Computers and Fluids 223 (2021) 104923

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

Test cases for inverse aerodynamic design

John C. Vassberg[a,1,*], Antony Jameson[b,2]

[a] Advanced Concepts Design Center, Boeing Commercial Airplanes, Long Beach, CA 90808, USA
[b] Department of Aerospace Engineering, Texas A & M University, College Station, TX 77843, USA

ARTICLE INFO

Article history:
Received 15 July 2020
Revised 22 February 2021
Accepted 10 March 2021
Available online 29 March 2021

Keywords:
Aerodynamics
Inverse Design
Optimization
CFD
Airfoils
Pressure distribution

ABSTRACT

What do you get a distinguished Englishman (who has everything) for his birthday? Another publication, but of course! The dilemma is in how to keep it a surprise, while developing a document which warrants his authorship. No problem, Antony and I have collaborated on several subjects over the years which we have yet to publish. The topic chosen for this gift is: Inverse Aerodynamic Design. The examples provided herein are from a subset of test cases we have utilized over the past three decades to ensure our methods are robust, accurate, and cost efficient. This paper takes a bird's eye view of the forest, so as to not get lost in the details of the trees. It is written in a more casual style. It is intended to provoke thought & spark discussion throughout the aerodynamic community, especially with regards to inverse design. It is loosely organized as follows. Any mis-statements, inaccuracies, controversial assertions, or praises bestowed upon the man-of-honor are the sole responsibility of the first author, while anything of any value can be attributed to the second. Antony, Happy 85th.

Introduction

Depending on your definition of computational fluid dynamics, CFD has helped designers shape airplanes over the past 5 or 6 decades, and analytical methods reach back more than a century. Nikolai Joukowsky published his historic conformal transformation [1] in 1910. This remarkably simple method continues to provide invaluable insight into the aerodynamics of airfoils and wings. Unfortunately, this sacred, fundamental building block of aerodynamic understanding is beginning to be dropped from the curriculum of aerodynamics classes across the nation, being replaced with more time spent on CFD tutorials. This shift in our educational system is flawed, for it is far better to understand the basics than it is to not. When things go wrong, and they will, designers need a solid aerodynamic foundation on which to fall. The beauty of Joukowsky is that with only two design variables (circle radius and offset), one can truly understand the effects of thickness and camber, as well as the details of pressure distributions, over a fairly wide set of airfoil shapes. Furthermore, when it comes to aerodynamic lift, it teaches us that *the tail wags the dog,* meaning, circulation is set by the trailing edge. Joukowsky is limited to cusp airfoils, and incompressible, inviscid flows, but the wealth of knowledge gained is well worth the minimal amount of effort required. And yes, there are many problems under study today (most in fact) that require numerical methods to solve. This is not an either-or proposition; the best designers continue to build their foundation with a diverse set of tools at their disposal, with inverse design being an important one.

We believe that the Art of Inverse Design is slowly fading from practice. In its place, the literature on aerodynamic shape optimizations suggests that the focus today is almost entirely on Drag Minimization exercises. We are concerned that inverse design could get lost to future generations of aerodynamicists. The intent of this publication is to help reinvigorate this art form.

Inverse Aerodynamic Design

For as long as the authors have been in this business, inverse aerodynamic design has played a crucial role in the development and refinement of aircraft, and can be traced back to 1945 when Sir James Lighthill [2] published his method based on conformal transformations. In this context, Inverse Design refers to the design of an airfoil, wing, or other aerodynamic shape by the specification of a desired pressure (or velocity) distribution about a geometry. However as stated, this is an ill-posed problem, as not all pressure distributions are achievable. Hence, the more formal definition of inverse design typically used is to minimize the integration of the squared differences between the target pressures and the achiev-

* Corresponding author.
E-mail address: john.c.vassberg@gmail.com (J.C. Vassberg).

[1] Boeing Technical Fellow, Fellow AIAA
[2] TEES Eminent Professor, Honorary Fellow AIAA

https://doi.org/10.1016/j.compfluid.2021.104923
0045-7930/

able set.

$$I = \int_S [C_P - C_P^*]^2 \, dS \tag{1}$$

Here, I is the cost function to minimize, S is the region over-which the target pressure distribution C_P^* is specified, and C_P is a distribution from the achievable set. This problem can be solved numerically by shape optimization techniques based on control theory [3–6].

An inverse-design optimization can be subjected to a variety of constraints. For example, one normally requires the airfoil contour to be a simple closed loop, but one can also require the final geometry to meet various thickness constraints. In addition, the optimization can be conducted at a fixed angle-of-attack (α) or at fixed lifting (C_l) conditions. A final note is that the region S (over-which the target pressures are specified) and the geometry to be designed, do not have to be one and the same. For example, assume one wants to design the floor and ceiling shapes of a two-dimensional wind tunnel such that an airfoil tested in this environment yields the same pressure distribution as it produces in free air. Here, one would run the airfoil in free air at the flow conditions of interest, and then use the resulting pressure distribution as the target for the inverse design of the wind-tunnel walls. In this example, the geometry designed is mutually exclusive of the region where the target pressures are sought.

Test Cases

The test cases provided herein follow more closely the classical inverse-design problem where the target pressures are specified (essentially) over the whole airfoil, and with the whole airfoil designed, subject to closed-contour shapes. There is nothing spectacularly important about the cases provided, and as such, the reader can instead insert his or her own favorite set of airfoils. And while all of the examples shown are two-dimensional Euler inverse designs, the authors also use three-dimensional and/or viscous analogs of these cases to test our full suite of aerodynamic shape optimization methods. (*e.g.,* layout a 3D wing based on the NACA0012 airfoil and the ONERA-M6 wing planform, and use the M6 wing pressure distributions as the target.) For the sake of consistency and simplicity, all of the given test cases begin with a sharp trailing-edge NACA0012 airfoil, are run at fixed lift, and require the designed shape to be a closed contour. The cases are generally ordered from simplest to the more difficult. They represent much larger departures between the initial state and target pressures than are typically encountered in practice. In all of our optimizations, we utilize a free-surface (or parameter free) design space, where every discrete point of the grid defining the airfoil is allowed to move independently of each other. Our standard C-mesh grid has dimensions of (192x32) cells, with 128 cells defining the airfoil. We also test at double and quadruple these dimensions, (384x64) and (768x128) cells, respectively. On our standard mesh, SYN83 completes about 4–5 design cycles per second on a single core of a MacBookPro laptop running Mojave 10.14.6 on a 2.8 GHz Intel i7 core, and the code compiled with gfortran.

ONERA-M6 Inverse Design

As a warm-up exercise, we begin with our simplest test case of designing the symmetric ONERA-M6 airfoil, starting with the symmetric NACA0012 shape. The slight complication we introduce is to choose a lifting condition, and hence, the target pressures are clearly not symmetric. The first step is to manufacture a realizable target pressure distribution by running the M6 airfoil at $M = 0.84$, and $\alpha = 1.25^\circ$. [Use your CFD method of choice, our illustrations are based on SYN83.] In our case, the resulting lift coefficient is $C_l = 0.479168$, and we impose this value to constrain the lift during the inverse-design optimization. In order to further simplify this test case, we do not impose any constraints on thickness, and we constrain the design space to the set of symmetric shapes. The left hand side of Fig. 1 depicts the manufactured pressures which are then used as the target distribution. The right hand side of this figure provides Mach contours in the near field. Fig. 2 illustrates the initial state of the solution about the NACA0012 at $M = 0.84$, and $C_l = 0.479$. Note the fairly-strong shock that exists on the lower surface, whereas the target pressures of the M6 has none. The pressures near the leading edge of the NACA0012 exhibit a monotonic rise into its shocks, whereas the M6 pressures develop a local peak on both surfaces. Fig. 3 shows the progress made after the first design cycle. In this figure, the dashed lines represent the initial state, the open circles are the target pressures, the plus (+) signs are the current upper-surface pressures, and the times ($\times$) signs are the current lower-surface pressures. The porcupine quills emanating from the airfoil geometry depict the current gradient of the optimization after being projected into a Sobolev space. Note that the changes made on the first design cycle are significant and seem reasonable. Figs. 4,5,6 illustrate the states after 10, 50, and 100 design cycles, respectively. For all practical purposes, this inverse design is done by 50 design cycles, where the $(+, \times)$ symbols fall inside the open circles. However, pay close attention to α; at 50 design cycles it is 1.25541°, while at 100 design cycles it is 1.25073°, thus continuing to migrate towards the known correct value of 1.25°. This test case only takes 24 seconds to complete 100 design cycles, but it is the simplest one.

RAE2822 Inverse Design

Our next, slightly harder test case is an inverse design of the RAE2822 airfoil. The added complication here, relative to the first test case, is that the RAE2822 is a non-symmetric airfoil. As with the M6 design, we manufacture a realizable target pressure distribution by analyzing the RAE2822 at its design point of $M = 0.75$, and $C_l = 0.6$. The optimization process is similar to that of the M6 case except that the design space is opened up to allow non-symmetric airfoil shapes. Again, no thickness constraints are applied. Fig. 7 illustrates the manufactured target pressures. Fig. 8 gives the initial state about the NACA0012 at $M = 0.75$, and $C_l = 0.6$, and

Figs. 9,10,11,12 provide the states after 1, 10, 100, and 1000 design cycles, respectively. Noticable progress is made after 10 design cycles, and substantial progress is recognized by 100 design cycles. This optimization is essentially done after 1000 design cycles, however even then, minor discrepancies persist near the upper-surface trailing-edge region. Another indication that the process is not completely converged is that α is still about 0.03° off the known correct value. This test case takes 232 seconds to run 1000 design cycles.

An interesting variation of this test case is to run it in reverse, using the RAE2822 as the seed airfoil, and the NACA0012 pressure distribution as the target. Here, we start with a non-symmetric shape, yet we know that the final designed shape should be symmetrical; this provides another metric to monitor during the convergence of the optimization process.

Crazy Inverse Design

Now is when the fun begins, and this is clearly a case of aerodynamicists gone wild. (It's all relative folks.) In this test case, we actually try to break our methods by crafting a very unrealizable target pressure distribution, just to see what happens, and yet we expect the worse. The set-up is simple. With $M = 0.8$, and $C_l = 0.5$, the upper-surface target is defined with a flat roof-top of $[C_P = -1.0]$ for $[0.025 \le X \le 0.5]$, and a flat ambient level of $[C_P = 0.0]$

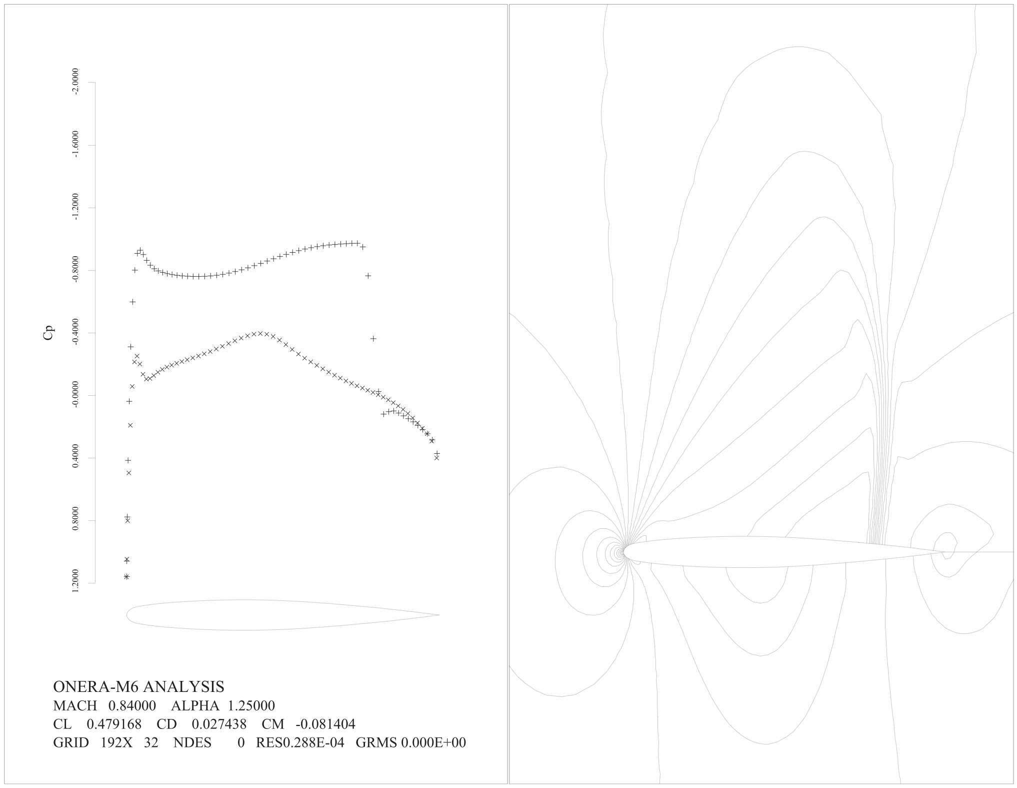

Fig. 1. ONERA-M6 airfoil pressure distribution, $M = 0.84$, $\alpha = 1.25^{\circ}$.

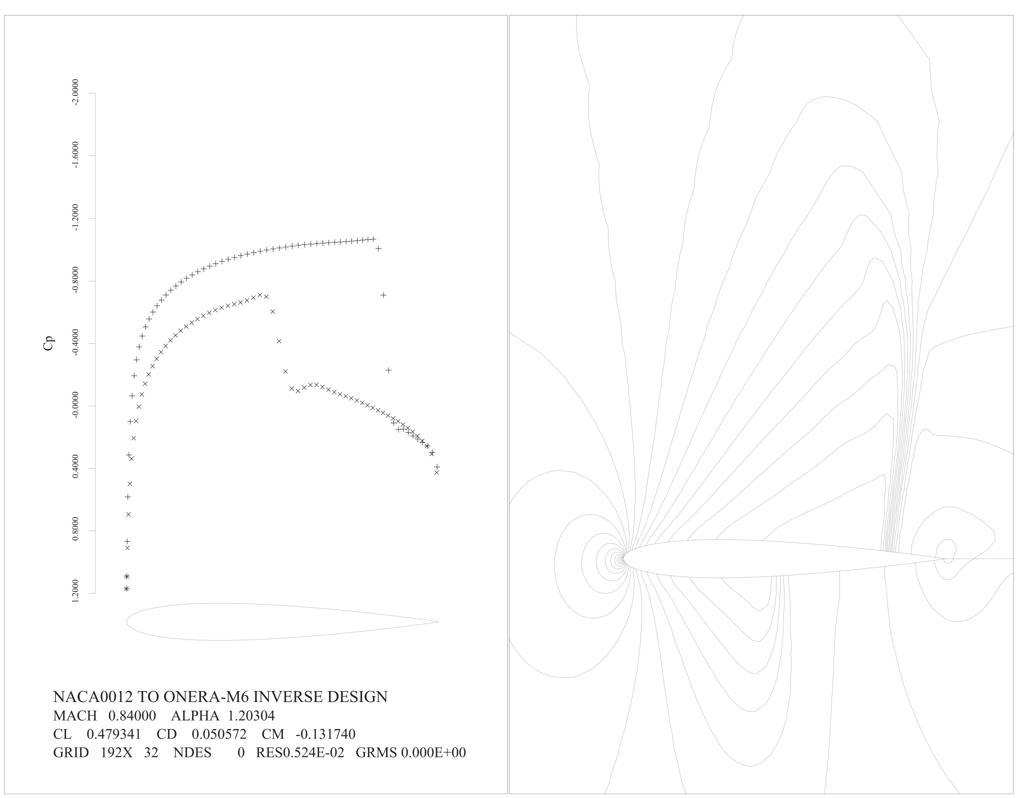

Fig. 2. NACA0012 to ONERA-M6 inverse design at cycle 0.

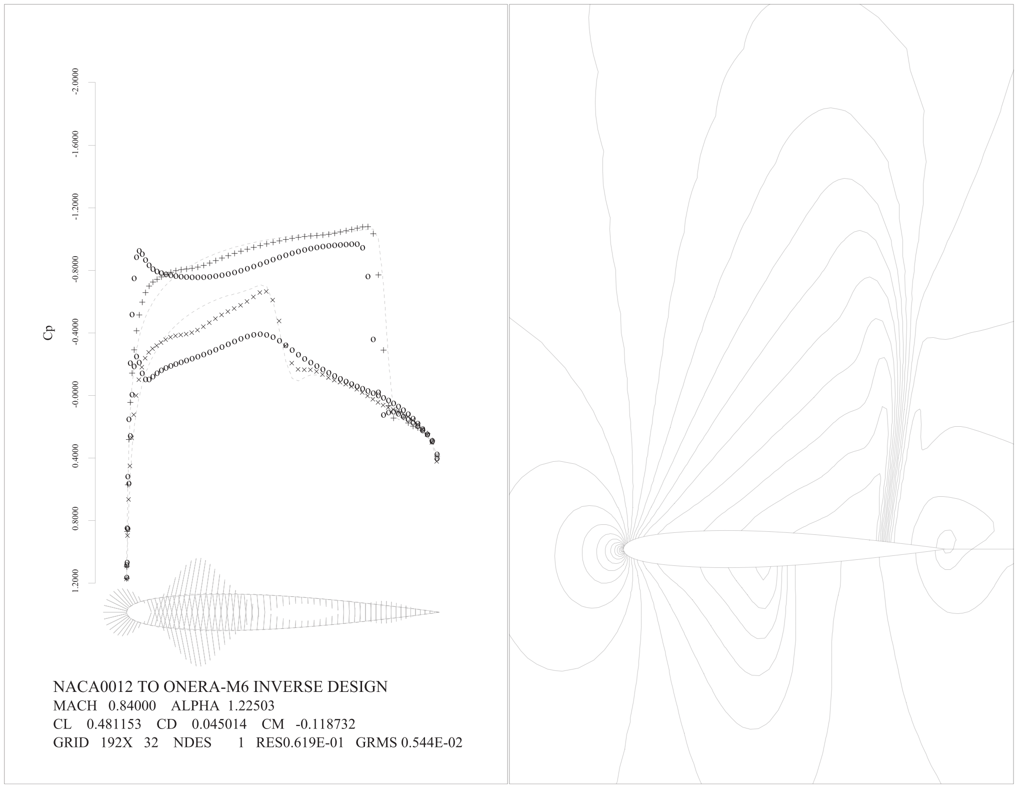

Fig. 3. NACA0012 to ONERA-M6 inverse design at cycle 1.

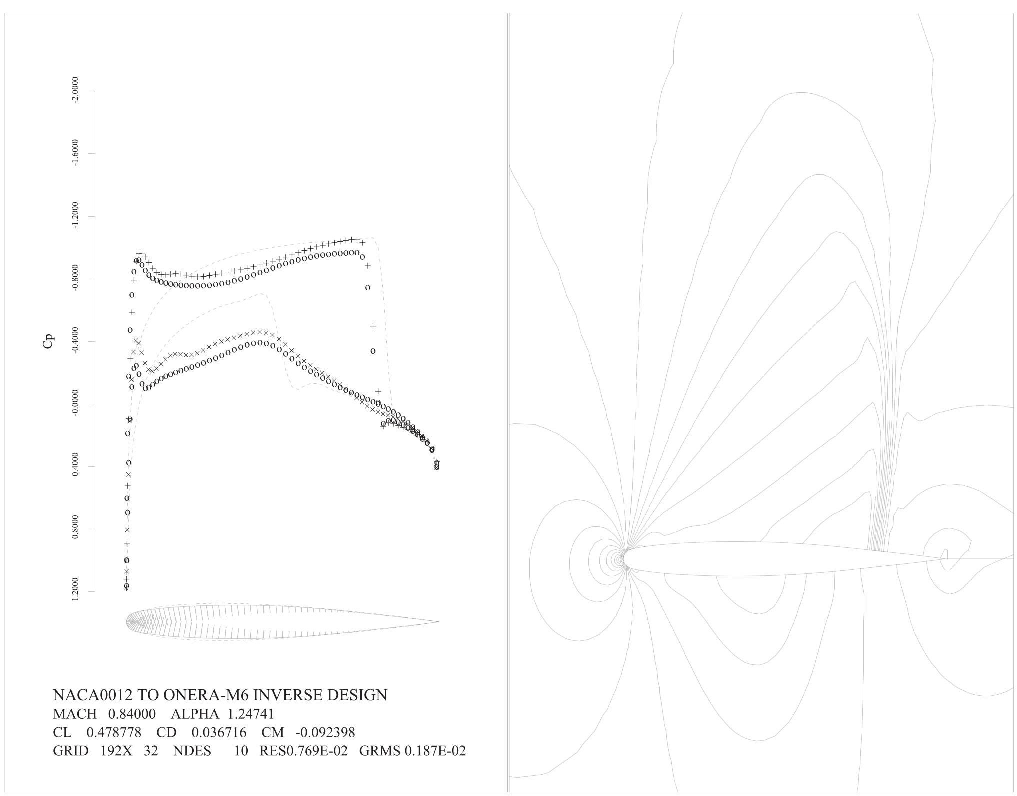

Fig. 4. NACA0012 to ONERA-M6 inverse design at cycle 10.

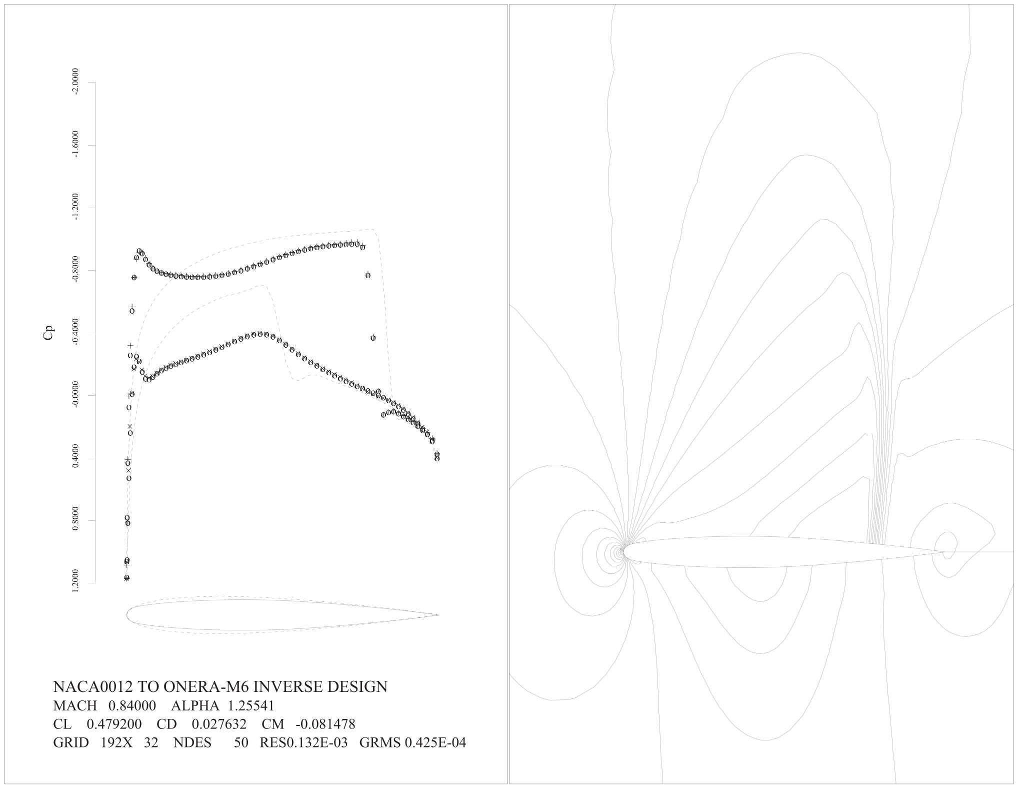

Fig. 5. NACA0012 to ONERA-M6 inverse design at cycle 50.

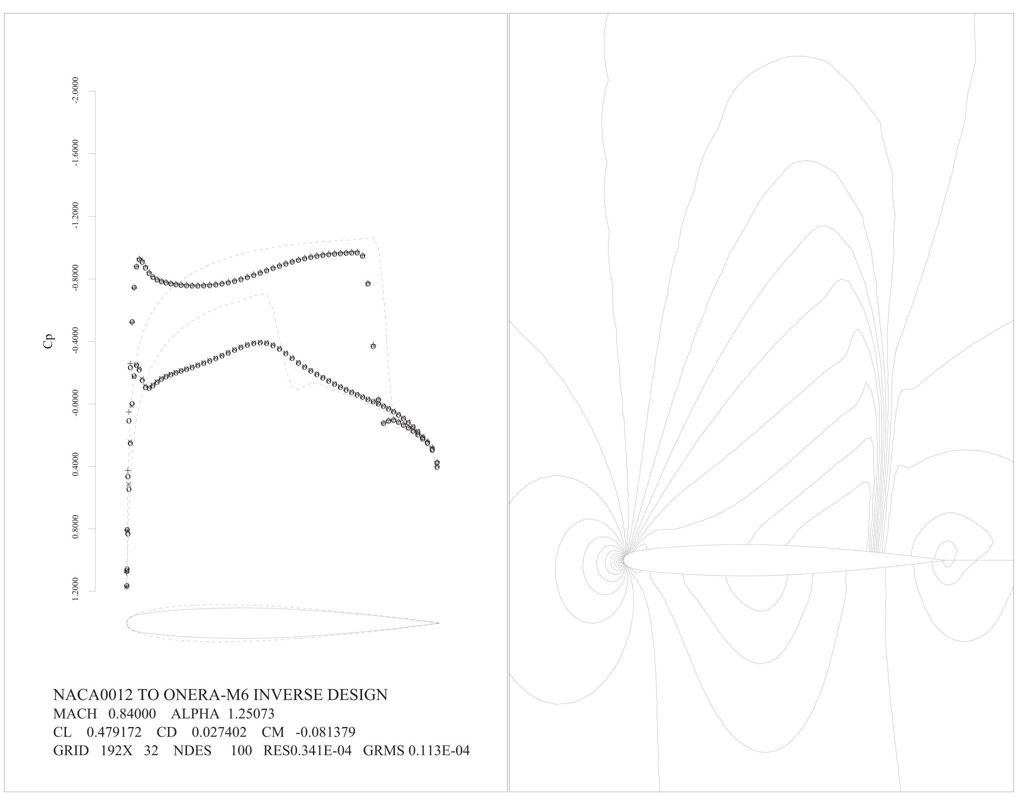

Fig. 6. NACA0012 to ONERA-M6 inverse design at cycle 100.

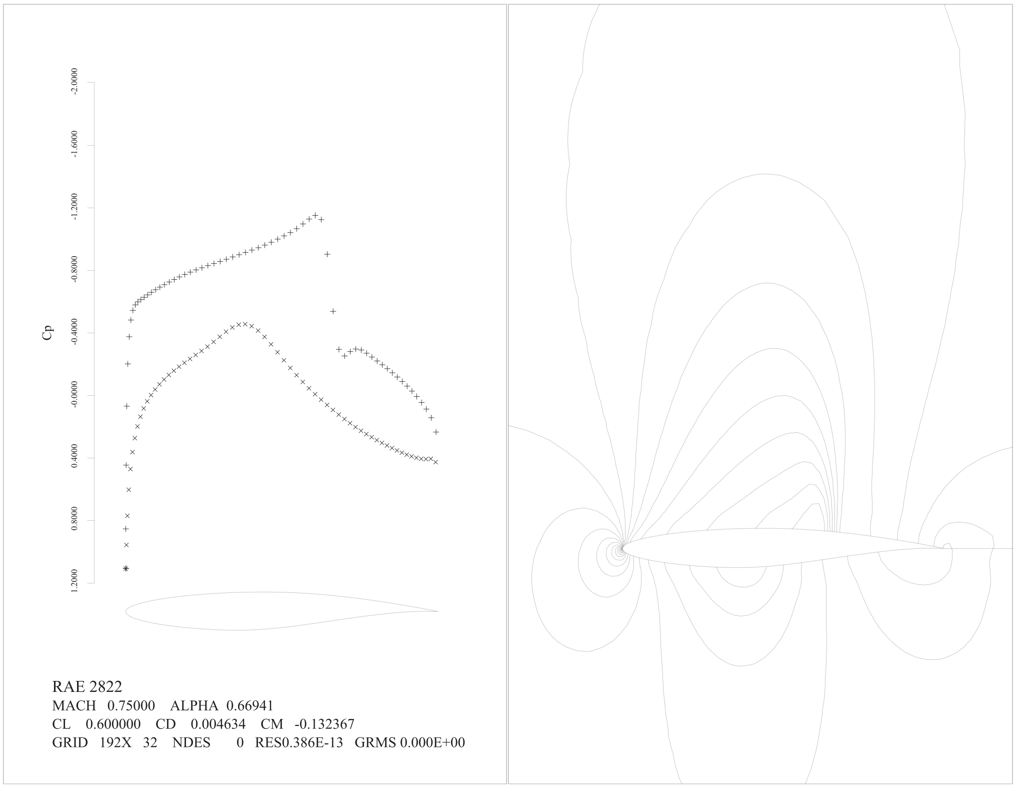

Fig. 7. RAE2822 airfoil pressure distribution, $M = 0.75$, $C_l = 0.6$.

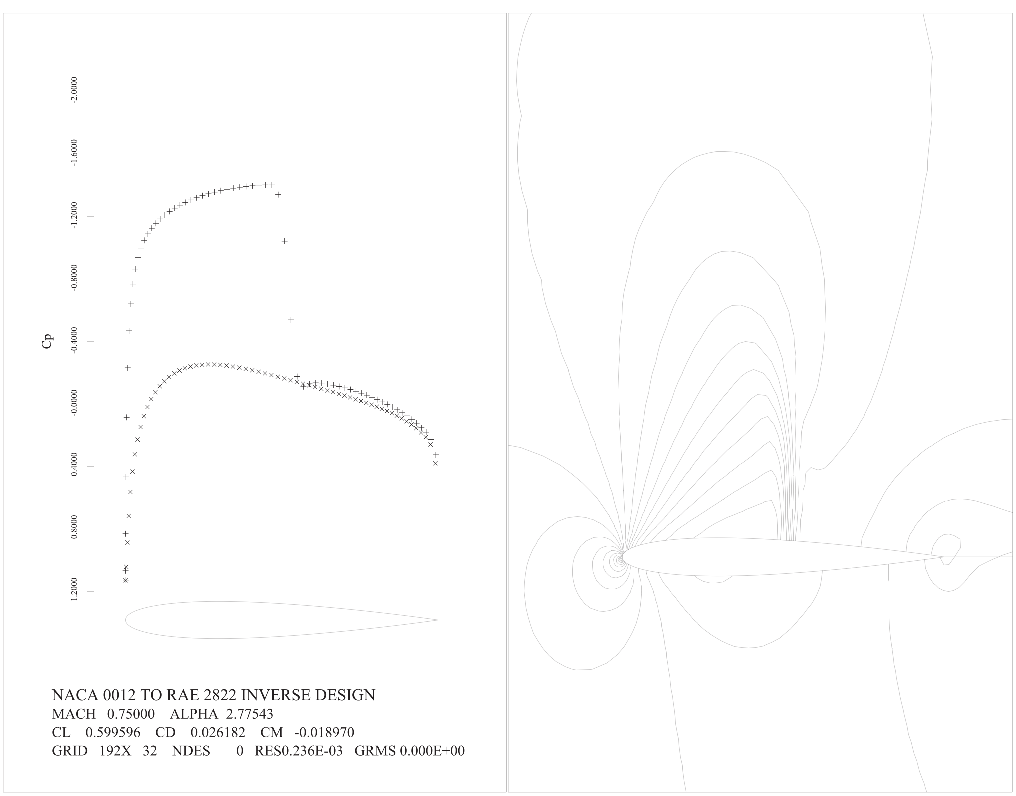

Fig. 8. NACA0012 to RAE2822 inverse design at cycle 0.

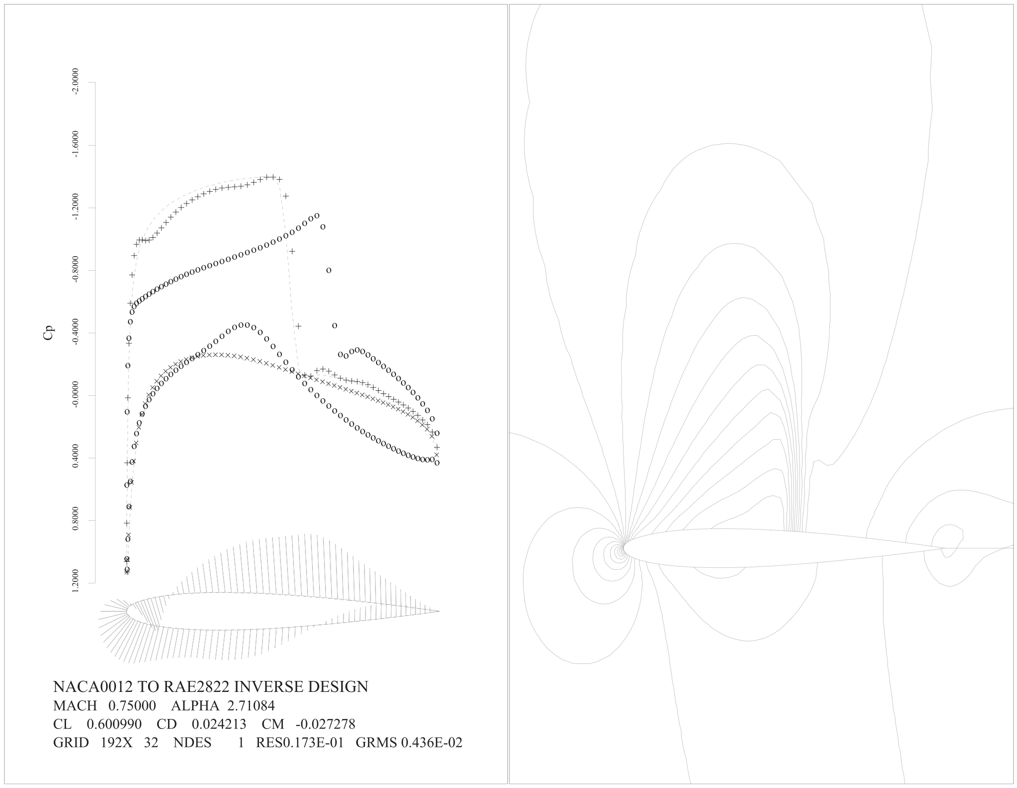

Fig. 9. NACA0012 to RAE2822 inverse design at cycle 1.

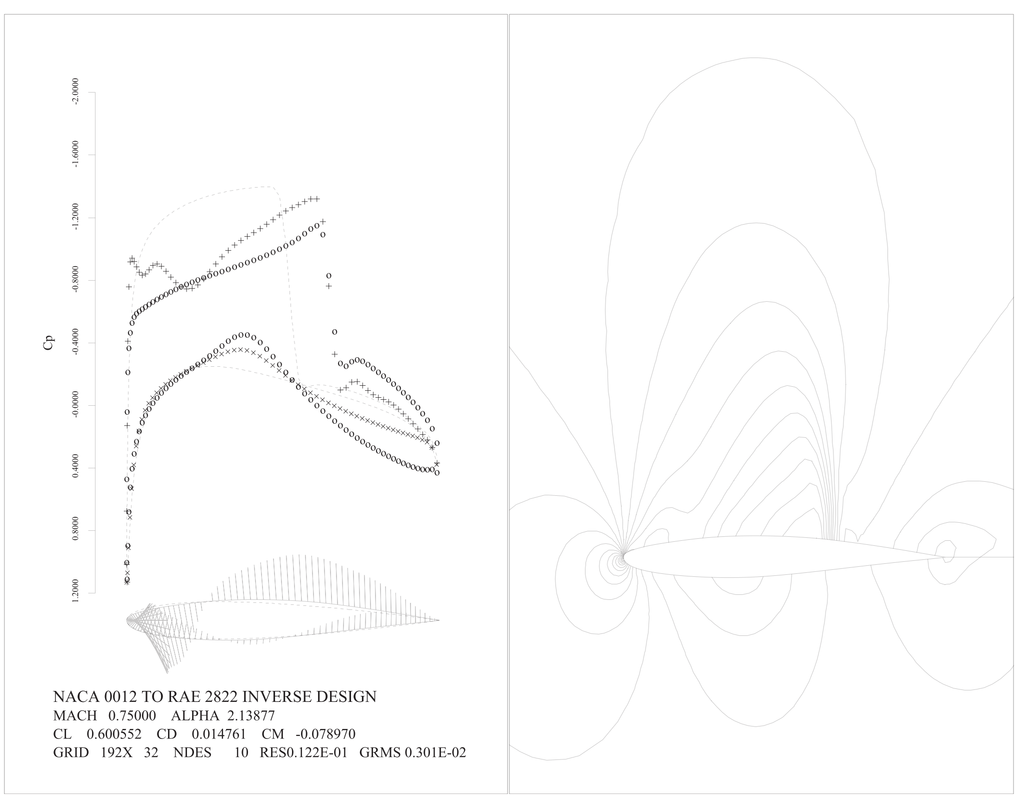

Fig. 10. NACA0012 to RAE2822 inverse design at cycle 10.

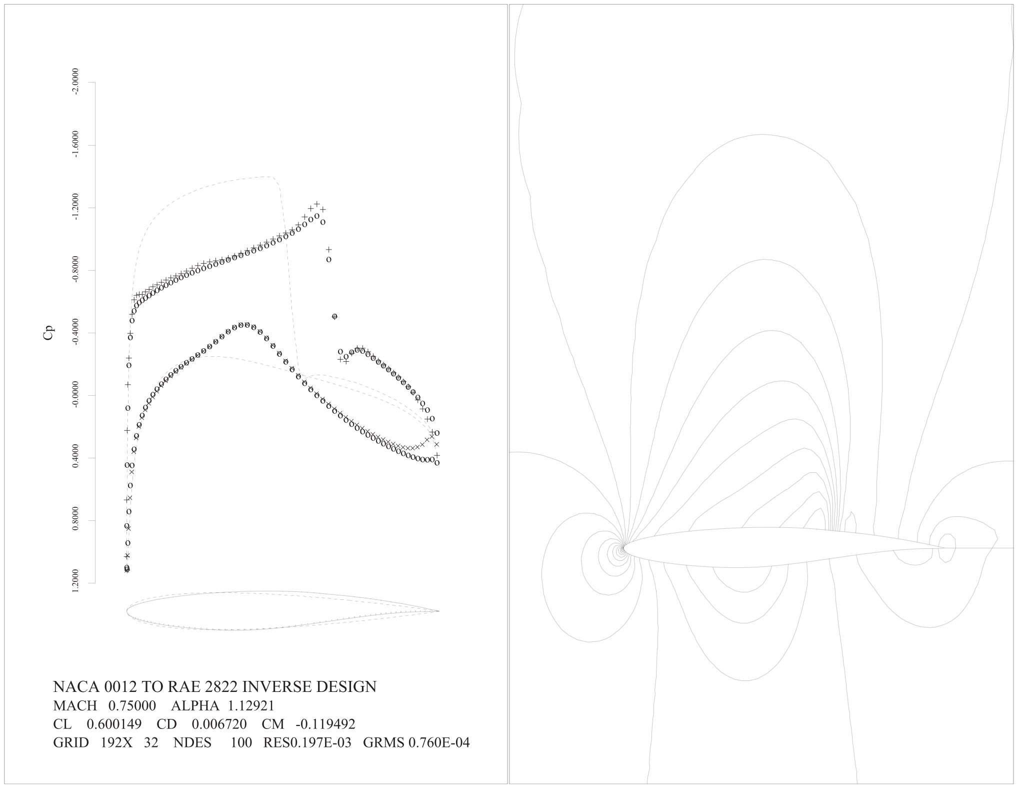

Fig. 11. NACA0012 to RAE2822 inverse design at cycle 100.

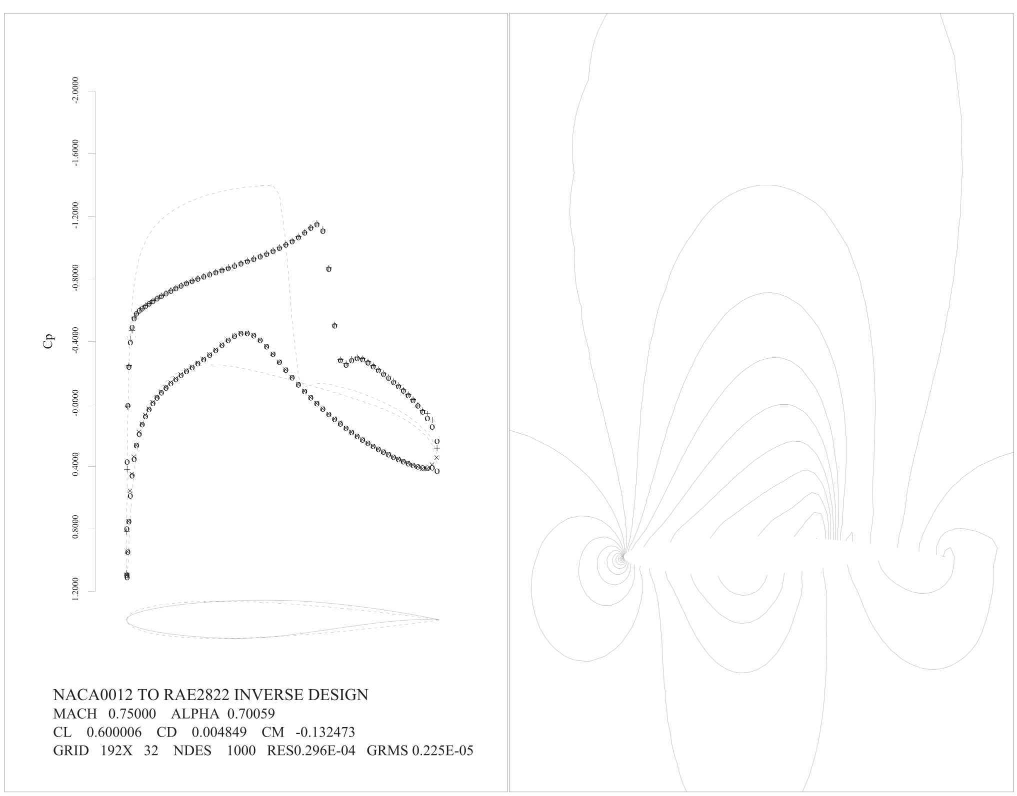

Fig. 12. NACA0012 to RAE2822 inverse design at cycle 1000.

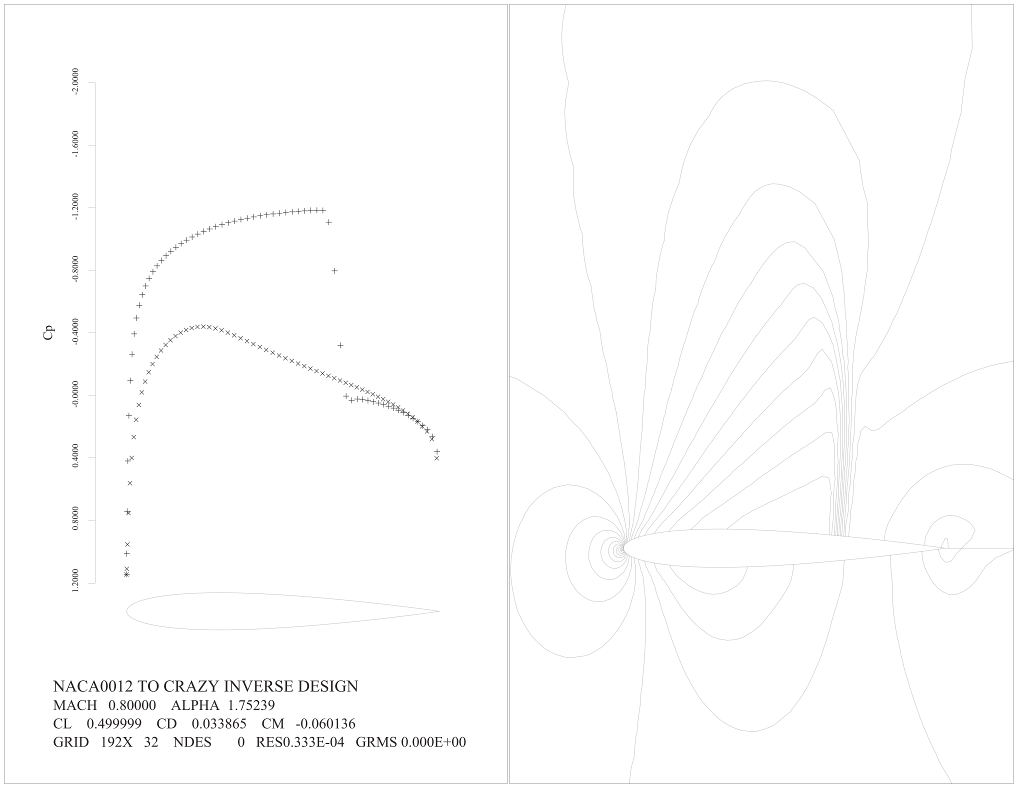

Fig. 13. NACA0012 airfoil pressure distribution, $M = 0.8$, $C_l = 0.5$.

for $[0.5 < X \leq 1.0]$. The specified target on the lower surface is a flat ambient level of $[C_P = 0.0]$ for $[0.025 \leq X \leq 1.0]$. Notice that we allow the first 2.5%-chord float to pressures that come naturally about the leading-edge stagnation region. Figs. 13,14,15,16,17,18,19 provide the convergence history of this inverse design at 0, 1, 10, 50, 100, 500, and 1000 design cycles, respectively. Notice that the porcupine quills are still somewhat discernible after 100 design cycles (requires close inspection), but beyond that the design is essentially done. Although we used Rankine-Hugoniot to guide us in setting-up an impossible test case, the optimization process outsmarted us by designing a tadpole-shaped airfoil which produces a curved-oblique shock at the prescribed mid-chord location. And although there is an over-shoot of the pressures approaching the shock, it is amazing how well the inverse design succeeded everywhere else. The ambient levels on both upper and lower surfaces are closely recovered with only an excursion very near the trailing edge, and the roof-top is almost captured true-to-form over the interval of $[0.025 \leq X \leq 0.35]$. All in all, this is a surprising outcome. For completeness, we also include the results of this test case, as run on our next finer mesh of (384x64) cell dimensions, in Fig. 20. Here the story remains much the same, just crisper, and with the region of over-shoot reduced by about a factor of two. With this finding, we propose the following question to the curious reader. If one continues to refine the mesh and enrich the design space, will the tadpole shape develop an aft-facing step at mid-chord, or will it remain a smooth ramp?

Liebeck-esque Inverse Design

Our next and final test case pays homage to our friend and colleague Bob Liebeck. In a nutshell, Liebeck's pioneering work from the1960′s was to design a class of airfoils for incompressible flows that maximized lift while just keeping the flow attached. He accomplished this by specifying a flat roof-top pressure on the forward portion of the upper surface, followed by a Stratford pressure recovery to the trailing edge. Bob used an inverse-design method developed by Malcolm James, also of the Douglas Aircraft Company. This research required trials-and-errors in specifying the target pressure distributions, yet yielded a lasting contribution to the development of very-high lift-to-drag-ratio airfoil designs. In keeping with the spirit of this approach, we define a roof-top pressure on the upper surface with $[C_P = -2.5]$ from $[0.0125 \leq X \leq 0.3517]$, followed by a Stratford-like recovery of $[C_P(X) = -15(1-X)^4 + 0.15]$ over the remainder of the upper surface. On the lower surface, we include a sharp rise from near the stagnation point defined by $[C_P(X) = 1.0135 - 24.5X]$ over the range of $[0.0020 \leq X \leq 0.0252]$, followed by a shallow rise to the trailing edge of $[C_P(X) = 0.4 - 0.25X]$, for $[0.0252 \leq X \leq 1.0]$. There is no reason to believe that this target is an achievable pressure distribution. This inverse design is conducted at $M = 0.2$, and at a fixed lift of $C_l = 1.315$. Again, we begin with the NACA0012 airfoil at this flow condition, as shown in Fig. 21. Here, the angle-of-attack is slightly over 11.3°. Notice that the C_P-peak of the NACA0012 at this lifting condition exceeds our specified roof-top level, but only over a small range. Fig. 22 illustrates the state after the first design cycle. It is quite obvious that the starting pressure distribution of the NACA0012 is significantly different than our specified target-pressure architecture. Consequently, one should also expect that the designed airfoil will be significantly different than the NACA0012. A quick review of the gradient (porcupine quills) in this figure reveals that the designed airfoil will be significantly thicker than the baseline, and will have large positive camber over the forward portion of the chord. Figs. 23,24,25,26,27,28,29,30 complete the snapshots of this inverse design's convergence history at design cycles of 5, 10, 50, 100, 500, 1000, 1500, and 2000, respectively. Note that this low-Mach test case exhibits a slower convergence than that of the transonic test cases. Also, while the target pressures are matched

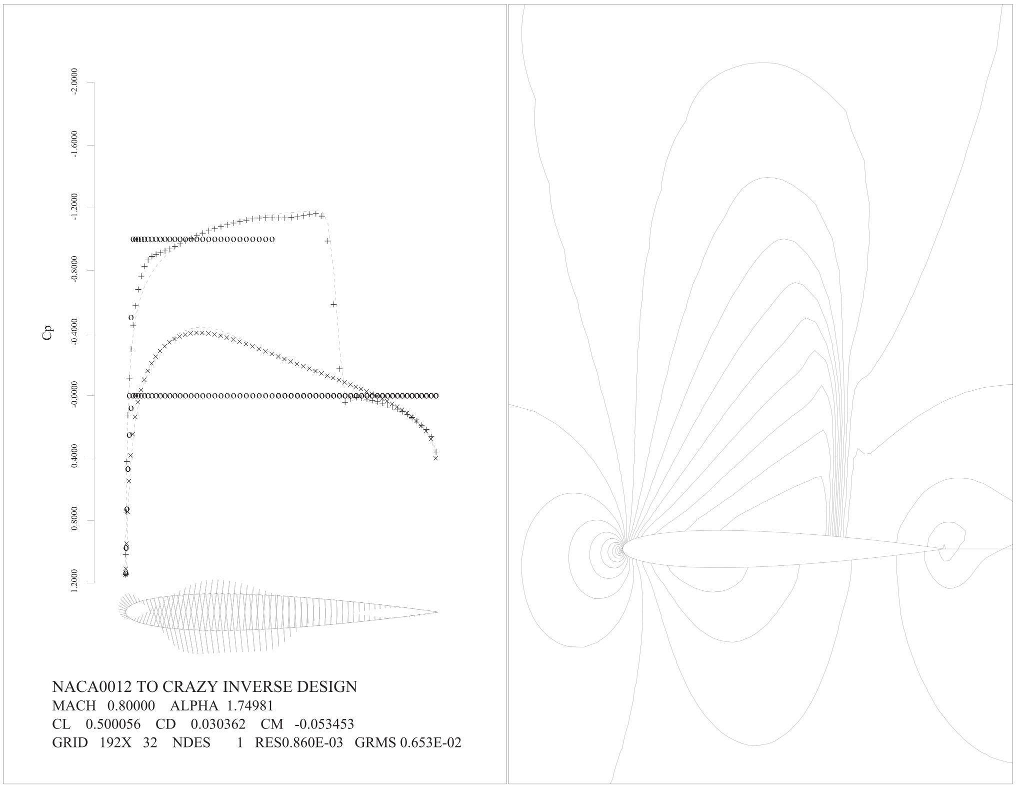

Fig. 14. NACA0012 to crazy inverse design at cycle 1.

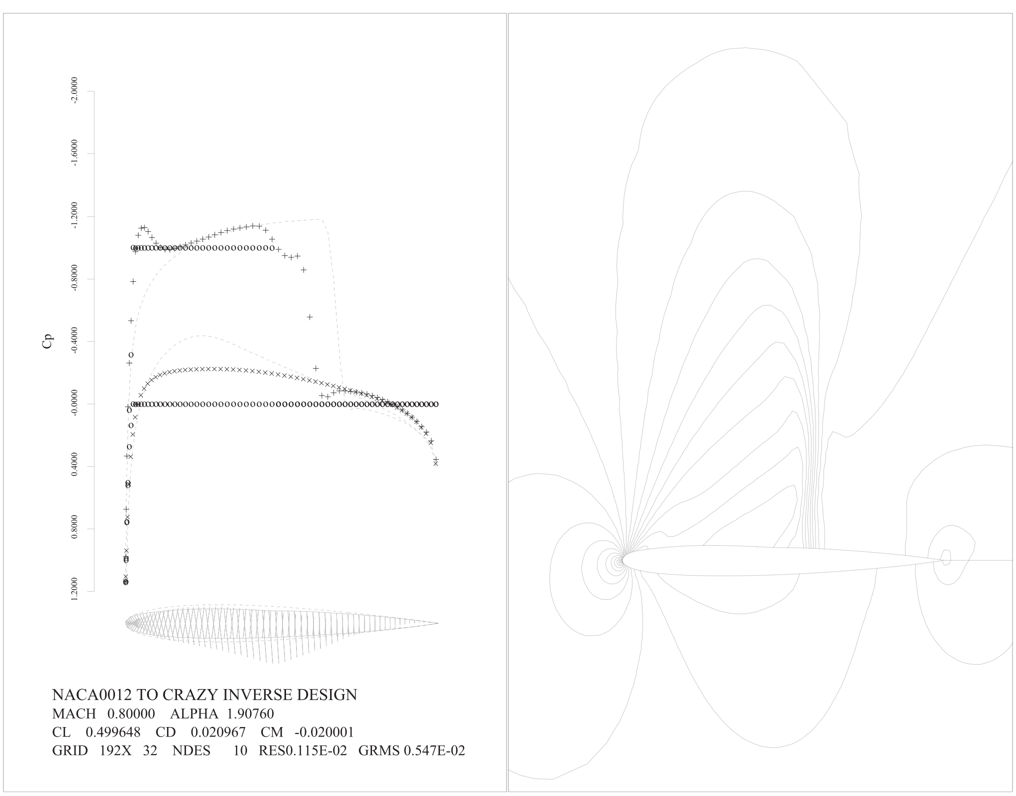

Fig. 15. NACA0012 to crazy inverse design at cycle 10.

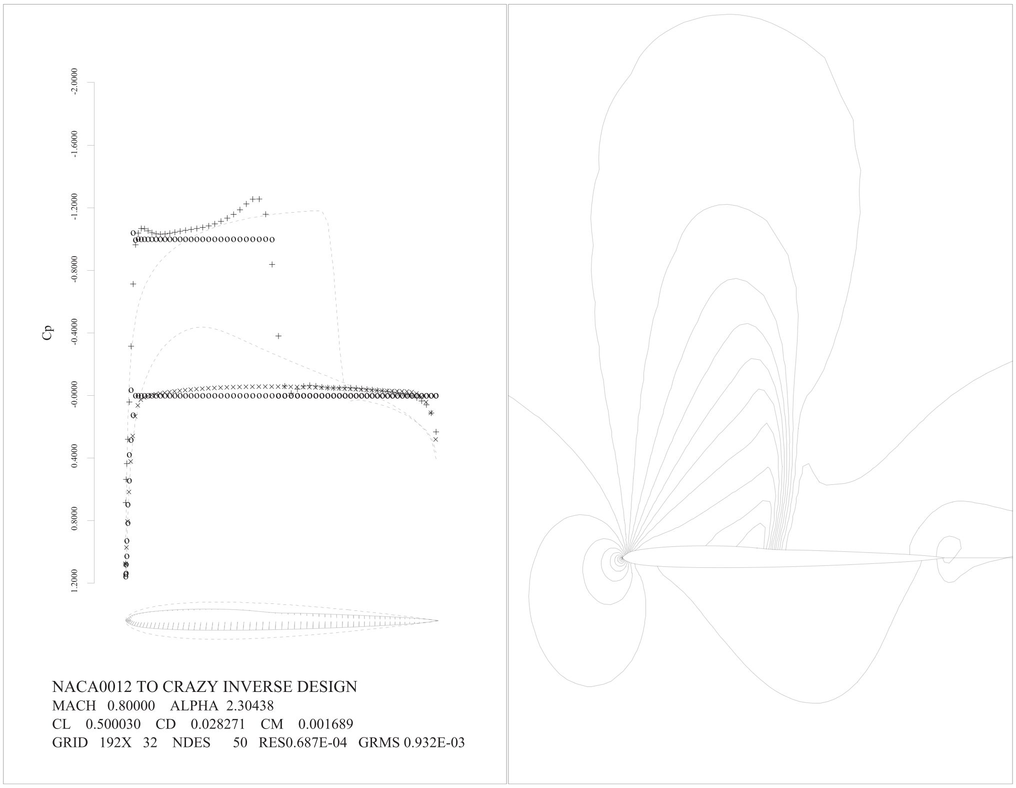

Fig. 16. NACA0012 to crazy inverse design at cycle 50.

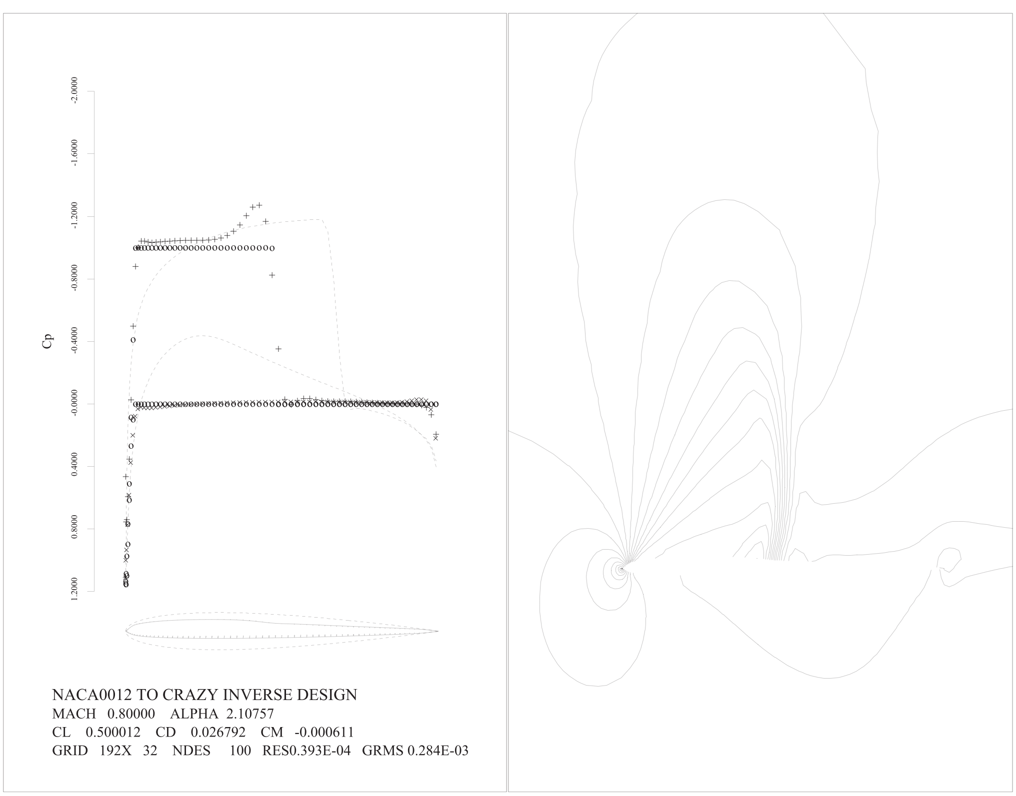

Fig. 17. NACA0012 to crazy inverse design at cycle 100.

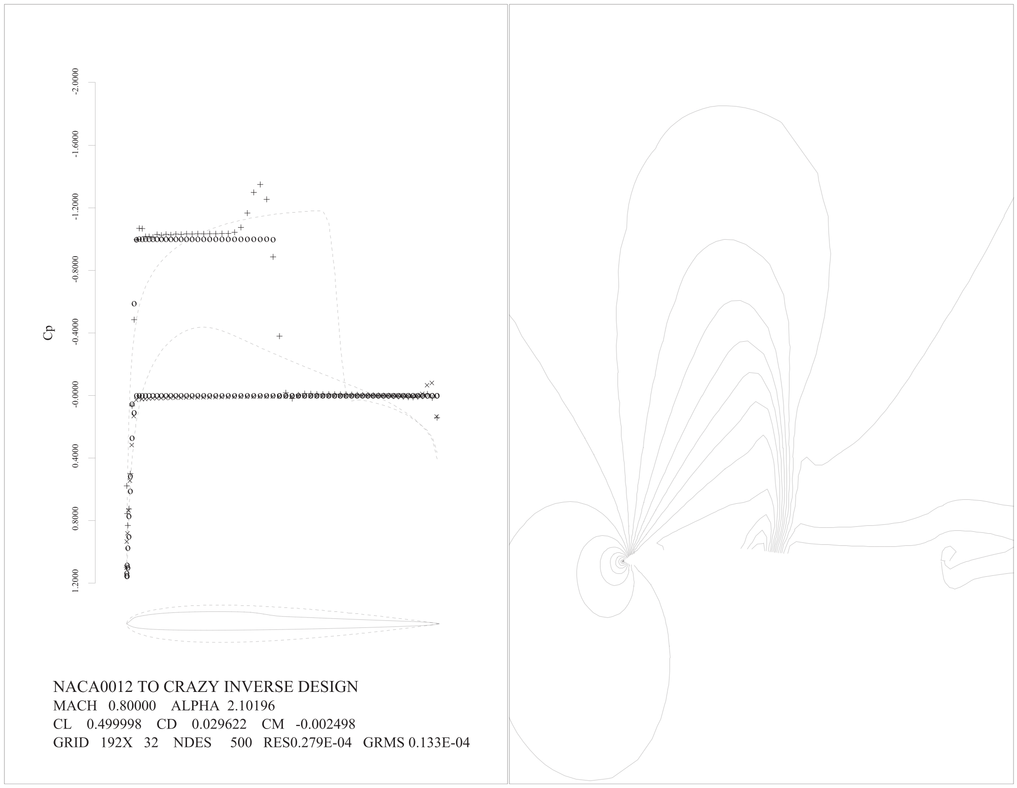

Fig. 18. NACA0012 to crazy inverse design at cycle 500.

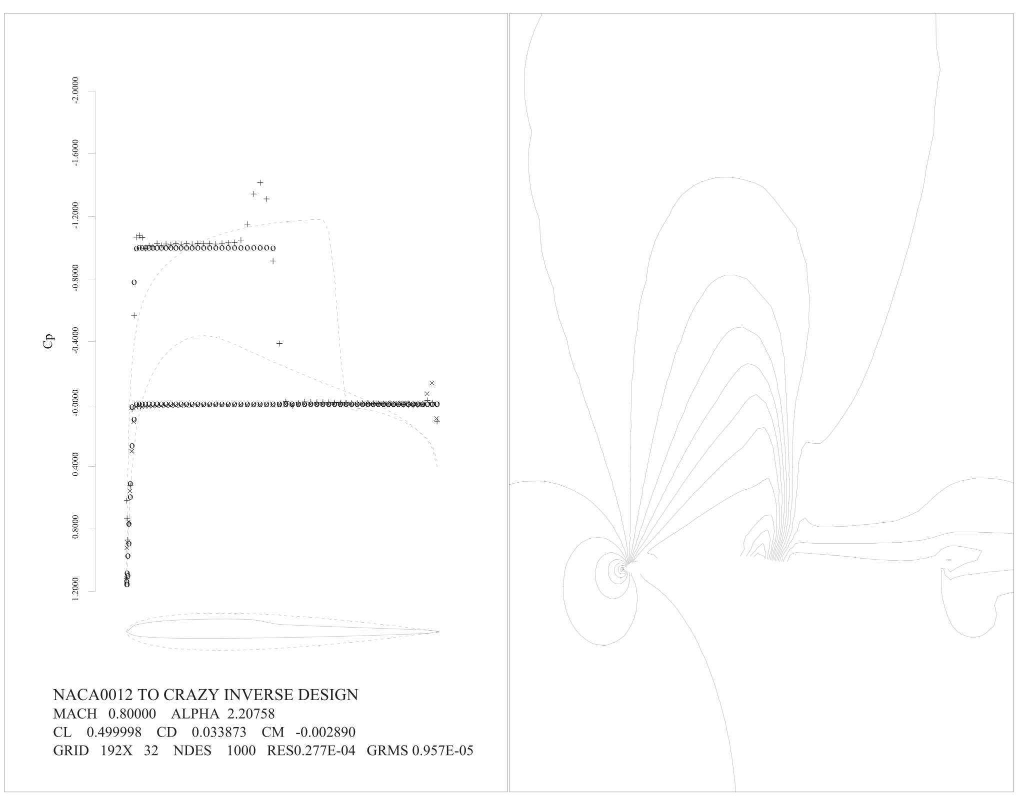

Fig. 19. NACA0012 to crazy inverse design at cycle 1000.

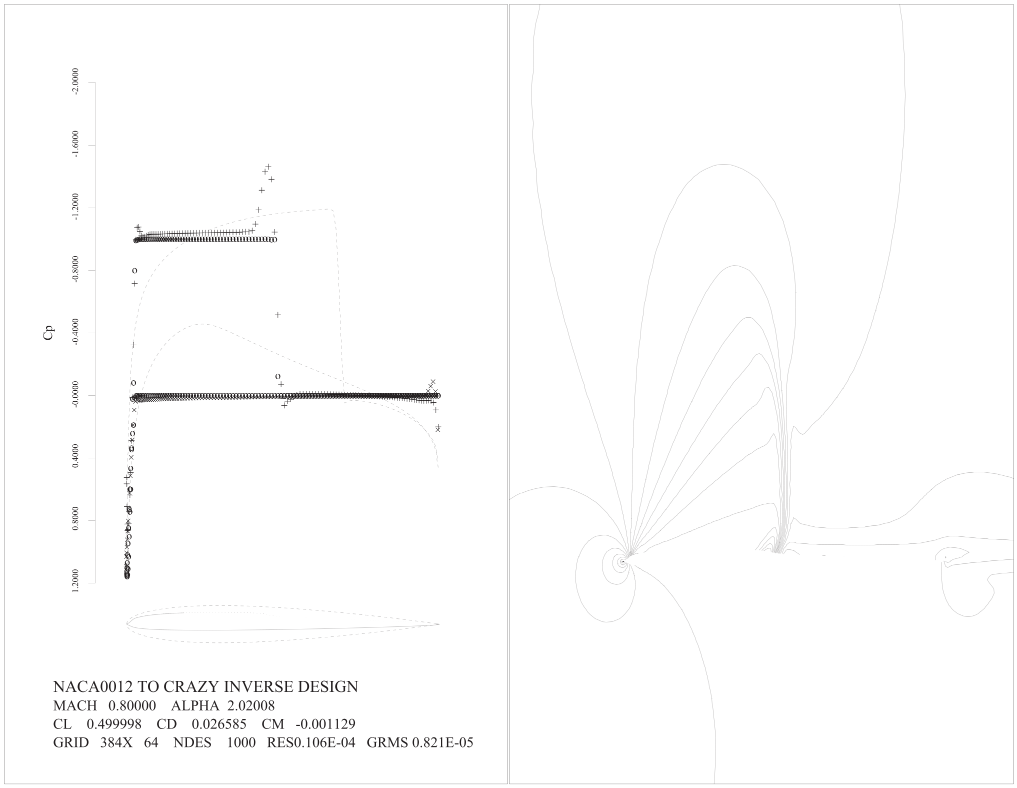

Fig. 20. NACA0012 to crazy inverse design at cycle 1000 on a finer mesh.

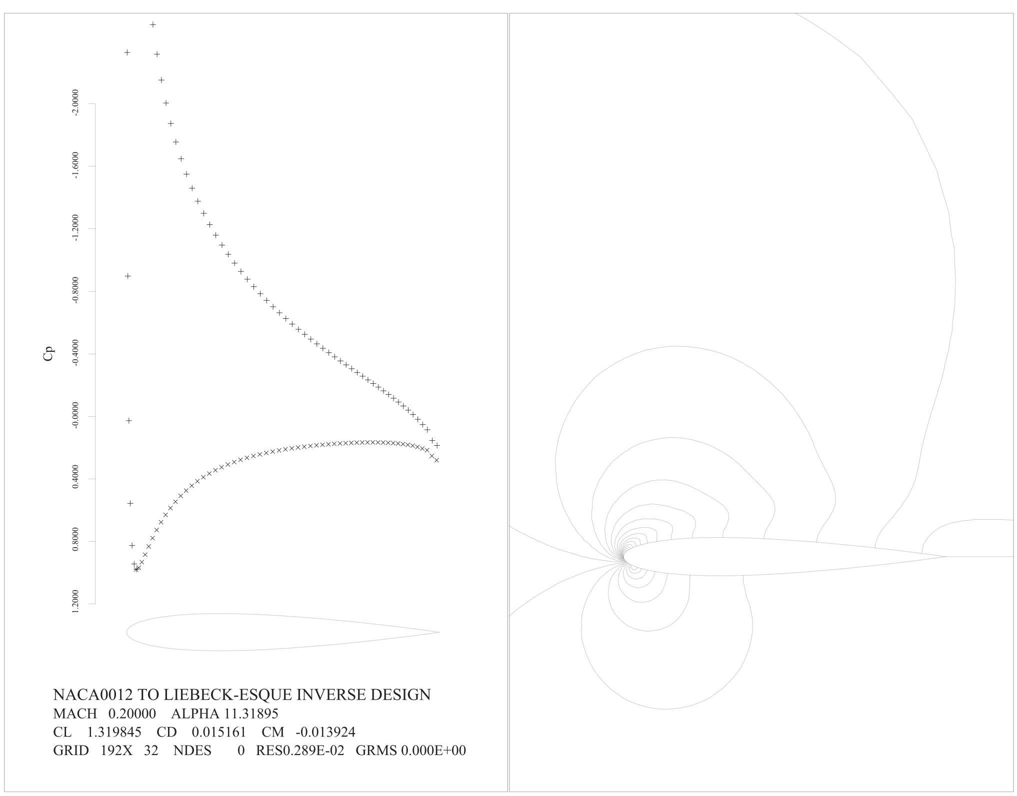

Fig. 21. NACA0012 airfoil pressure distribution, $M = 0.2$, $\alpha = 11.3^{\circ}$, $C_l = 1.32$.

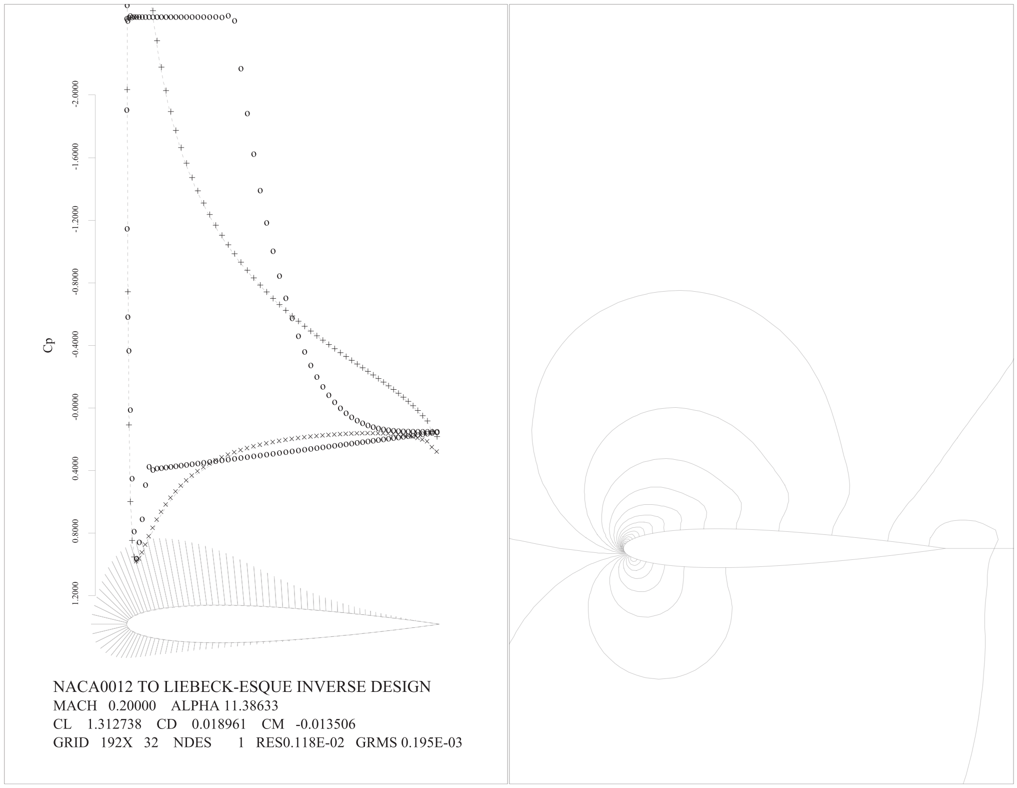

Fig. 22. NACA0012 to Liebeck-esque inverse design at cycle 1.

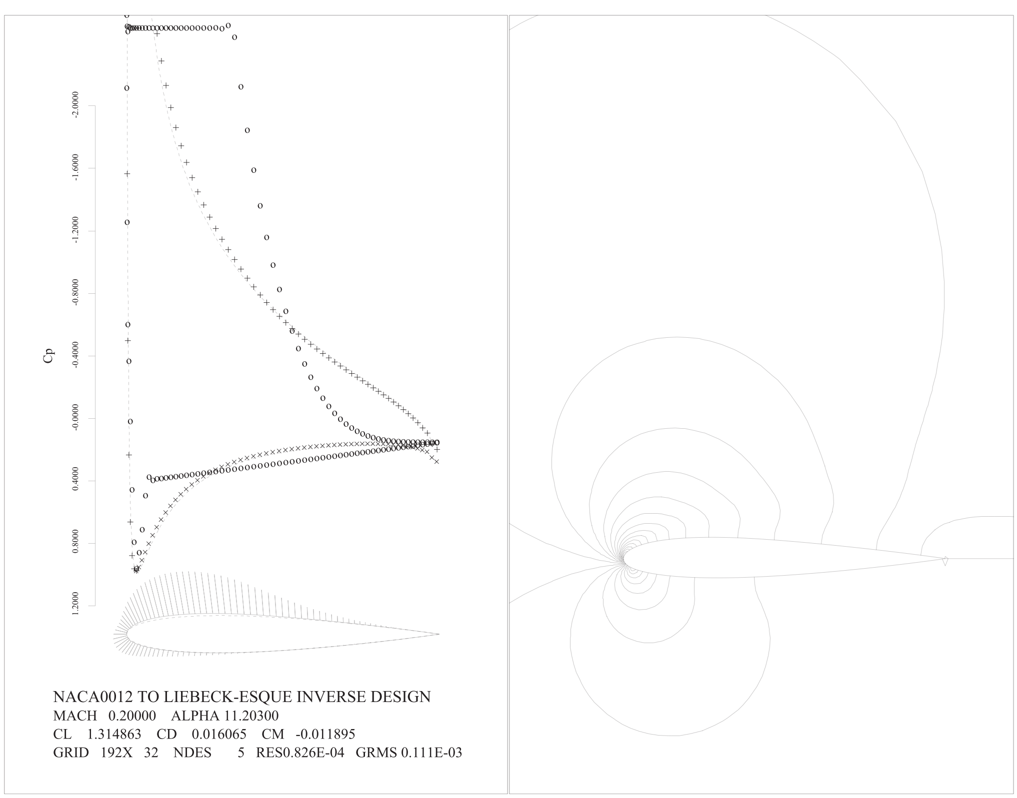

Fig. 23. NACA0012 to Liebeck-esque inverse design at cycle 5.

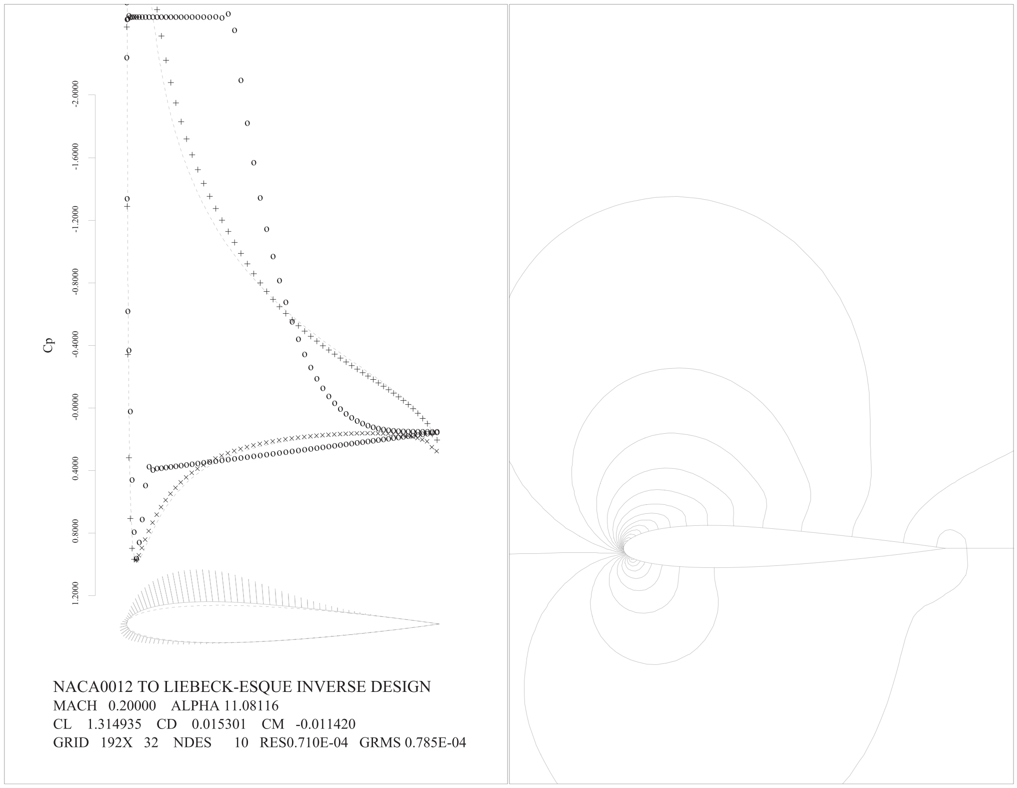

Fig. 24. NACA0012 to Liebeck-esque inverse design at cycle 10.

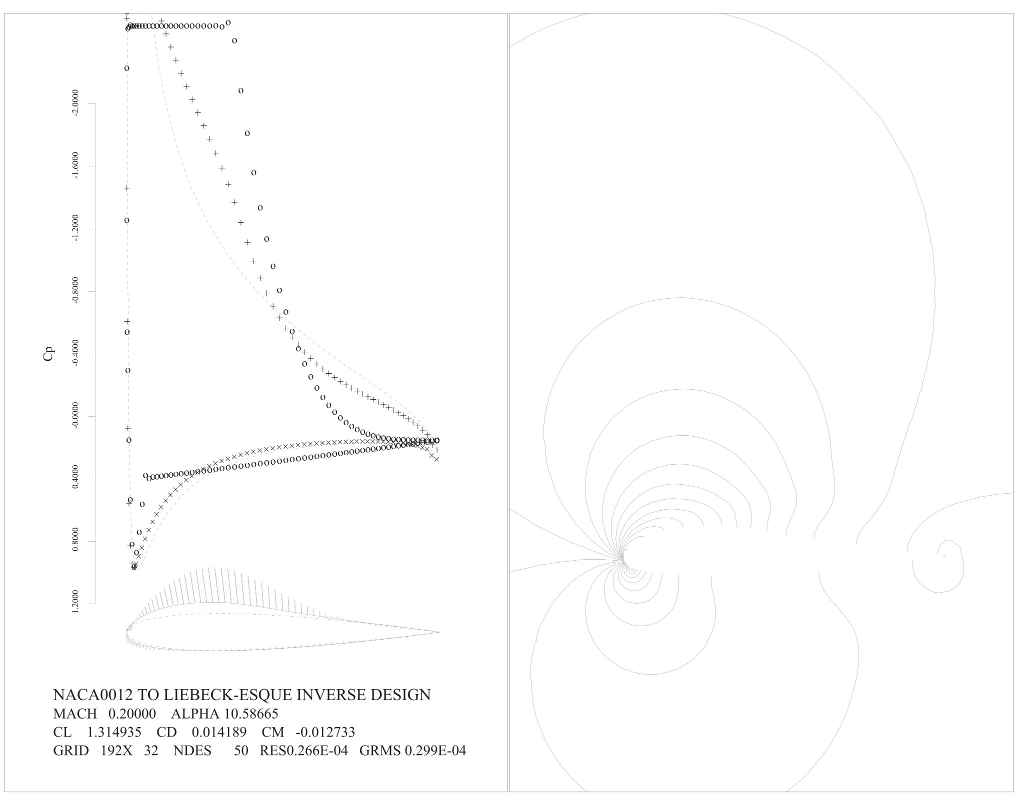

Fig. 25. NACA0012 to Liebeck-esque inverse design at cycle 50.

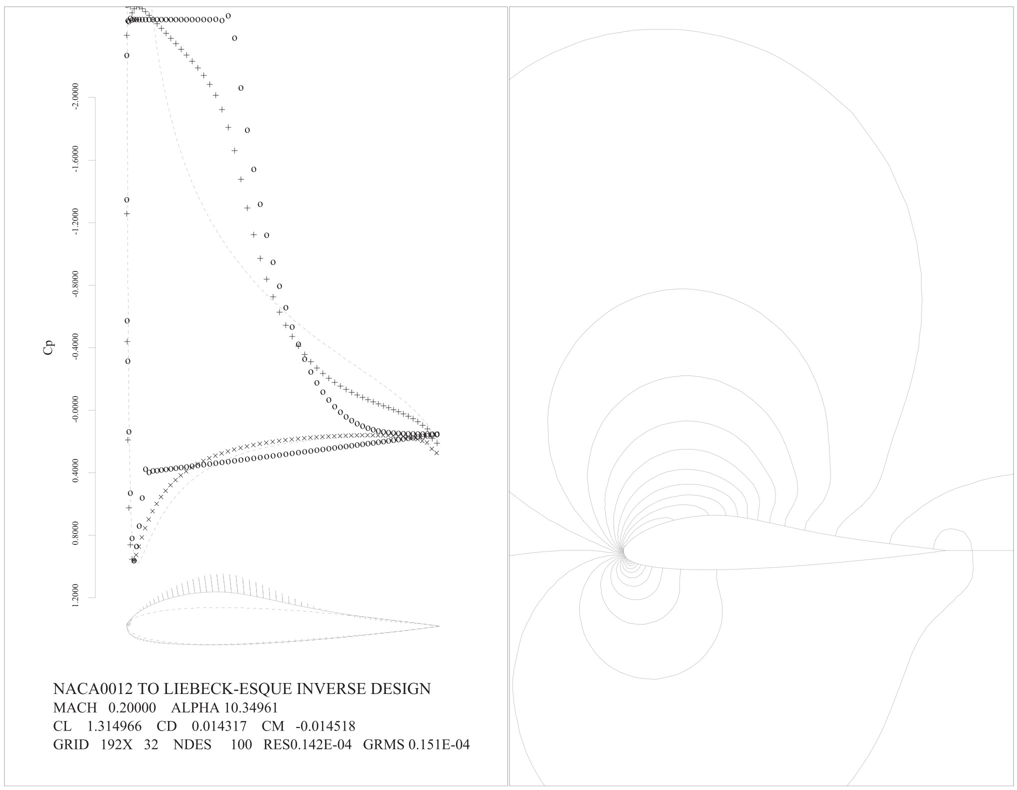

Fig. 26. NACA0012 to Liebeck-esque inverse design at cycle 100.

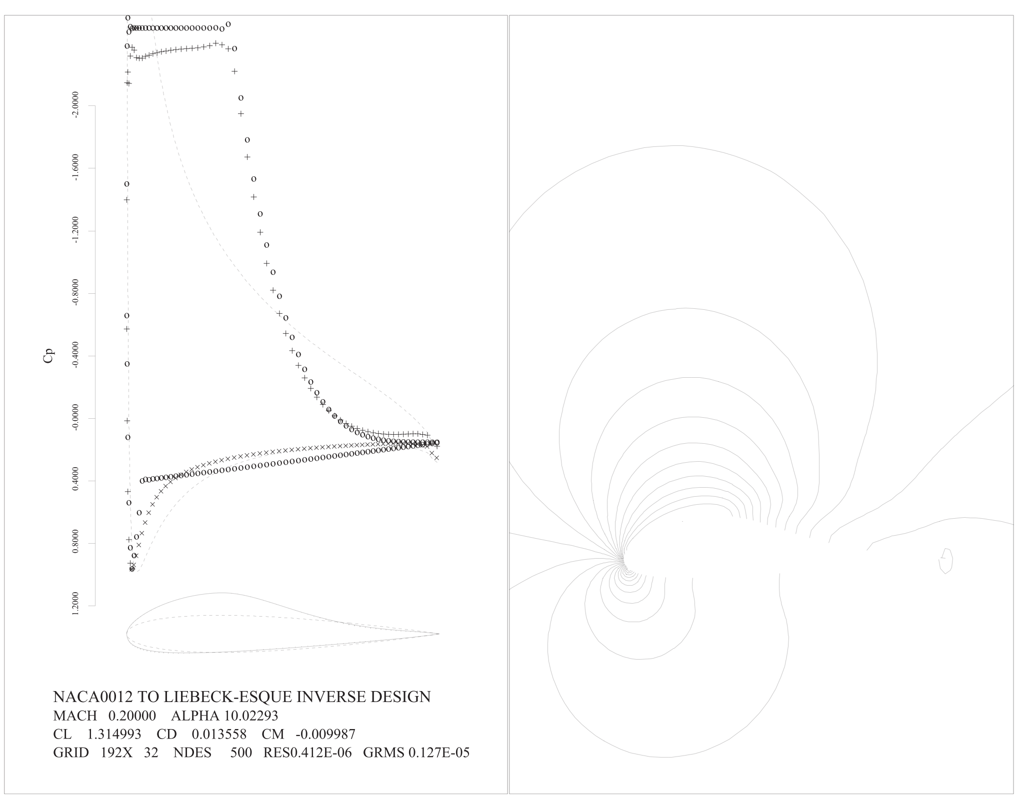

Fig. 27. NACA0012 to Liebeck-esque inverse design at cycle 500.

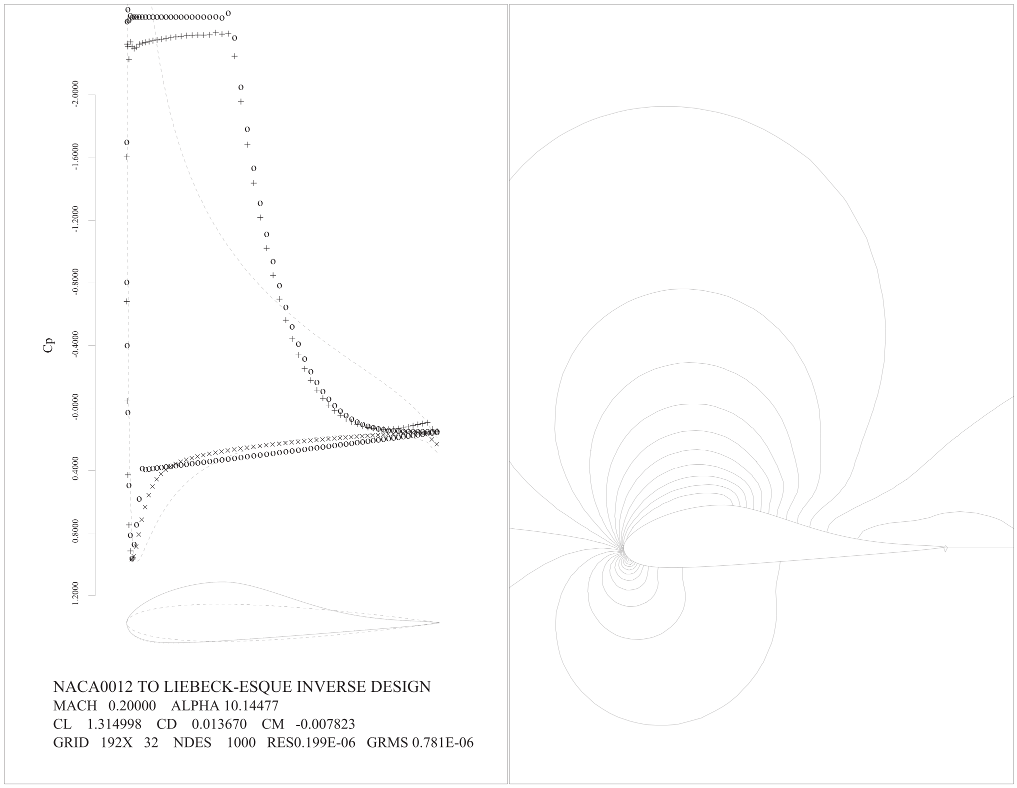

Fig. 28. NACA0012 to Liebeck-esque inverse design at cycle 1000.

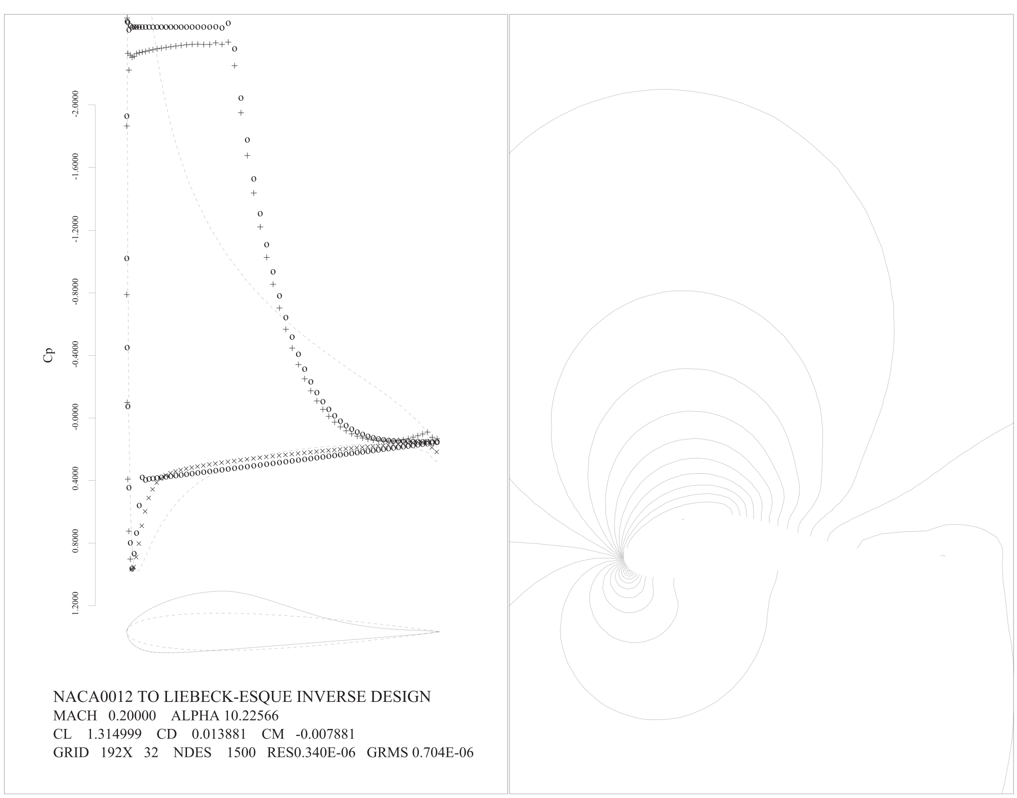

Fig. 29. NACA0012 to Liebeck-esque inverse design at cycle 1500.

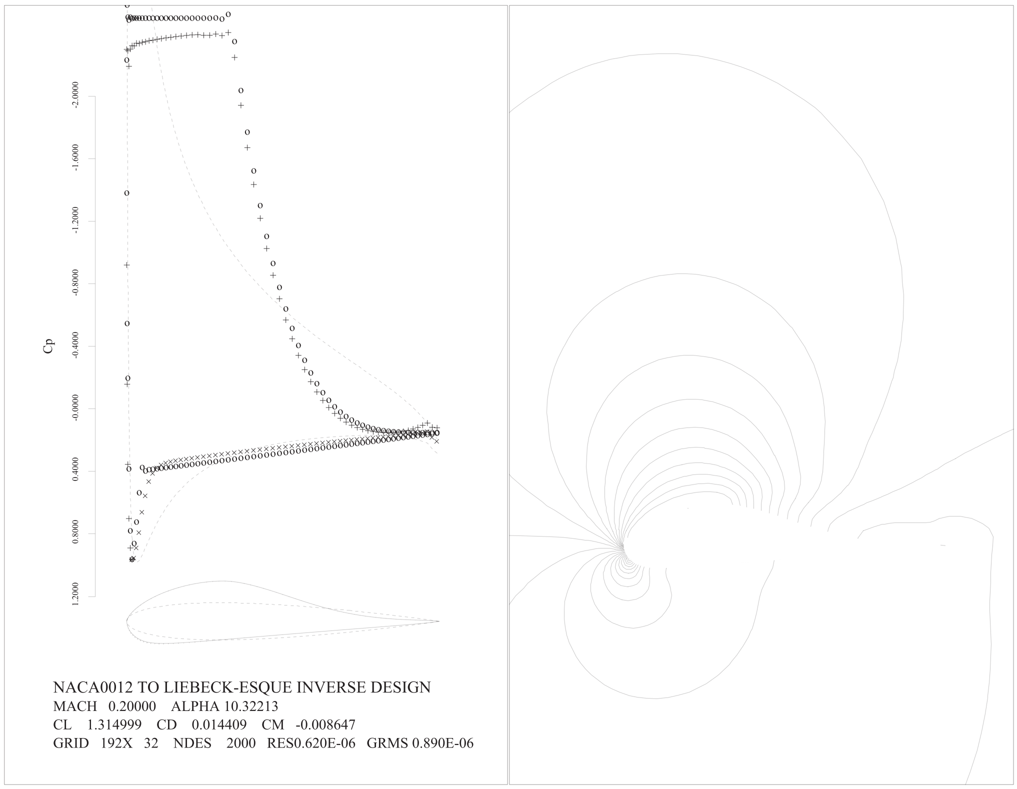

Fig. 30. NACA0012 to Liebeck-esque inverse design at cycle 2000.

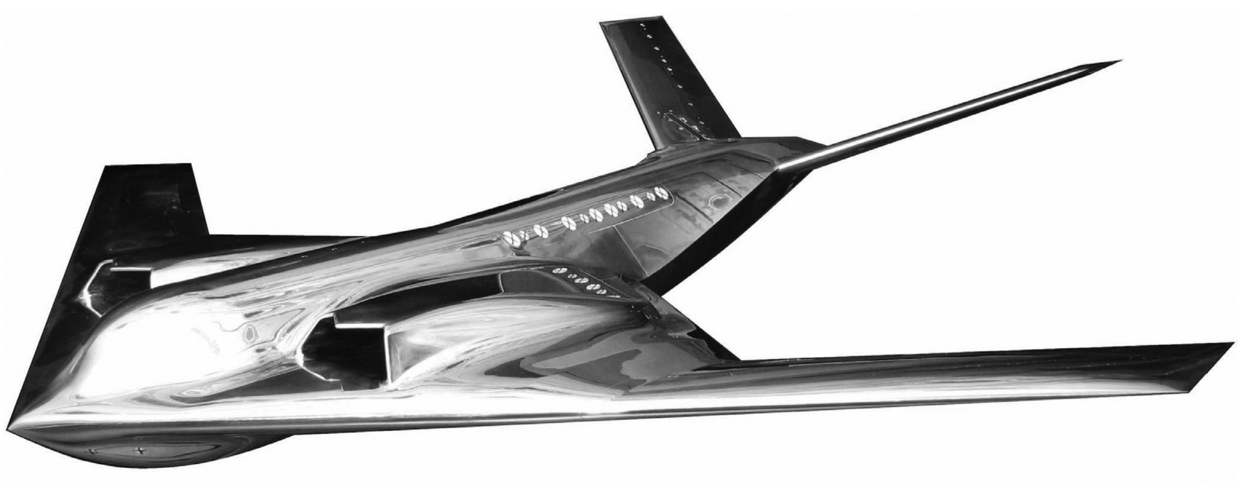

Fig. 31. Speed-Agile NTF cryogenic model, 3% Scale.

fairly well over most of the airfoil, the roof-top level is not quite recovered.

For those familiar with Liebeck's work, notice the similarity of the final design with that of his L1003 airfoil. This is to be expected as we fashioned our target pressures after a simplified caricature of the L1003 pressure distribution.

Summary

A small set of test cases for inverse aerodynamic design are presented herein. These cases range from very simple to fairly difficult, and even venture into the crazy absurd. The first two cases include manufactured, realizable target pressure distributions. The last two cases are based on unrealizable target pressures, yet yield reasonable results nonetheless. These test cases, in and of themselves, do not teach much of anything in the art of inverse design, but they do provide benchmarks to compare against. In order to learn this art form, one must play with this technique to find out what works (and what does not) to continually expand one's foundation.

A word of caution to the reader; when manufacturing a realizable target pressure distribution, be sure that the airfoil used to do so is supported by the design space utilized. This is especially important if only a small number of design variables define the space. Otherwise, a near-perfect recovery of the target pressures will be highly unlikely.

We hope to provoke thought and spark discussions to help reinvigorate the study of inverse design so that this art form is not lost to our future generations of aerodynamicists. On the surface, and to the uninitiated, inverse design may seem pointless when accurate drag minimization capabilities are now readily available. However, this is not the case, and again, it is not an either-or proposition. Just as CFD analyses and experimental testing provide complimentary strengths to understanding the performance of an aerodynamic design, inverse design and drag minimization provide complimentary capabilities to achieving optimum designs. An example of designing wind-tunnel walls was discussed earlier; obviously this problem cannot be solved with drag minimization.

Problems do not have to be difficult to be educational. As such, we offer these test cases to be taken under consideration by the Aerodynamics Design Optimization Discussion Group (ADO DG). These exercises are pertinent to researchers at all levels, and nicely compliment the current suite of ADO DG test cases. They are relatively inexpensive to conduct, are easy to set up, and do not evolve into a pathologically difficult problem to solve. We recommend formalizing the set of metrics to track, as well as to include a convergence on the dimension of the design space.

As a bonus, and a peek into how we use both inverse design and drag minimization techniques in collaboration with each other, refer to Fig. 31, which is an image of the Speed-Agile cryogenic model the authors designed about a decade ago. This wing-body-horizontal-tail model includes embedded flow-through nacelles which are designed to naturally aspirate at the powered mass-flow condition. In order to accurately analyze the aerodynamic performance of this configuration, a very large grid system was required, and with powered effects simulated. Performing a drag minimization on the full configuration as analyzed, would have taken several weeks to perform and months to get right. Instead, we applied a different approach which we had developed and had used before. We replaced the fuselage and nacelles with a simple root-plug extension of the inboard wing. This results in a wing-out-of-a-wall configuration, but it lacks the proper influence of the fuselage and flow-through nacelles on the exposed wing outboard of the nacelle-wing intersection region. To correct this, we redesigned the simple root-plug to provide an equivalent disturbance on the exposed wing by using the pressure distribution of the full configuration as the target on the exposed wing, and held the exposed wing geometry frozen. Now, with the new equivalent simple body (root plug) held frozen, we redesigned the exposed wing with drag minimization. The complete process from start to finish was performed over the course of a weekend. An analysis of the full configuration, with the optimized exposed wing, confirmed an equivalent drag reduction as observed during the optimization of the exposed wing installed on the equivalent simple body. This final analysis took longer than two days.

Whether you are just beginning your career in aerodynamic design, or are a seasoned veteran, studying the works of our classic pioneers is essential to building a solid foundation. All of us should have heroes to aspire to, some may be from a time centuries past, while others can be a colleague working side-by-side with you today. Take inspiration and knowledge from wherever and whenever it is available. With that, we conclude our discussion on the importance of inverse aerodynamic design, and we look forward to feedback from the broader aerodynamic design community on this subject.

Long Live the Art of Inverse Aerodynamic Design!

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Joukowsky NE. Uber die konturen der tragflachen der drachenflieger. Motorluflschiffahrt (German) 1910(1:281-284).
[2] Lighthill MJ. A new method of two dimensional aerodynamic design. Aeronautical Research Council London, 1945(R&M, 2112).
[3] Jameson A. Aerodynamic design via control theory. J Sci Comput 1988;3:233–60.
[4] Vassberg JC, Jameson A. Theoretical Background for Aerodynamic Shape Optimization. VKI lecture series. Brussels, Belgium: Von Karman Institute; 2018.
[5] Vassberg JC, Jameson A. Industrial Applications of Aerodynamic Shape Optimization. VKI lecture series. Brussels, Belgium: Von Karman Institute; 2018.
[6] Vassberg JC, Jameson A. Influence of Shape Parameterization on Aerodynamic Shape Optimization. VKI lecture series. Brussels, Belgium: Von Karman Institute; 2018.