous in most cases and few, if any, of these metrics have been adopted universally. The situation is further complicated by the complex interdependence between flow solver discretization and mesh topology. Nevertheless, if particular grid topologies or cell shapes could be shown to produce large or inconsistent discretization errors, then anomalous grid convergence results could be predicted and eventually avoided through grid-quality measures and a priori error estimation. Hence, the development of objective grid-quality metrics remains an important issue for advancing the state of the art in computational aerodynamics.

V. Conclusions

Because discretization error has been identified as one of the leading sources of error for computational aerodynamics, improvement in the state of the art will require advances in grid generation technology, in addition to the development of new metrics and diagnostic tools. At the same time, the complex interdependence between grid quality/resolution and flow solver performance and accuracy will dictate increased collaboration between the mesh generation and flow solution communities.

The continued development of robust and efficient high-quality mesh generation strategies remains an important objective. However, with increasing resolution demands, capabilities for generating meshes of several hundred million mesh points and beyond should be targeted. This may involve increased adoption and implementation of parallel mesh generation strategies, as well as more elaborate techniques such as in situ parallel mesh refinement, using coarser template meshes (block-structured or unstructured) to alleviate disk storage and I/O bottlenecks.

For grid convergence studies, improved tools for generating consistent families of uniformly refined meshes are required, particularly for unstructured meshes. Simultaneously, tools for assessing the compliance of generated grids with the necessary grid convergence metrics are required. For example, a grid interpolation scheme could be developed in which the average local resolution between two similar grids is compared at each cell and in which areas that exceed prescribed tolerances are flagged. Such a tool could also be extended to evaluate grid compliance with user-specified guidelines, such as the guidelines issued by the DPW committee. In fact, a good first test case would be to run some of the publicly available grids generated for the DPW workshops through such a tool.

Although prediction of absolute drag remains a difficult problem, accurate prediction of incremental drag, such as the change in drag between two closely related configurations, is of the utmost importance for guiding design decisions. Although accurate incremental drag prediction is often less demanding than absolute drag prediction, the DPW results have shown that for cases with separated flow, even reliable incremental drag prediction can be difficult to achieve. To obtain high-quality incremental predictions on two related configurations, similar grid topologies and resolution must be maintained as much as possible. Thus, grid generation techniques that guarantee minimal changes between configurations are desirable. For cases in which only shape changes occur, with no changes in geometric topology, mesh deformation techniques may be most appropriate for constructing self-similar meshes on configurations of interest.

The remaining areas of interest can be divided into a priori error estimation and a posteriori error estimation. A priori error estimation involves the examination of the mesh before flow solution calculations to estimate the suitability of the mesh for the required calculations. This phase involves the development of suitable mesh metrics for predicting solution error. At a minimum, this means computing and validating specific metrics, such as aspect ratio, dihedral angle, etc. A suite of mesh metrics could be built up and published or kernels for implementation could be made available online to speed adoption. More elaborate a priori error-estimation techniques would include devising numerical tests for specific grids and discretizations, such as gradient reconstruction techniques, etc., which could be run on the subject grid in a single pass and report approximation accuracy.

Finally, the area of a posteriori error estimation holds great long-term promise for improving overall solution accuracy, particularly when combined with adaptive meshing strategies. As mentioned previously, solution gradient-based error-estimation techniques have met with limited success over the years in this area. On the other hand, adjoint-based approaches have shown great promise in this area and have been investigated in simplified settings for some time now [31–33]. In fact, an overview of this approach was given at the DPW I meeting in 2001 [37], citing some recent two-dimensional work at the time [31] and outlining the promise this approach holds for aerodynamic drag prediction. However, nearly 7 years later, there is still no industrial-strength adjoint-based adaptive meshing strategy in production use. This slow rate of progress is a testament to the complex issues involved in bringing such ideas to fruition, which include not only the mechanics of mesh refinement and dynamic load balancing, but other issues such as adjoint solution technology and CAD geometry representation in a dynamic parallel environment. The way forward will require increased collaboration among all of these different disciplines.

References

[1] Levy, D. W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R. A., Pirzadeh, S., and Hemsch, M. J., "Summary of Data from the First AIAA CFD Drag Prediction Workshop," AIAA Paper 2002-0841.

[2] Mavriplis, D. J., and Levy, D. W., "Transonic Drag Prediction Using an Unstructured Multigrid Solver," *Journal of Aircraft*, Vol. 42, No. 4, 2005, pp. 887–893. doi:10.2514/1.8233

[3] "2nd AIAA CFD Drag Prediction Workshop." Orlando, FL, 2003, http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/Workshop2.

[4] "3rd AIAA CFD Drag Prediction Workshop." San Francisco, 2006, http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/Workshop3.

[5] Vassberg, J. C., Sclafani, A., and DeHaan, M., "A Wing-Body Fairing Design for the DLR-F6 Model: A DPW-III Case Study," AIAA Paper 2005-4730.

[6] Levy, D. W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R. A., Pirzadeh, S., and Hemsch, M. J., "Data Summary from the First AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 40, No. 5, 2003, pp. 875–882. doi:10.2514/2.6877

[7] Laflin, K., Klausmeyer, S. M., Zickuhr, T., Vassberg, J. C., Wahls, R. A., Morrison, J. H., Brodersen, O. P., Rakowitz, M. E., Tinoco, E. N., and Godard, J., "Data Summary from Second AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 42, No. 5, 2005, pp. 1165–1178. doi:10.2514/1.10771

[8] Vassberg, J. C., Tinoco, E. N., Mani, M., Brodersen, O. P., Eisfeld, B., Wahls, R. A., Morrison, J. H., Zickuhr, T., Laflin, K. R., and Mavriplis, D. J., "Summary of the Third AIAA CFD Drag Prediction Workshop," AIAA Paper 2007-0260.

[9] Hemsch, M., "Statistical Analysis of CFD Solutions from the Drag Prediction Workshop," *Journal of Aircraft*, Vol. 41, No. 1, 2004, pp. 95–103. doi:10.2514/1.1796

[10] Vassberg, J. C., Buning, P. G., and Rumsey, C. L., "Drag Prediction for

the DLR-F4 Wing/Body Using OVERFLOW and CFL3D on an Overset Mesh," AIAA Paper 2002-0840.

[11] Pirzadeh, S. Z., and Frink, N. T., "Assessment of the Unstructured Grid Software TetrUSS for Drag Prediction of the DLR-F4 Configuration," AIAA Paper 2002-0839.

[12] Rakowitz, M., Eisfeld, B., Schwamborn, D., and Sutcliffe, M., "Structured and Unstructured Computations on the DLR-F4 wing-body configuration," *Journal of Aircraft*, Vol. 40, No. 2, 2003, pp. 256–264. doi:10.2514/2.3117

[13] Lee-Rausch, E. M., Buning, P. G., Morrison, J. H., Park, M. A., Rivers, S. M., Rumsey, C. L., and Mavriplis, D. J., "CFD Sensitivity Analysis of a Drag Prediction Workshop Wing/Body Transport Configuration," AIAA Paper 2003-3400.

[14] Lee-Rausch, E. M., Frink, N. T., Mavriplis, D. J., Rausch, R. D., and Milholen, W. E., "Transonic Drag Prediction on a DLR-F6 Transport Configuration Using Unstructured Grid Solvers," AIAA Paper 2004-0554.

[15] Sclafani, A., DeHaan, M., and Vassberg, J. C., "OVERFLOW Drag Prediction for the DLR-F6 Transport Configuration: A DPW-II Case Study," AIAA Paper 2004-0393.

[16] Rumsey, C., Rivers, S., and Morrison, J., "Study of CFD Variation on Transport Configurations from the Second Drag Prediction Workshop," AIAA Paper 2004-0394.

[17] Yamamoto, K., Ochi, A., Shima, E., and Takaki, R., "CFD Sensitivity of Drag Prediction on DLR-F6 Configuration by Structured Method and Unstructured Method," AIAA Paper 2004-0398.

[18] Brodersen, O. P., Rakowitz, M., Amant, S., Larrieu, P., Destarac, D., and Suttclie, M., "Airbus, ONERA and DLR Results from the Second AIAA Drag Prediction Workshop," *Journal of Aircraft*, Vol. 42, No. 4, 2005, pp. 932–940. doi:10.2514/1.8662

[19] Langtry, R. B., Kuntz, M., and Menter, F., "Drag Prediction of Engine-Airframe Interference Effects with CFX-5," *Journal of Aircraft*, Vol. 42, No. 6, 2005, pp. 1523–1529. doi:10.2514/1.12728

[20] Mavriplis, D. J., "Grid Resolution Study of a Drag Prediction Workshop Configuration Using the NSU3D Unstructured Mesh Solver," AIAA Paper 2005-4729.

[21] Tinoco, E., and Su, T., "Drag Prediction with the Zeus/CFL3D System," AIAA Paper 2004-0552.

[22] Sclafani, A., Vassberg, J., Harrison, N., Rumsey, C., DeHaan, M., and Morrison, J., "Drag Prediction for the DLR-F6 Wing/Body and DPW Wing Configurations Using CFL3D and OVERFLOW on an Overset Mesh," AIAA Paper 2007-0257.

[23] Murayama, M., and Yamamoto, K., "Comparison Study of Drag Prediction for the Third CFD Drag Prediction Workshop by Structured and Unstructured Mesh Method," AIAA Paper 2007-0258.

[24] Brodersen, O., Eisfeld, B., Raddatz, J., and Frohnapfel, P., "DLR Results from the Third AIAA CFD Drag Prediction Workshop," AIAA Paper 2007-0258.

[25] Tinoco, E., Winkler, C., Mani, M., and Venkatakrishnan, V., "Structured and Unstructured Solvers for the Third CFD Drag Prediction Workshop," AIAA Paper 2007-0255.

[26] Morrison, J. H., and Hemsch, M., "Statistical Analysis of CFD Solutions from the Third AIAA Drag Prediction Workshop," AIAA Paper 2007-0254.

[27] Mavriplis, D. J., "Results from the Third Drag Prediction Workshop Using the NSU3D Unstructured Mesh Solver," AIAA Paper 2007-0256.

[28] Vassberg, J. C., deHaan, M. A., and Sclafani, T. J., "Grid Generation Requirements for Accurate Drag Predictions Based on OVERFLOW Calculations," AIAA Paper 2003-4124.

[29] Anderson, W. K., Bonhaus, D. L., McGhee, R. J., and Walker, B. S., "Navier-Stokes Computations and Experimental Comparisons for Multielement Airfoil Configurations," *Journal of Aircraft*, Vol. 32, No. 6, 1995, pp. 1246–1253. doi:10.2514/3.46871

[30] Mavriplis, D. J., "Unstructured Mesh Discretizations and Solvers for Computational Aerodynamics," AIAA Paper 2007-3955.

[31] Venditti, D., and Darmofal, D., "Grid Adaptation for Functional Outputs: Application to 2-D Inviscid Compressible Flow," *Journal of Computational Physics*, Vol. 176, No. 1, 2002, pp. 40–69. doi:10.1006/jcph.2001.6967

[32] Mueller, J. D., and Giles, M. B., "Solution Adaptive Mesh Refinement Using Adjoint Error Analysis," AIAA Paper 2001-2550.

[33] Park, M., "Adjoint-Based, Three-Dimensional Error Prediction and Grid Adaptation," AIAA Paper 2002-3286.

[34] Salas, M. D., "Some Observations on Grid Convergence," *Computers and Fluids*, Vol. 35, No. 7, 2006, pp. 688–692. doi:10.1016/j.compfluid.2006.01.003

[35] Pirzadeh, S., "Three-Dimensional Unstructured Viscous Grids by the Advancing-Layers Method," *AIAA Journal*, Vol. 34, No. 1, 1996, pp. 43–49. doi:10.2514/3.13019

[36] Lee-Rausch, E. M., Park, M., Nielsen, E., Jones, W., and Hammond, D., "Parallel Adjoint-Based Error Estimation and Anisotropic Grid Adaptation for Three-Dimensional Aerospace Applications," AIAA Paper 2005-4842.

[37] *AIAA CFD Drag Prediction Workshop*, Anaheim, CA, 2001, http://www.aiaa.org/tc/apa/dragpredworkshop/dpw.html.