26th AIAA Applied Aerodynamics Conference
18 - 21 August 2008, Honolulu, Hawaii

AIAA 2008-6226

A Pilot Project in Preparation of an Aerodynamic Optimization Workshop with Lessons Learned

John C. Vassberg,[*] *Antony Jameson,*[†]
Sergey Peigin,[‡] *Boris Epstein,*[§]
Dino L. Roman[**] *and Neal A. Harrison*[††]

Abstract

This paper is an addendum to a recent publication on a Comparative Study of 3D Wing Drag Minimization by Different Optimization Techniques; with a focus on a path forward towards the development of a larger-scale Aerodynamic Optimization Workshop. Three independent optimization efforts were conducted concurrently, each without knowledge of the others' results. The resulting "optimal" wings were then cross analyzed by each of the CFD methods used in the optimization phase, and by a fourth CFD method which was not used to develop any of the geometries. The incremental improvements, as predicted by these cross analyses, are compared to establish a level of uncertainty in the performance improvements and identify any unusual trends. The baseline geometry of this pilot project is the DPW-W1 wing, which was placed into the public domain by the Drag Prediction Workshop as a test case for the third workshop of June 2006. Lessons learned from this collaborative effort are documented and suggestions are provided to aid in the preparation of a possible future workshop on aerodynamic shape optimization.

1.0 Introduction

During AIAA conferences of the last several years, the first four authors have been discussing the idea of collaborating on a project related to aerodynamic shape optimization. Specifically, we were interested in devising a means by which vastly different aerodynamic optimization techniques could be evaluated against each other in an open, honest and fair manner. We also wanted this process to be able to accommodate up to thirty participants from around the world. The motivation for targeting this scale of participation is based on the level of response demonstrated by the AIAA CFD Drag Prediction Workshop (DPW) series. DPW has shown that a grass-roots campaign can generate substantial cooperation across the aeronautics industry to study a problem of

[*] Boeing Technical Fellow, The Boeing Company, Huntington Beach, CA, Associate Fellow AIAA
[†] Thomas V. Jones Professor of Engineering, Stanford University, CA, Fellow AIAA
[‡] Professor, Aerodynamic Department, Israel Aircraft Industries, Israel, Senior Member AIAA
[§] Professor, Computer Sciences Department, Academic College of Tel-Aviv, Israel, Member AIAA
[**] Associate Technical Fellow, The Boeing Company, Huntington Beach, CA, Senior Member AIAA
[††] Engineer/Scientist, The Boeing Company, Huntington Beach, CA, Member AIAA

common interest, and that there is great value in developing public-domain databases for method validation and verification efforts. We anticipate that there will be similar interest in an Aerodynamic Optimization Workshop (AOW).

During the summer of 2006 at the AIAA conferences in San Francisco, we agreed that we would broach the subject of a pilot project on aerodynamic optimization with our respective organizations, and to our delight, our managements were very much interested in proceeding with such collaboration. During the fall of 2006, we formalized our intent to collaborate with agreements between The Boeing Company and Israel Aircraft Industries. Our next face-to-face meeting was in Reno, NV, January 2007. At this gathering, a technical approach was formulated and finalized during the next few weeks.

This paper is organized as follows. The technical approach of this project is discussed in Section 2. The test cases are defined in Section 3. Description of the numerical methods are provided in Section 4. The results of the optimization efforts and cross analyses are given in Section 5. Lessons learned from this experience and suggestions for the planning of a possible AOW are discussed in Section 6. Tables are embedded in the body of text, while all figures are appended at the end of this document.

2.0 Technical Approach

In order to establish an open, honest, and fair evaluation of several vastly different aerodynamic optimization methods, at set of general guidelines were defined for this pilot project. First, a fairly simple geometric problem of interest would be selected. Second, the optimizations would be conducted completely independent of one another. Third, when all optimizations are finalized, all resulting geometries would be simultaneously distributed to each other. Fourth, independent cross analyses of all geometries would be conducted using all the CFD methods used in the optimization studies. Fifth, an additional CFD method would be utilized for another independent cross analysis. Sixth, when the cross analyses are completed, the results would then be simultaneously distributed to each participant. Seventh, all data generated during the collaborative pilot project would be owned by all parties involved. Finally, the results of the pilot project would be documented and published to the public domain.

As noted above, this collaborative study was comprised of three concurrent and independent optimization activities, followed by four concurrent and independent cross-analysis efforts. The first author was responsible for the coordination and synchronization of the independent efforts. He acted as the hub in the collection and redistribution of geometries and other pertinent information for the duration of this collaboration. Although having access to most of the data during the independent activities, special attention was given to not contaminate any of the studies by prematurely releasing any of this information between parties. The three optimization activities were conducted by: 1) the third and fourth authors, 2) the fifth author, and 3) the second author. The four post-optimization cross analyses were also conducted by these same teams, plus a fourth independent analysis being performed by the last author.

The next section describes the test cases selected for this collaboration.

3.0 Test Cases

For this pilot project, two very simple test cases are defined; the first being a single-point optimization, while the second is a two-point optimization. The baseline geometry is the DPW-W1 wing-only configuration from DPW-III; see Reference [1]. The flow condition of the single-point optimization is a C_L of 0.5, a Mach number of 0.76M, and a chord Reynolds number of Re=5 million. This condition is the design point of the DPW-W1 wing. The two-point optimization augments the design point with a similar C_L and Re condition at a Mach of 0.78M. The geometric constraints are established with respect to the baseline DPW-W1 wing. First, the wing planform is to remain fixed. Second, the maximum thickness must be maintained; however, its location is allowed to float. Third, two beam constraints are included. To mimic beam constraints, the thickness of the modified wing must be greater than or equal to the baseline wing at nondimensional chordwise positions of x/c = [0.20 , 0.75]. The objective function for the single-point optimization is the wing drag augmented with a pseudo-trim penalty. The trim penalty is given as a function of incremental pitching moment relative to the baseline DPW-W1 wing. If the incremental pitching moment is positive (more nose up), then this penalty is set to zero. If this increment is negative (more nose down), then the trim penalty is +1 count per -0.01 in ΔC_M. The drag with trim penalty is referred heretofore as an adjusted drag. The objective function for the two-point optimization is the summation of the adjusted drags at both flight conditions.

Inclusion of a minimum constraint on the maximum lift coefficient at 0.2M was also considered. However, this was eventually dropped since determining maximum lift by CFD is not only unreliable, but also expensive as an alpha sweep is required.

Brief descriptions of the optimization & analysis methods of this work are presented next.

4.0 Optimization and Analysis Methods

Three Reynolds-Averaged Navier-Stokes (RANS) optimizations methods were utilized in this study; they are: OPTIMAS, MDOPT, and SYN107.

OPTIMAS uses a floating-point genetic algorithm (GA) as its search engine; it was developed by the third and fourth authors. The design space utilized here is defined with 35 design variables comprised of camber, thickness and twist distributions at three defining airfoil stations. The GA optimization process cycles through a number of generations (~8) in search of a "global" optimum design. The population of each generation is comprised of approximately 200 geometries. The corresponding set of CFD flow solutions for each generation is run very efficiently in parallel. The analysis method of OPTIMAS is the NES CFD code, which utilized a grid of about 250,000 points in this

work. The OPTIMAS optimizations conducted here cost about three orders-of-magnitude more than an objective-function evaluation. For more information on OPTIMAS and NES, see References [2-5].

MDOPT is a response-surface method developed by The Boeing Company; see Reference [6]. The design space utilized here is defined with 35 design variables of camber, thickness and twist distributions at four defining airfoil stations. In order to establish an interpolated response surface (IRS), the design space is initially seeded with a number of geometries as defined by a design of experiments (DOE). The IRS is created for objective functions and constraints. This IRS is interrogated to locate a "global" optimum. However, due to the fact that the objective-function IRS does not exactly represent the CFD data, except at the pre-computed locations, the IRS is refined in the neighborhoods of local optimums. This process of IRS refinement is repeated until sufficient convergence is obtained. To augment the IRS optimization, a direct driven (DD) approach is also available; design-variable sensitivities are directly approximated using a finite-difference method. DD can be used to locate a local optimum once the IRS model identifies the neighborhood of the "global" optimum. The analysis method employed by this application of MDOPT is the TLNS3D RANS CFD method from NASA LaRC; see Reference [7]. In this work, the TLNS3D solutions utilized a grid with dimensions of (305x145x81); a total of 3,582,225 points. The MDOPT optimizations here cost about three orders-of-magnitude more than an objective-function evaluation.

SYN107 is a gradient-based method developed by the second author. The gradient is computed by solving an adjoint equation to the continuous RANS equations. The design space is automatically established by treating the geometry as a free surface, with thousands of design variables linked to the discrete surface points of the computational grid. This design space is then navigated towards a local optimum by recasting the gradient in a Sobolev space. This significantly reduces the number of design cycles required to reach convergence of the optimum shape. The analysis method of SYN107 is the FLO107 CFD code, also developed by Jameson. The grid utilized for this work consists of 818,545 points. A SYN107 optimization costs about one order-of-magnitude more than an objective-function evaluation. For more information on SYN107 and FLO107, see References [9-11].

During the post-optimization cross-analysis phase of this collaborative effort, an additional CFD analysis method was utilized to provide an unbiased prediction of the performance increments for the various optimal geometries. This was the OVERFLOW code developed at NASA ARC/LaRC; see Reference [12]. OVERFLOW is a RANS solver based on overset grids, and is used throughout the industry for accurate drag predictions. OVERFLOW is applied here using best practices documented by Vassberg, *et al*. in Reference [13]. For the applications performed herein, the overset grid system is comprised of about 4,000,000 points. It is generally accepted by the authors that the OVERFLOW analyses on the high-density grids yielded the most accurate predictions of incremental drags obtained by any CFD method in this study.

The next section discusses the results and findings of this pilot project.

5.0 Results

This section provides a subset of the results originally documented by the authors in Reference [14]. The focus here, however, is to better understand the information encapsulated in the data from the cross analyses. First, the adjusted incremental improvements from the cross analyses are presented. Second, this information is reorganized into matrix layouts to highlight contradictory trends in the data. The suspect trends are explained such that similar data mining can be useful in the future.

Figures 1-3 illustrate the adjusted incremental drag reductions of seven optimal geometries as predicted by four different CFD methods; adjusted by the trim penalty. In these figures, O4 and O5 are the single-point and dual-point optimization results of the OPTIMAS method, respectively. Similarly, M5 and M7 correspond to the MDOPT results, while S4, S5, and S6 are for SYN107. Note that S5 is an additional single-point optimization performed at the M=0.78 condition. In these figures, the Gold, Blue, Green, and Red bars represent results from NES, TLNS3D, FLO107, and OVERFLOW, respectively. Note that the scatter between the four CFD incremental drag predictions is up to five counts at M=0.76, up to ten counts at M=0.78, and almost 13 counts for the aggregate of both conditions. Given these levels of scatter, it is difficult to determine for certain that any of the optimal geometries stand out above the others; nonetheless, more discussion on this will follow.

Figure 1 should be viewed with focus on the O4, M5, and S4 optimal geometries as their optimizations correspond directly to the flow condition of M=0.76 and C_L=0.5. According to OVERFLOW, the O5 geometry has realized a drag reduction of about 14.5 counts, while both M5 and S4 yield almost 14 counts. Again, the variation of predicted improvement between the CFD methods on these three geometries is 3-5 counts.

Figure 2 is included for completeness; these data are used for aggregate drag values to be presented later. Here, it is interesting to review the S5 optimal geometry generated by a single-point optimization at this flow condition of M=0.78 and C_L=0.5. OVERFLOW has predicted that S5 yields an adjusted drag improvement of about 32 counts over the baseline DPW-W1 wing. It is also interesting to investigate the results of O5, M7, and S6. These multipoint optimal geometries include the 0.78M flow condition as a driver in their designs. It is clear that their adjusted drag reductions, at this condition, are significantly larger than the corresponding designs of O4, M5, and S4. Another nice comparison is between S5 and S6; where S5 is a single-point optimum at this condition, while S6 is a multipoint optimum which includes this condition. At the 0.78M condition, the improvement of S5 should be larger than that of S6. In fact it is, however only by 3-4 counts, depending on the analysis method.

Figure 3 provides the aggregate adjusted drag improvements of the M=[0.76, 0.78] and C_L=0.5 flow conditions. The optimal geometries that correspond to this dual-point case are the O5, M7, and S6 designs. [The adjusted drag increments of each flow condition are depicted in Figures 1-2.] For example, the OVERFLOW shows the aggregate adjusted drag improvement of O5 is just over 40 counts, of M7 is about 37 counts, and of

S6 is about 39 counts. It is also interesting to note that the multipoint optimization of S6 should yield a performance improvement that is bound from above by the summation of improvements of the two single-point optimizations of S4 and S5, as estimated by FLO107. Summing the 0.76M drag reduction of S4 (12.9 cts.) with the 0.78M drag reduction of S5 (29.8 cts.) gives an aggregate improvement of 42.7 counts. This summation establishes an upper bound for the aggregate improvement for S6, which yielded a 36.6 drag reduction, or about 86% of that possible. However, what is most interesting about these dual-point optimizations is not so much in the final aggregate drag increments, but rather in the composite of these objective functions. For example, the 40 counts of O5 is comprised of 15.5 counts at M=0.76 and 24.5 counts at M=0.78 (a 39% / 61% split). Conversely, the aggregate 37 counts of M7 is a combination of 7.5 counts at M=0.76 and 29.5 counts at M=0.78 (a 20% / 80% split). Whereas the S6 aggregate of 39 counts is the sum of 10.5 counts at M=0.76 plus 28.5 counts at M=0.78 (a 27% / 73% split). These discrepancies are due to the fact that the three optimal geometries are three distinctly different geometries, despite being the results of an equivalent objective function. Figure 4 punctuates the differences in these geometries by providing their corresponding pressure distributions, compared with those of the baseline DPW-W1 wing at the M=0.76, C_L=0.5 flow condition. Note that the O5 (green chain-dot) pressures are characterized by a flat roof top terminated by the shock located at about mid-chord. On the lower surface, the O5 exhibits extra loading near the leading edge (LE) which is generated by a local LE droop in the airfoil's camber. Furthermore, the O5 carries more aft loading than the other wings. Both the M7 (red chain-double-dot) and the S6 (blue dash) have sloping roof top pressures into the shock located somewhat forward of mid-chord, with the M7 shock furthest forward.

Although the aerodynamic performance improvements from the three optimization methods fall within the scatter band of the CFD analyses, it does appear that O4 and O5 exhibit larger adjusted drag reductions at their respective design conditions than the other geometries. Refer to Figure 1 and compare O4, M5, and S4. Similarily in Figure 3 compare O5, M7, and S6. To better investigate these specific comparisons, the associated information has been reorganized and presented in Tables 1-4 in matrix form. The geometries are represented by the rows of the matrix, while the analysis methods are captured by the matrix columns.

Table 1 provides the adjusted incremental drags of the single-point optimal geometries O4, M5, and S4. Note that the cells of this matrix are color-coded to illustrate any bias that may exist between geometry and analysis method; gray cells indicate an unbiased result. For example, TLNS3D may favor M5, but should be unbiased towards O4 and S4. OVERFLOW should be unbiased across all of the geometries. In spite of these biases, it is obvious that the CFD methods unanimously concur that the O4 improvement exceeds that of either M5 or S4. This trend may be better seen in Table 2, where the deltas between geometry improvements are shown; the values of the top two rows are negative across all columns. This situation can be possible for a number of reasons. First, the O4 could be a more global optimum than either M5 or S6. Second, the O4 geometry is not supported in the design spaces of either MDOPT or SYN107. Third, an

error or inconsistency has been introduced into one or more of the optimization processes. Furthermore, there could be other causes.

In the case of this pilot project, a small inconsistency was inadvertently introduced in the baseline geometry of the OPTIMAS study, and then further compounded by a rounding error in the beam constraints. As noted in Section 3, the baseline geometry for this collaboration is the DPW-W1 wing. One characteristic of this wing is that its trailing edge is not sharp; it has a 0.5%-chord blunt base. Another characteristic is that it is defined by a 13.5%-chord thick airfoil section. Unfortunately, the automated gridding process of the NES flow solver requires a sharp trailing edge. In removing the blunt base of the DPW-W1 wing, and relofting the surface, the baseline wing used in the OPTIMAS study is only 13.3% thick. This discrepancy in maximum thickness cascaded into other discrepancies in the beam constraints. At x/c=0.20, the beam height should be greater than 12.306% chord; instead it was only limited by 12.0% chord. The beam height at x/c=0.75 should be greater than 6.748% chord; rather it was only bound by 5.9% chord.

In order to minimize the effect of the issue discussed in the previous paragraph, both baseline geometries (with and without trailing-edge bluntness) were analyzed by TLNS3D, FLO107, and OVERFLOW. The adjusted drag increments presented herein are based on using the sharp TE baseline for the O4 and O5 geometries, whereas the original DPW-W1 wing was used to establish the improvements of the M5, M7, S4, S5, and S6 configurations. However, this technique will only capture a portion of the differences that exist between the independent studies.

Tables 3-4 provide similar data as that of Tables 1-2, but for the multipoint optimization studies. These data reveal similar trends as discussed above.

Table 1: Adjusted Drag Improvements for the Single-Point Optimization Case.

0.76M Adjusted ΔCDs (cts)				
CASE	NES	TLNS3D	FLO107	OVERFLOW
O4	-13.5	-16.0	-13.0	-14.5
M5	-9.8	-11.0	-10.4	-13.8
S4	-8.7	-13.2	-12.9	-13.8

Table 2: Delta Improvements for the Single-Point Optimization Case.

0.76M Adjusted ΔCDs Deltas				
DELTA	NES	TLNS3D	FLO107	OVERFLOW
O4-M5	-3.7	-5.0	-2.6	-0.7
O4-S4	-4.8	-2.8	-0.1	-0.7
M5-O4	3.7	5.0	2.6	0.7
M5-S4	-1.1	2.2	2.5	0.0
S4-O4	4.8	2.8	0.1	0.7
S4-M5	1.1	-2.2	-2.5	0.0

Table 3: Adjusted Drag Improvements for the Multi-Point Optimization Case.

[0.76M + 0.78M] Adjusted ΔCDs (cts)				
CASE	NES	TLNS3D	FLO107	OVERFLOW
O5	-35.0	-42.4	-32.0	-40.3
M7	-33.1	-29.6	-26.4	-37.1
S6	-24.9	-37.5	-36.6	-38.8

Table 4: Delta Improvements for the Multi-Point Optimization Case.

[0.76M + 0.78M] Adjusted ΔCDs Deltas				
DELTA	NES	TLNS3D	FLO107	OVERFLOW
O5-M7	-1.9	-12.8	-5.6	-3.2
O5-S6	-10.1	-4.8	4.6	-1.5
M7-O5	1.9	12.8	5.6	3.2
M7-S6	-8.2	7.9	10.2	1.7
S6-O5	10.1	4.8	-4.6	1.5
S6-M7	8.2	-7.9	-10.2	-1.7

6.0 Lessons Learned

The class of optimizations included in this pilot project can consume a considerable amount of computational resources as well as elapsed time. From a practical perspective, redoing even one of the optimization studies can become problematic. Because of the synchronizations adopted in this work to redistribute data, the issue of a baseline mismatch was not discovered until it was too late to correct. Nonetheless, the impetus of this pilot project was to identify unanticipated issues (such as those discussed above) before embarking on a larger-scale planning effort for an International Aerodynamic Optimization Workshop. This pilot project has been quite successful in this regard. We have begun to understand some of the unexpected issues that can arise between teams spread across multiple organizations, using vastly different optimization techniques. In constructing our technical approach, we enabled competition and cooperation to coexist in a collaborative working environment, with everyone in pursuit of an honest evaluation of the different methods being applied to an aerodynamic design problem of common interest. Participating in this pilot project has been a very rewarding and enlightening experience. The insights gained through this study help formulate the suggestions provided below in preparation of a future workshop on aerodynamic optimization.

Some of the differences in the resulting optimal designs addressed in Section 5 are attributable to variations in the techniques used to enforce the maximum thickness and beam constraints, aerodynamic penalties (trim drag), and even to the beginning baseline geometries used. Other causes of these differences are due to encountering a relatively flat design space for the flow condition of M=0.76, C_L=0.5, and Re=5 million. Here, the Mach-Sweep-C_L-Thickness characteristics of the baseline DPW-W1 wing allowed the

optimizers to essentially eliminate the shock drag, while twisting the wing to give near elliptic span loadings, only to deviate from this optimal load distribution to help manage the trim-drag penalty. The optimal objective for this flight condition, under the constraints specified, was simply too easy for the optimizers to find; there exists a large range of geometries that yield nearly the best that can be done. For example, Figure 5 illustrates the resulting geometries of two different optimizations by the same method, with the only difference being a variation on the pitching-moment constraint. In this side study, the pitching moment of the O1 design was not constrained, while that of O2 was constrained to match the pitching-moment of the baseline wing. As can be seen in this figure, the geometries are significantly different from each other. Yet their drags are quite comparable, being only 0.6 counts apart. Hence, the single-point optimization case of 0.76M, C_L=0.5 and Re=5 million is not well posed. On the other hand, the 0.78M condition did provide sufficient challenge.

Fundamentally, the technical approach chosen for this pilot project can provide a sound foundation for an Aerodynamic Optimization Workshop. The cross-analysis step of this approach is deemed especially necessary in order to establish an honest and fair evaluation of the various methods, including levels of uncertainty in the CFD predictions. However, in a large-scale setting, the costs of requiring every participant to analyze every other participant's geometry will escalate beyond reason; some compromise must be considered. For example, when the master coordinator of the workshop redistributes all the geometries to all the participants, he could assign four specific geometries to each participant that are mandatory for their independent analyses; analyzing the remaining shapes can be made optional. It might even be beneficial to mandate that each participant reanalyze his own geometry from the set broadcast. This can expose differences that may arise in the elapsed time between optimization and cross analyses, such as different versions of CFD software, different people running the optimization and analysis codes, different levels of convergence, or even slight variations in geometry definitions that can be introduced by relofting or reformatting steps. Choosing which four additional cases can be as straightforward as follows. Assume there are N participants, then for participant I, the master coordinator can assign geometries from participants I-2, I-1, I, I+1, and I+2; where -1, 0, N+1, and N+2 correspond to N-1, N, 1, and 2, respectively. This will generate a penta-diagonal matrix of assessments similar to those of Tables 1-4; Figure 6 depicts a template for a workshop of 26 participants. Recall that off-diagonal results are unbiased evaluations, while diagonal terms can have a bias between geometry and analysis method. Finally, anonymity of the geometries' origins should be maintained throughout the cross-analysis phase.

The test case should be of sufficient interest, yet not too complex. In this collaboration, a wing-only geometry was considered. However, a wing/body configuration may be more appropriate. Having a sharp (or essentially sharp) trailing edge on the wing could help avoid issues such as those encountered in this work. There is also benefit to adopting a configuration that is the subject of other public-domain efforts, such as the Common Research Model (CRM) being developed by NASA and US Industry. The CRM wing/body/horizontal-tail configuration is also the subject geometry for the next DPW.

This geometry will be released to the public domain and made available through the DPW website in August 2008. For more information on the CRM see Reference [15].

Geometric constraints such as maximum thickness or beam heights should be explicitly stated, as opposed to implied definitions such as those presented in Section 3. An independent check on the adherence to these constraints should also be part of the cross-analysis process. If a geometric constraint is violated, it should be reported. However, correcting for this type of discrepancy can be difficult at best.

In this effort, the objective function was drag augmented with a pitching-moment penalty at a given flow condition. Of specific interest is the constraint on lift coefficient. In the pilot project it was C_L=0.5. It is well known that drag is a strong function of lift. Hence, if the computed lift is close to, but not exactly the constrained value, an error can be introduced in the aerodynamic performance assessment. To account for this, an aerodynamic optimization should not only include the constraint on lift, it should incorporate a correction that accurately accounts for a small deviation of the actual computed lift to that of the target value. For example, the objective function could be based on the lift-to-drag ratio (L/D) of the configuration, or based on its idealized profile drag (C_D-C_L^2/(πAR)). However, with the understanding that only a small deviation from the targeted lift coefficient will be tolerated (*e.g.,* +/-0.001).

Pitching moment can be considered in the definition of the optimization test case. However, if pitching moment is constrained to either an explicit absolute quantity or an implied value relative to the baseline, a dilemma can be introduced during cross analyses. For example, this class of constraint may be found to be violated by some of the analyses and shown to be in adherence by other analyses. Hence, it may be better to include the effect of a pitching moment change as a pseudo-trim penalty, as was done in this pilot project. This type of penalty will seamlessly account for discrepancies in computed pitching moments between analysis methods.

The test case should challenge the optimizations. Hence, the baseline geometry should be of reasonably good design, and the combination of flow conditions and geometry constraints should not permit the optimizations to easily achieve shock-free designs.

In a workshop setting, it should be explicitly stated that each participant can only submit one optimized geometry. They can develop as many as they wish, but they must choose one to submit. Then, they must be prepared and willing to analyze the mandatory geometries assigned to them. If they want to cross analyze geometries beyond their assigned ones, they should be free to do so. Further filling in the empty cells of the cross-analysis matrix is beneficial, and this additional data should be willingly accepted.

As can be gleaned from the discussions above, there are numerous issues to consider in the planning of the larger-scale endeavor of an International Aerodynamic Optimization Workshop. The suggestions given here are by no means considered complete, or cast in stone, but rather are intended to spark further conversations throughout the community on this subject; we welcome the dialogue.

7.0 Summary

A pilot project in preparation of a larger-scale aerodynamic optimization workshop has been conducted by three teams working in collaboration. Three independent optimization efforts were conducted concurrently, each without knowledge of the others' results. The baseline geometry of this pilot project is the DPW-W1 wing. The resulting optimal wings were cross analyzed by each of the CFD methods used in the optimization phase, and also by a fourth CFD method not used to develop any of the geometries. The incremental improvements, as predicted by these cross analyses, are investigated in depth to identify any unusual aspects in the data. Data mining exposed an unexpected trend, and its root-cause was revealed as an inconsistency in the baseline geometries being used across the teams. Lessons learned from this collaborative effort are documented. Based on this experience, detailed suggestions are provided to aid in the preparation of a possible future workshop on aerodynamic shape optimization.

8.0 Acknowledgements

This work was supported by internal funding from The Boeing Company – Phantom Works, Israel Aircraft Industries, and Intelligent Aerodynamics, Inc. In particular, the authors thank Arnold Nathan, Adrian Rosenberg and Stefan Seror of Israel Aircraft Industries, and Tony Antani and Robb Gregg of The Boeing Company – Phantom Works for supporting this pilot project, as well as encouraging the technical collaborations between organizations.

9.0 References

1. AIAA CFD Drag Prediction Workshop
   website: http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/
   email: dpw@cessna.textron.com/
2. B Epstein & S Peigin. Constrained Aerodynamic Optimization of Three-Dimensional Wings Driven by Navier-Stokes Computations. *AIAA Journal*, **43**, 1946-1957, 2005.
3. B Epstein, T Rubin & S Seror. Accurate Multiblock Navier-Stokes Solver for Complex Aerodynamic Configurations. *AIAA Journal*, **41**, 582-594, 2003.
4. S Peigin, B Epstein, T Rubin & S Seror. Parallel Large Scale High Accuracy Navier-Stokes Computations on Distributed Memory Clusters. *Journal of Supercomputing*, **27**, 49-68, 2004.
5. S Seror, T Rubin, S Peigin & B Epstein. Implementation and Validation of the Spalart-Allmaras Turbulence Model for a Parallel CFD Code. *Journal of Aircraft*, **42**, 179-188, 2005.
6. ST LeDoux, WW Herling, J Fatta & RR Ratcliff. Multidisciplinary Design and Optimization System using Higher Order Analysis Code. 10th AIAA/ISSMO MAO Conference, *AIAA Paper 2004-4567*, Albany, NY, 2004.
7. VN Vatsa & DP Hammond. Viscous Flow Computations for Complex Geometries on Parallel Computers. 4th NASA Symposium on Large-Scale Analysis and Design on High-Performance Computers and Workstations, July, 1997.
8. A Jameson. Optimum Aerodynamic Design using Control Theory. *CFD Review*, 495-528, Wiley, 1995.
9. A Jameson. Aerodynamic Design via Control Theory. *Journal of Scientific Computing*, **3**, 233-260, 1988.
10. A Jameson & JC Vassberg. Computational Fluid Dynamics for Aerodynamic Design: Its Current and Future Impact. *AIAA Paper 2001-0538*, Reno, NV, 2001.
11. A Jameson, L Martinelli & JC Vassberg. Using Computational Fluid Dynamics for Aerodynamics – A Critical Assessment. *ICAS Paper 2002-1.10.1*, Toronto, 2002.
12. PG Buning, DC Jespersen, TH Pulliam, WM Chan, JP Slotnick, SE Krist, & KJ Renze. OVERFLOW User's Manual, Version 1.81. NASA Report, NASA Langley Research Center, Hampton, VA, 1999.
13. JC Vassberg, MA DeHaan & AJ Sclafani. Grid Generation Requirements for Accurate Drag Predictions based on OVERFLOW Calculations. 16th AIAA CFD Conference, *AIAA Paper 2003-4124*, Orlando, FL, June, 2003.
14. B Epstein, A Jameson, S Peigin, DL Roman, NA Harrison & JC Vassberg. Comparative Study of 3D Wing Drag Minimization by Different Optimization Techniques. AIAA Paper 2008-0326, Reno, NV, January, 2008.
15. JC Vassberg, MA DeHaan, SM Rivers & RA Wahls. Development of a Common Research Model for Applied CFD Validation Studies. AIAA Paper 2008-6919, Honolulu, HI, August, 2008.

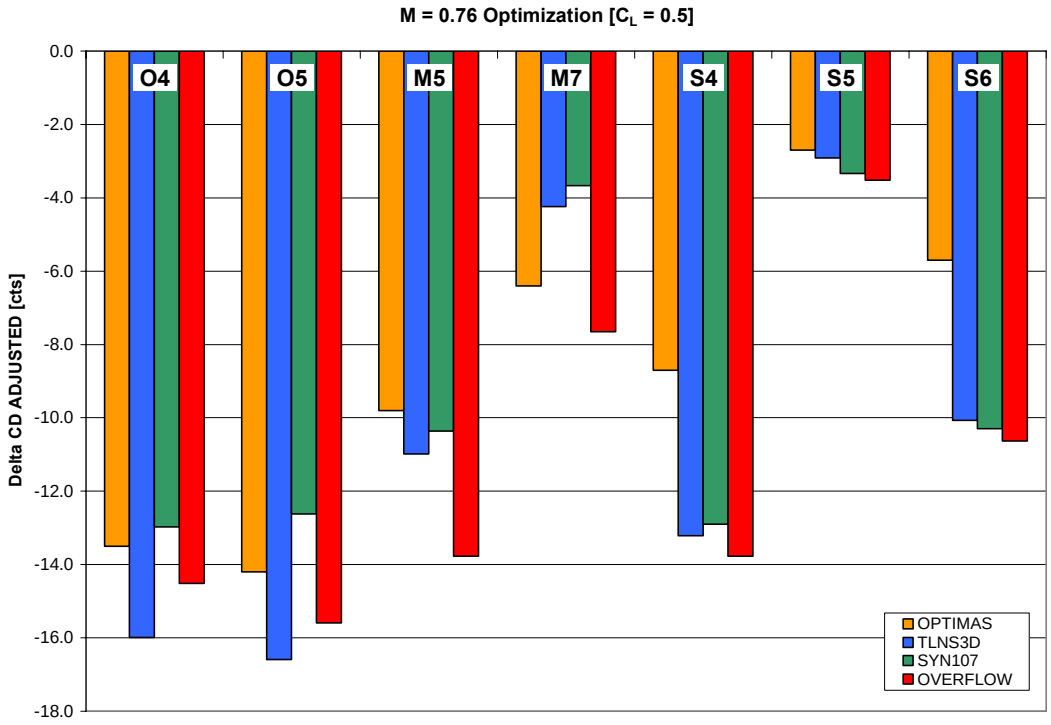

Figure 1: Adjusted Drag Reductions at M = 0.76, CL = 0.5, Re = 5 million.

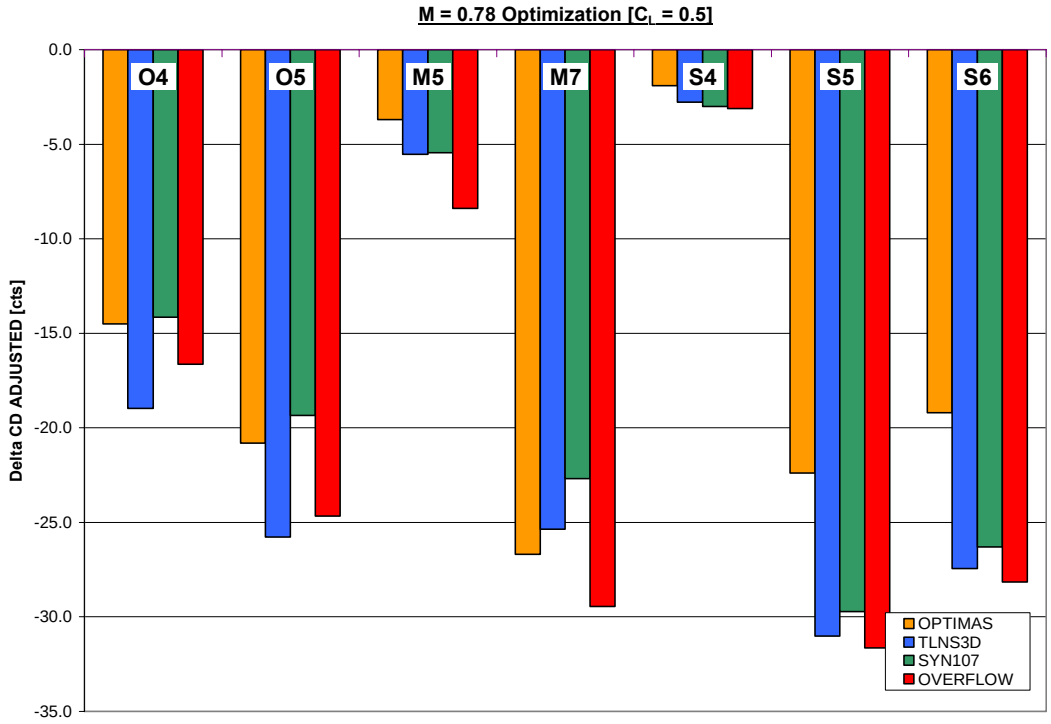

Figure 2: Adjusted Drag Reductions at M = 0.78, CL = 0.5, Re = 5 million.

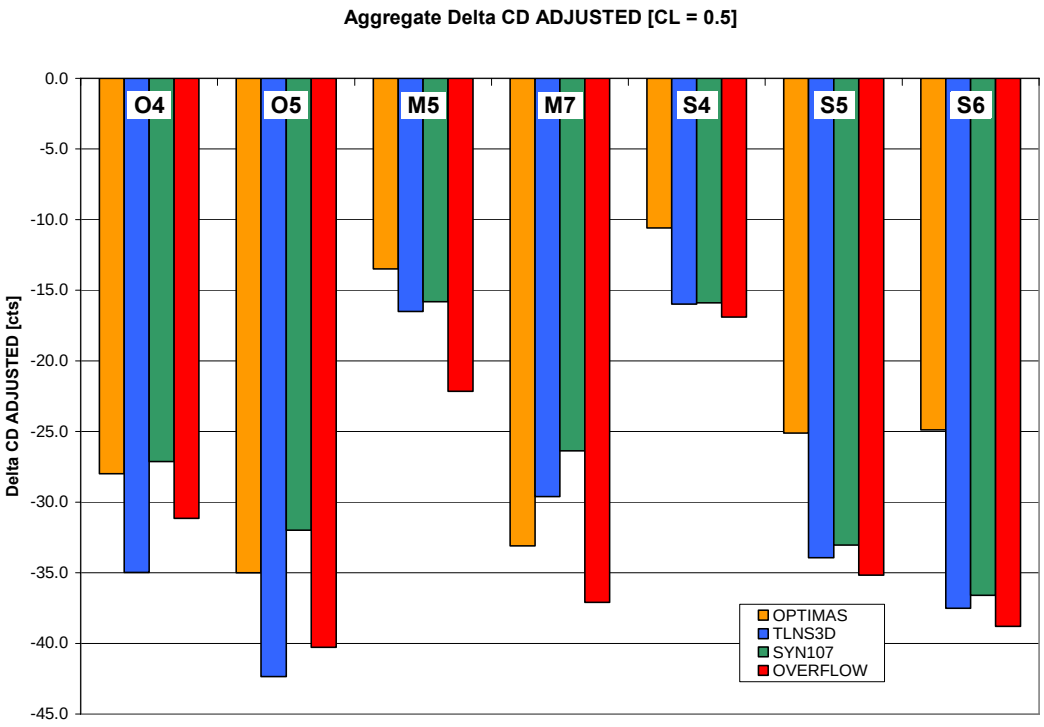

Figure 3: Aggregate Adjusted Drag Reductions at M = [0.76, 0.78], CL = 0.5.

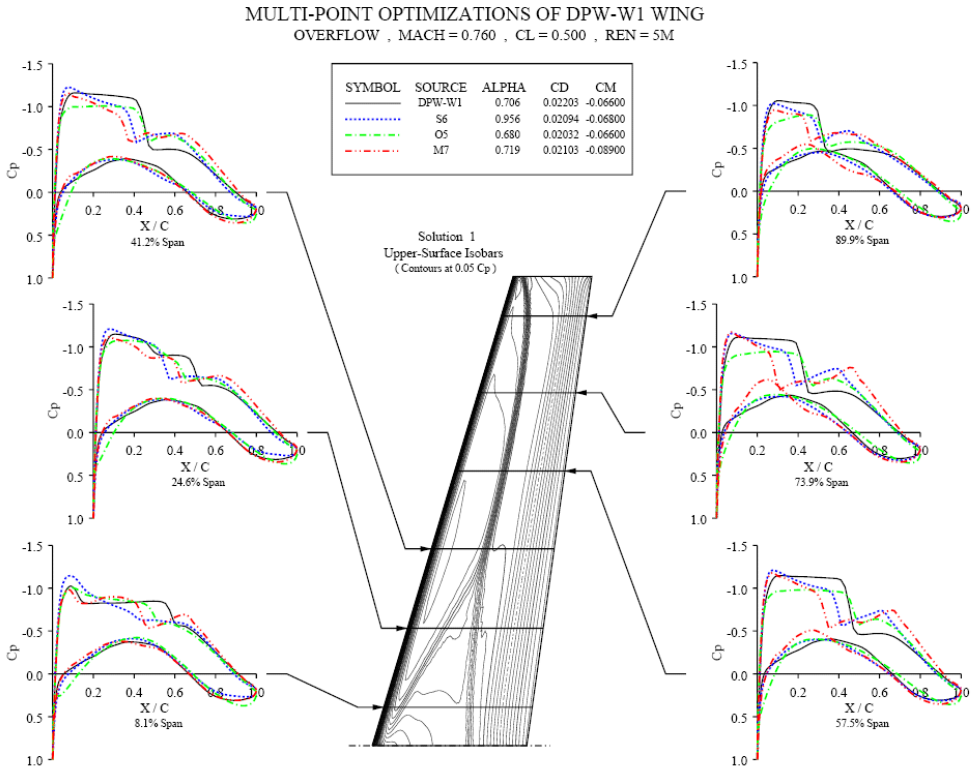

Figure 4: Comparison of Pressure Distributions at M = 0.76, CL = 0.5, Re = 5 million.

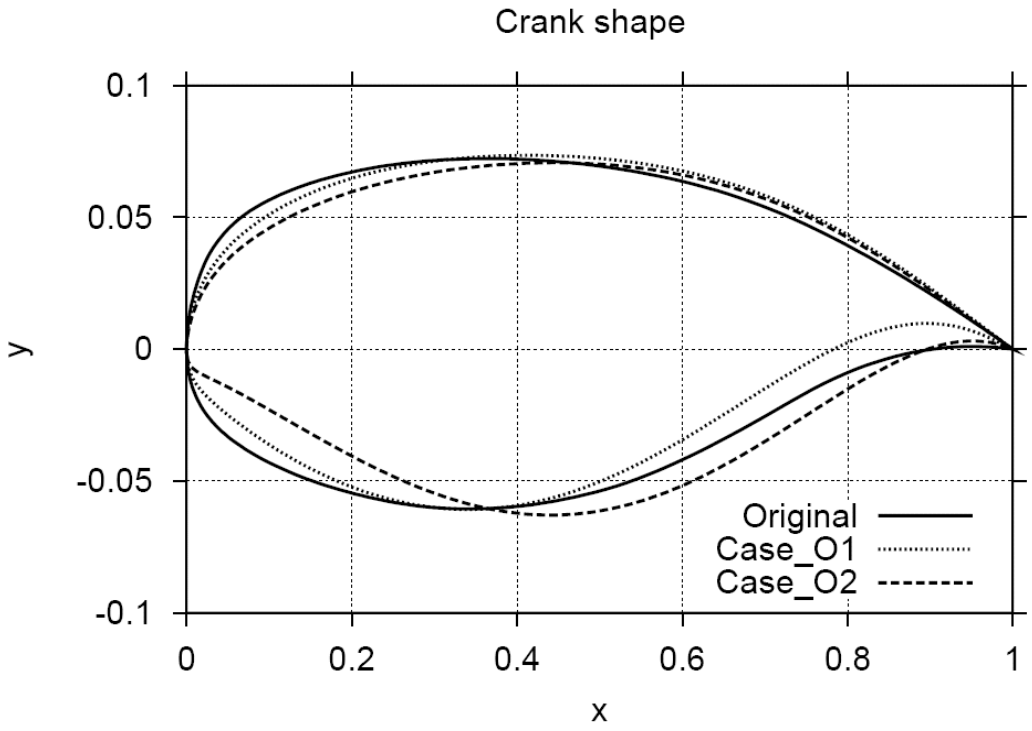

Figure 5: Comparison of Baseline, O1, and O2 Airfoil Sections at 37.5% Semispan.

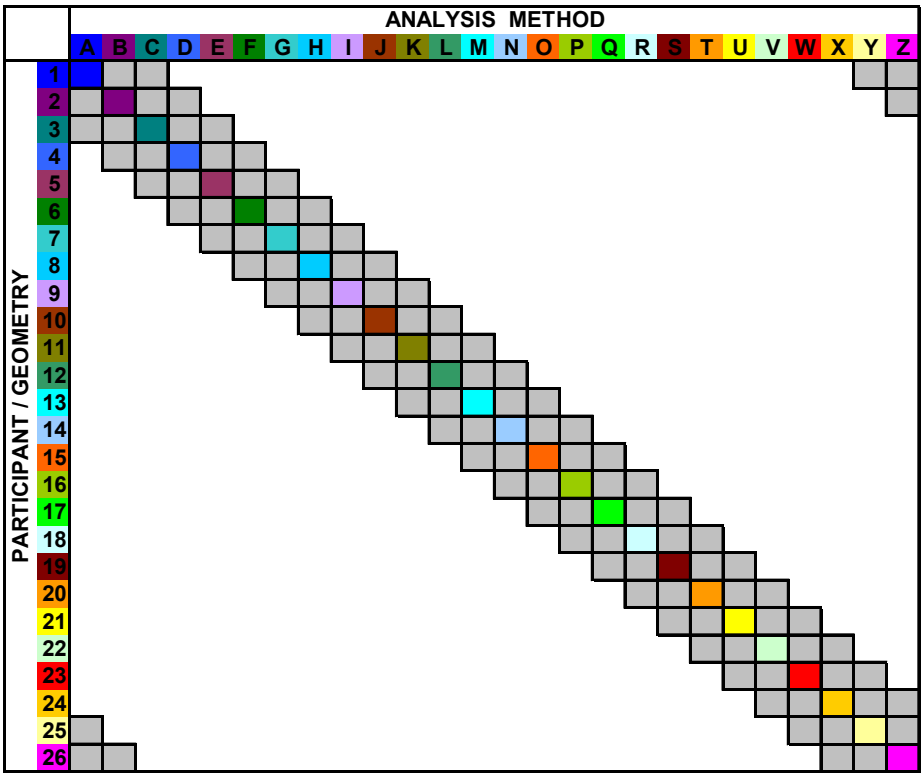

Figure 6: Template for Cross-Analysis Assignments for a Future Workshop.