

Expected Drag Minimization for Aerodynamic Design Optimization Based on Aircraft Operational Data

Rhea P. Liem¹

Hong Kong University of Science and Technology, Department of Mechanical and Aerospace Engineering, Hong Kong

Joaquim R. R. A. Martins², Gaetan K. W. Kenway³

University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI

Abstract

Aerodynamic shape optimization must consider multiple flight conditions to obtain designs that perform well in a range of situations. However, multipoint studies have relied on heuristic choices for the flight conditions and associated weights. To eliminate the heuristics, we propose a new approach where the conditions and weights are based on actual flight data. The proposed approach minimizes the expected drag value given by a probability density function in the space of the flight conditions, which can be estimated based on data from aircraft operations. To demonstrate our approach, we perform drag minimizations of the Aerodynamic Design Optimization Discussion Group Common Research Model wing, for both single-point and multipoint cases. The multipoint cases include five- and nine-point formulations, some of which approximate the expected drag value over the specified flight-condition probability distribution. We conclude that if we focus on the resulting design, a five-point optimization with points based on the flight-condition distribution and equal weights is sufficient to obtain an optimal shape with respect to the expected drag value. However, if it is important to retain the accuracy of the expected drag integration at each optimization iteration, we recommend the proposed approach.

Keywords: Aerodynamic shape optimization, multipoint optimization, expected value, numerical integration, mission analysis, surrogate models, aircraft design

¹Assistant Professor

²Professor

³Research Investigator

1. Introduction

Prior to 1960, the aircraft design process relied mainly on flow visualization techniques and wind-tunnel experiments using pressure and force measurements [18]. Computational aerodynamics was brought about by radical improvements in numerical algorithms coupled with advances in computing technologies. Early investigations of aerodynamic shape optimization include those by Hicks et al. [15], Hicks and Henne [14], and Constantino and Holst [5] in the late 1970s and early 1980s. In these early efforts, full potential flow solvers were coupled with conjugate gradient optimization algorithms to enable the automated design of airfoils and wing shapes. Major advances in aerodynamic shape optimization occurred in the late 1980s and early 1990s, when adjoint methods made it possible to efficiently compute shape gradients [18, 16, 40, 41]. With adjoint methods, the cost of computing the gradients became independent of the number of design variables, which enabled detailed optimization based on high-fidelity models.

The first adjoint-based investigations focused on drag minimization at a single flight condition [2, 33, 37, 39]. Single-point optimized designs suffer performance degradations at off-design conditions [18, 4]. The drag polar for single-point designs features a cusp because the optimization eliminates the shock at the nominal flight condition, while making the shock much stronger at off-design conditions. This cusp tends to become more prominent as the number of design variables increases [9].

Because of the limitations of single-point optimizations, it is necessary to consider multiple flight conditions in aerodynamic shape optimization. Jameson [18] pioneered this effort, seeking a compromise design by taking the sum of the cost functions for several design points. The design problem was formulated as a control problem, where the cost function measured the wave drag and the deviation from a desired pressure distribution. The wave drag was added to the cost function with a multiplier, which could be varied to alter the trade-off between drag reduction and deviation from the desired distribution.

The most common approach in multipoint formulations has been a composite objective function, typically expressed as a weighted sum of the drag coefficient over several flight conditions [34]. To emulate the different flight conditions, Reuther et al. [38] varied the lift coefficient in an unconstrained transonic optimization, and Drela [9] varied lift coefficient for a low Reynolds number airfoil optimization. In some airfoil optimization problems, the Mach numbers in the multipoint formulation were varied [9, 34]. Other authors went a step further and varied both Mach number and lift coefficient [18, 11, 4]. Multipoint optimization results have been shown to be more robust and thus more practical than those of single-point optimizations [9, 18, 38, 4, 31, 22]. Furthermore, the optimizer typically increases the drag-divergence Mach numbers [9, 22].

Motivated by the desire to compare the various aerodynamic design optimization methods, researchers formed the Aerodynamic Design Optimization Discussion Group (ADODG)⁴, which is sponsored by the American Institute of Aeronautics and Astronautics (AIAA). The ADODG cases range from a 2D airfoil
35 inviscid drag minimization to a full-configuration multipoint drag minimization based on the solution of the Reynolds-averaged Navier–Stokes equations (RANS). For the inviscid airfoil optimization problem, Méheut et al. [32] performed gradient-based single-point and multipoint aerodynamic optimizations. They performed cross-validation to compare six different optimized shapes that were produced using different grid types, solvers, and postprocessing procedures. For the Common Research Model (CRM) wing drag mini-
40 mization using RANS, Méheut et al. [32] showed that the two three-point optimization problems yielded a consistent performance improvement, and Lee et al. [24] showed similar trends. Lyu et al. [31] and Kenway and Martins [22] compared the single-point and multipoint optimizations of the CRM wing with five to nine points; they quantified the robustness of the various cases by plotting contours of the performance over the flight-condition space.

45 The goal of this paper is to address two major questions in multipoint aerodynamic design optimization: (1) Which flight conditions should be considered, and (2) how much weight should be attributed to each of these conditions? Although designers generally know the nominal flight condition based on the particular aircraft mission, as well as the likely range of the flight conditions, it is not clear how to translate this information into a multipoint drag minimization formulation. In previous work, the multipoint formulation
50 has been based on common sense or prior design experience, which is somewhat arbitrary [4, 9]. Lyu et al. [31] and Kenway and Martins [21] assigned equal weights in a multipoint drag minimization with five flight conditions, which has been a popular approach in multipoint optimization. Buckley and Zingg [3] employed an integration rule to formulate a weighted-integral objective function in their multipoint airfoil optimization. Each flight condition was assigned a weighting function based on design experience. For
55 high-fidelity aerostructural design optimization, Liem et al. [25] chose the flight conditions to be considered and the associated weights based on a histogram of the aircraft missions. Gallard et al. [12] analyzed the linear dependencies between the shape gradients and computed a minimal set of flight conditions.

We propose a new formulation that selects the flight conditions based on a flight-condition probability density function (PDF). When this PDF is used, the expected performance provides a first-moment measure

⁴AIAA Aerodynamic Design Optimization Discussion Group, <https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/default.aspx> (accessed 11 July 2016)

60 of the real-world performance. The exact PDF, however, is typically unknown. Therefore, we replace the PDF with a distribution generated based on publicly available flight mission data. To demonstrate the proposed approach, we perform our drag minimizations for the ADODG CRM wing geometry, for the ADODG single-point and multipoint cases. We also perform a series of studies to compare the various multipoint approaches and to analyze the effect of the number of flight conditions considered. In the proposed formula-
65 tion, we minimize the expected performance over all flight conditions, accounting for the time that is spent at each condition.

We describe the basic optimization problem and the numerical tools used in Section 2. Then, in Section 3 we explain the optimization formulations to be compared, including the proposed expected drag minimization approach. We then discuss our numerical results and findings in Section 4, and Section 5 provides
70 concluding remarks.

2. Aerodynamic Design Optimization Approach

This section describes the aerodynamic design optimization approach, including the numerical tools, the overall optimization problem formulation (objective function, design variables, and constraints), and the baseline geometry that is optimized. The overall approach and geometry have been previously presented
75 by Lyu et al. [31] and Kenway and Martins [22], but we add the new formulation for the expected drag minimization. The baseline wing geometry and specifications follow those given by the ADODG, consisting solely of the wing from the CRM full configuration. The single-point optimization benchmark has been previously solved by various groups [31, 32, 43], and the baseline and optimized geometries and meshes are provided by Lyu et al. [31]⁵. Multipoint optimization results for this case are presented by Kenway and
80 Martins [22].

2.1. Multilevel Optimization Acceleration

To reduce the computational cost of the high-fidelity multipoint optimization, we first perform the optimization on a coarse grid and then find the final solution on a fine grid, following the work of Lyu et al. [31] and Kenway and Martins [22]. The coarse result is used as the starting point for the optimization on the fine
85 grid. This multilevel approach accelerates the optimization process because iterations on the coarse grid are

⁵<http://mdolab.engin.umich.edu/content/aerodynamic-design-optimization-workshop> (Accessed January 14, 2017)

faster and cheaper. However, it is important to ensure that the coarse grid captures the main characteristics of the flow, e.g., the shock strength and location. The geometries for the coarse and fine meshes used in this optimization are briefly described next.

2.2. Computational Meshes

90 We use two different meshes in the optimization: a coarse mesh with 450 k cells (L2) and a fine mesh with 3.5 M cells (L1), as shown in Figure 1. The multiblock meshes are generated using a hyperbolic mesh generator and exhibit an O-type topology. Kenway and Martins [22] presented a convergence study for these meshes as well as a finer mesh with over 28 M grid cells (L0) and an extrapolated zero-mesh spacing. The L0 mesh is not considered in this work because the optimization is too expensive [31]. The convergence study
95 showed that the drag-count difference between the L1 mesh and the zero-mesh spacing is $\mathcal{O}(1)$, an error of 1.5%. The difference in drag (when comparing the baseline and optimized designs) is $\mathcal{O}(10^{-2})$ drag counts, corresponding to a relative error of 0.4%. Thus, the L1 mesh was considered a good compromise between computational cost and accuracy.




Figure 1: Two O-meshes are used in a two-level optimization.

2.3. Optimization Problem Formulation

100 The optimization problem is formulated as a drag minimization, subject to aerodynamic and geometric constraints. The formulation is similar to the ADODG CRM wing single-point and multipoint bench-

marks [31, 22], and it is summarized in Table 1. It differs from the ADODG benchmarks only in the formulation of the objective function.

Specifically, the objective function is the expected drag coefficient, $\mathbb{E}[C_D]$, over the operating conditions. The expectation integral is approximated numerically using a number of points in the flight-condition space, giving a multipoint optimization. The flight-condition weights w_j are determined based on the integration scheme and the flight-condition PDF. The multipoint formulations are further detailed in Section 3.

	Function/variable	Description	Quantity
minimize	$\mathbb{E}[C_D] \approx \sum_{j=1}^M w_j C_{D_j}$	Expected value of C_D	
with respect to	$\mathbf{x}_{\text{shape}}$	z perturbation of FFD control points	768
	α_j	Angle of attack for each flight condition j	M
	Total design variables		768 + M
subject to	$C_{L_j} - C_{L_j}^* = 0$	Lift constraint at each flight condition j	M
	$C_M \geq -0.17$	Moment constraint at nominal condition	1
	$V \geq V_{\text{baseline}}$	Volume constraint	1
	$t_i \geq 0.25t_{i_{\text{baseline}}}$	Thickness constraint at each point i	750
	Total constraints		752 + M

Table 1: Multipoint aerodynamic shape optimization formulation

The design variables consist of the wing shape design variables, $\mathbf{x}_{\text{shape}}$, and the angle of attack for each flight condition, α_j . The wing shape is parameterized using the free-form deformation (FFD) volume approach with trivariate B-spline volumes [22]. Figure 2 shows the two FFD volumes used in this work: the coarse (Figure 2a) and fine (Figure 2b) FFDs. The fine FFD has 768 control points that move in the vertical (z) direction. This enables the optimizer to change the airfoil shapes and twist distribution, while keeping the wing planform shape constant. In total, there are $768 + M$ design variables, where M is the number of flight conditions.

The aerodynamic constraints include the lift-coefficient constraint for each flight condition, and a single pitching-moment inequality constraint enforced at the nominal condition ($M = 0.85$, $C_L = 0.5$). The

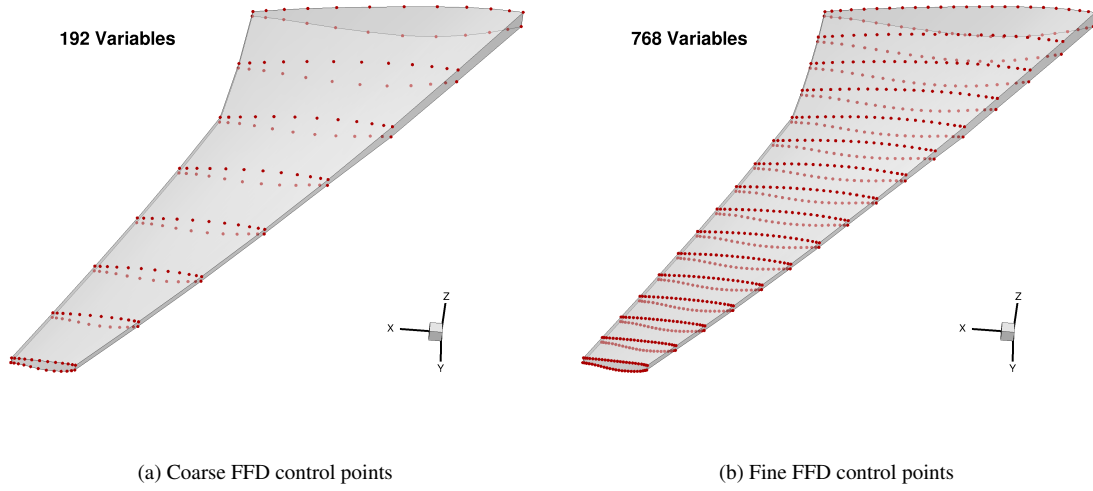


Figure 2: The wing geometry is parameterized by moving FFD control points in the z direction.

remaining constraints are geometric and consist of a volume constraint and a series of thickness constraints. The volume constraint ensures that the internal volume of the wing is greater than or equal to the baseline volume. The wing thicknesses, which are computed at 750 locations, must be greater than or equal to 25% of the baseline thickness. There is a total of $752 + M$ constraints in each multipoint problem: M lift constraints, 1 moment constraint, 1 volume constraint, and 750 thickness constraints,

The flight conditions considered in the multipoint objective function correspond to the cruise segment of the flight, and no off-design conditions are considered. This optimization is meant to be an academic problem focusing on the multipoint objective function formulation. For a more realistic aircraft design problem, we would need to include off-design conditions such as high-speed dive and low-speed performance [23].

2.4. CFD Solver

The flow is modeled by the compressible RANS equations using ADflow [44]. ADflow is a finite-volume, cell-centered multiblock flow solver. It uses the Jameson–Schmidt–Turbel (JST) artificial dissipation scheme [17] and the Spalart–Allmaras (SA) turbulence model. Both the mean flow and turbulence equations use a diagonally dominant alternating direction implicit (DDADI) scheme to solve the Navier–Stokes equations. To efficiently compute the gradients of the functions of interest with respect to the large number of shape variables, we use a discrete adjoint method implemented with algorithmic differentiation,

which includes the effect of the turbulence model. More details on this adjoint implementation are provided by Lyu et al. [29].

135 2.5. Optimization Algorithm

Multipoint aerodynamic shape optimizations based on RANS CFD solutions are costly. Therefore, it is desirable to use a gradient-based optimization method to reduce the required number of function evaluations. Studies comparing different algorithms in the context of aerodynamic design optimization have shown that genetic algorithms and other gradient-free approaches are prohibitively expensive and not sufficiently accurate [45, 30]. For optimization problems similar to that presented in this paper, Lyu et al. [31] showed that local minima did not pose any significant issues for the CRM wing case, which justified the use of gradient-based optimization. In this paper, the problems are solved with SNOPT, which is an optimizer based on the sequential quadratic programming approach [13], through the Python interface pyOpt [36].

3. Multipoint Objective Function Formulations

145 As mentioned in the previous section, the optimization problem is identical to the ADODG CRM wing benchmark except for the objective function formulation and the flight conditions considered. In this section we describe all the multipoint formulations that we investigated, including the proposed expected drag coefficient minimization.

3.1. Multipoint and Expected Drag Formulations

The objective function of a multipoint aerodynamic optimization is typically formulated as a weighted sum of the drag coefficient evaluated at different points,

$$f_{\text{obj}} = \sum_{j=1}^M w_j C_{D_j}, \quad (1)$$

150 where w_j is the weight assigned to the j^{th} flight condition, and M is the number of flight conditions considered.

One of the keys to obtaining a design that is optimal in practice is the choice of the flight conditions and their weights in the multipoint optimization formulation. As previously mentioned, the conditions and weights are usually based on design intuition. To automate the process, we minimize the expected value

of the drag coefficient based on the distribution of the flight conditions. Suppose we have a PDF in the flight-condition space $p(M, C_L)$ given over a domain Ω . The expected value of the drag coefficient is

$$\mathbb{E}[C_D] = \iint_{\Omega} C_D(M, C_L) p(M, C_L) dM dC_L. \quad (2)$$

When the flight-condition space is rectangular, this value can be computed using numerical integration with $m \times n$ quadrature points

$$\mathbb{E}[C_D] \approx \sum_{i=1}^n \sum_{k=1}^m \tau_{ik} C_D(M_i, C_{L_k}) p(M_i, C_{L_k}), \quad (3)$$

where τ_{ik} is the quadrature weight assigned to each point. Comparing this approximation to the multipoint objective (1), we see that the products $\tau_{ik} \cdot p(M_i, C_{L_k})$ correspond to the weights w_j , and $C_D(M_i, C_{L_k})$ corresponds to C_{D_j} . Therefore, an expected drag minimization can be written as a multipoint drag minimization with an appropriate choice of weights.

The expectation integral can be approximated using various numerical integration rules [8, 6, 7]. We mainly use the midpoint integration rule. For the 9-point ADODG case (Section 3.3.1) we follow the original formulation that uses the trapezoidal rule. When we assume that the PDF is a normal distribution, we use the Gauss–Hermite quadrature rule [1] (Section 3.3.4).

3.2. Flight-Condition Probability Density Function

We develop a more realistic objective function by computing the expected drag value, where the PDF is derived from actual aircraft operation data. To construct the PDF, we obtained historical flight data from the Bureau of Transportation Statistics⁶ detailing the missions of all Boeing 777-200ER flights departing from and arriving in the United States in 2011 [25]. The construction of the flight-condition histogram based on these missions has been previously presented by the authors [25], and it involves a surrogate-based mission analysis for the cruise segments of each flight mission [26]. A mission analysis evaluates the range equation via numerical integration, by discretizing the mission profile into a number of integration intervals. The analysis gives the flight conditions (Mach number, altitude, angle of attack, tail rotation angle) and the aerodynamic data (C_L , C_D , C_M) at each integration interval i . When all the flight missions have been analyzed, the complete database is generated. A nominal Mach number of 0.85 is assumed in the analysis.

⁶TranStats, Bureau of Transportation Statistics <http://www.transtats.bts.gov/> (Accessed 14 July 2016)

This Mach number is shifted from the nominal value of the Boeing 777 aircraft (0.84) to adjust for the CRM nominal values.

The mission analysis generated a total of 21 430 flight-condition data points. We generated the histogram by assigning these data points to 30×30 grid cells. The corresponding bin intervals for Mach and C_L are 0.0014 and 0.0049, respectively. The flight-condition histogram shown in Figure 3 shows the number of flight conditions in each Mach- C_L bin. From this distribution, we identified a high-frequency region, which is highlighted by the red rectangle. This region is within a range of Mach = $[0.835, 0.865]$ and $C_L = [0.4525, 0.5375]$, beyond which the frequency quickly drops.

This formulation is equivalent to an optimization under uncertainty problem, where the flight-condition histogram can be replaced by a PDF. When historical data are not available, the PDF can be derived based on market research or a survey of customer needs. Our method is applicable regardless of how the flight-condition distribution is generated.




Figure 3: Flight-condition histogram with high-frequency region (red rectangle). The color-mapped frequency values are the number of evaluated flight conditions in each Mach- C_L bin.

In the next section we describe the optimization cases and how we approximate the expected-value objective function based on the flight-condition histogram with numerical integration using both five and

185 nine integration points.

3.3. Optimization Cases

We now describe all the optimization cases that we solved, including multipoint drag minimizations with different point stencils and expected drag minimization cases. The cases are listed in Table 2, and the results are presented in Section 4.

Table 2: The optimization cases include a reference single-point case, three standard multipoint cases, and three cases based on PDF integration

Case	Number of points	Mach [min, max]	C_L [min, max]	Integration rule	PDF
1	1	0.85	0.5	–	–
9tu-ADODG	9	[0.820, 0.880]	[0.420, 0.591]	Trapezoidal	Uniform
5	5	[0.840, 0.860]	[0.475, 0.525]	–	–
5h	5	[0.840, 0.860]	[0.475, 0.525]	–	Flight histogram
9mu	9	[0.840, 0.860]	[0.475, 0.525]	Midpoint	Uniform
9mh	9	[0.840, 0.860]	[0.475, 0.525]	Midpoint	Flight histogram
9gn	9	[0.838, 0.862]	[0.465, 0.535]	Gauss–Hermite	Normal

190 The single-point optimization (Case 1) is performed for comparison purposes. This case is identical to the single-point ADODG CRM wing case presented by Lyu et al. [31]. We investigate a total of six multipoint cases: three that correspond to our proposed expected-drag minimization formulation and three for comparison purposes. The nine-point ADODG case, labeled Case 9tu-ADODG, has previously been solved by Kenway and Martins [22] and Lee et al. [24]. Case 5 uses a five-point stencil similar to that used
195 by Lyu et al. [31] and Kenway and Martins [21]. The other three cases (9mu, 9mh, and 9gn) correspond to the proposed approach detailed in Section 3.2.

The optimization cases considered differ in the number, selection, and weights of the flight conditions. The five-point cases (5 and 5mh) do not perform a strict 2D integration since they do not consider a rectangular stencil of points. The nine-point cases have a complete 3×3 stencil for the numerical integration.
200 The cases labeled with a “u” (9tu-ADODG and 9mu) assume a uniform flight-condition PDF. In the cases

with an “h” (5mh and 9mh), the flight conditions and weights are determined based on the flight-condition histogram. We use “m” and “t” in the case names to refer to the midpoint and trapezoidal integration rules, respectively. Finally, Case 9gn approximates the flight conditions as a normal PDF (“n”), which is then integrated exactly using the Gauss–Hermite integration rule (“g”).

205 3.3.1. Conventional Multipoint Optimization

Case 9tu-ADODG is defined by the ADODG as Case 4.6 [22]. The flight conditions and weights, which are normalized by the sum, are listed in Table 3. These weights are proportional to those of the trapezoidal rule for numerical integration, even though the objective was not formulated strictly as an expectation integral approximation. The formulation implicitly assumes that the whole region is of equal importance, which
 210 is equivalent to assuming a uniform PDF for the flight conditions. The cruise Mach number range (0.82–0.88) and C_L range (0.420–0.591) considered cover a larger region in the flight-condition space than those used in the other cases. Figure 4 clearly shows that the distribution of the selected flight conditions extends beyond the histogram considered in this case. These points include extreme flight conditions at which the aircraft would not usually operate, and thus the drag values at these points are not important in the design.
 215 However, the points are included in the optimization formulation to provide a margin for buffet onset. Buffet occurs when shock-induced separation interacts with the boundary layer to cause unsteady flow and causes vibration of the airframe, which is undesirable. Including these points in the optimization helps prevent strong shocks, and this indirectly prevents buffet from occurring. Ideally, these off-design points would be implemented as constraints [23].

220 Case 5 uses a cross stencil in Mach- C_L space about the nominal condition with constant weights; a similar case was solved by Lyu et al. [31] and Kenway and Martins [21]. The nominal point is perturbed in each direction, and equal weights are assumed for all flight conditions. The range in cruise Mach is the same (0.85 ± 0.01), but the range in C_L is narrower (0.5 ± 0.025). The narrower C_L range is more consistent with the flight-condition distribution relevant to this problem, which is also used in cases 5mh, 9mu, and 9mh.
 225 These five points are illustrated in Figure 5, and Table 4 lists the points and weights (w_j).

3.3.2. Nine-Point Configuration with Uniform Distribution

The nine integration points are chosen from the high-frequency region shown in Figure 3. We use a 3×3 regular grid stencil to distribute the integration points uniformly in this region. This point distribution is shown in Figure 6, where the red rectangles show the cells corresponding to each integration point, and ΔM
 230 and ΔC_L denote the size of the cell in each direction. When a uniform distribution is assumed, the PDF at




Figure 4: Flight conditions and corresponding weights for Case 9tu-ADODG; the sizes of the dots are proportional to the weight values.

Mach	C_L	Re	w_j
0.82	0.483	4.82×10^{-6}	0.0625
0.82	0.537	4.82×10^{-6}	0.1250
0.82	0.591	4.82×10^{-6}	0.0625
0.85	0.450	5.00×10^{-6}	0.1250
0.85	0.500	5.00×10^{-6}	0.2500
0.85	0.550	5.00×10^{-6}	0.1250
0.88	0.420	5.18×10^{-6}	0.0625
0.88	0.466	5.18×10^{-6}	0.1250
0.88	0.513	5.18×10^{-6}	0.0625

Table 3: Flight conditions for Case 9tu-ADODG

points across the integration domain is $1/(3\Delta M \cdot 3\Delta C_L)$, since there are three cells in each direction for this configuration. Using this PDF in the midpoint integration rule, we obtain the weights listed in Table 5.

When we use the flight-condition distribution to approximate the PDF, we use the same integration domain and point distribution as for Case 9mu. However, in this case the PDF is approximated based on the flight-condition histogram. This approximation is illustrated in Figure 7 in grayscale: the darker the shade, the higher the PDF value.

We derived the nine weight values based on the flight-condition histogram (Figure 3) by integrating the frequencies within each of the nine subdomains in the 3×3 stencil. By definition, the PDF function in the expectation integral computation (2) must satisfy the following equality:

$$\iint_{\Omega} p(M, C_L) dM dC_L \approx \sum_{i=1}^n \sum_{k=1}^m p(M_i, C_{L_k}) \Delta M \Delta C_L = 1. \quad (4)$$

If we denote the total frequency in each of the subdomains centered at (M_i, C_{L_k}) by $f_H(M_i, C_{L_k})$, we have

$$\sum_{i=1}^3 \sum_{k=1}^3 \frac{f_H(M_i, C_{L_k})}{\sum_i \sum_k f_H(M_i, C_{L_k})} = \sum_{i=1}^3 \sum_{k=1}^3 \bar{f}_H(M_i, C_{L_k}) = 1, \quad (5)$$

where $\bar{f}_H(M_i, C_{L_k})$ is the relative frequency. The relation $\bar{f}_H(M_i, C_{L_k}) = p(M_i, C_{L_k}) \Delta M \Delta C_L$ can




Figure 5: Integration points for Case 5.

Mach	C_L	w_j
0.85	0.475	0.2
0.84	0.500	0.2
0.85	0.500	0.2
0.86	0.500	0.2
0.85	0.525	0.2

Table 4: Integration points and weights for Case 5

then be substituted into the two-dimensional midpoint integration rule, and $\mathbb{E}[C_D]$ can be approximated as

$$\mathbb{E}[C_D] \approx \sum_{i=1}^n \sum_{k=1}^m C_D(M_i, C_{L_k}) \bar{f}_H(M_i, C_{L_k}), \quad (6)$$

where $\bar{f}_H(M_i, C_{L_k})$ yields the values for each weight w_j . The computed weights are listed in Table 6 and represented in grayscale in Figure 7.

3.3.3. Five-Point Configuration with Flight-Condition Histogram

In addition to Case 9mh, we use the discrete histograms derived from the generated flight-condition distribution in a five-point configuration case. As shown in Table 6, the weights corresponding to the corner points are small (less than 0.1). We hypothesize that removing these points does not significantly change the optimization results. Using this five-point stencil, we hope to obtain results similar to those for the nine-point case at a fraction of the computational cost.

The distribution of the integration points for this case (5h) is similar to that for Case 5, as shown in Figure 8. The derivation of the weights is similar to that for Case 9mh, and the weights are again normalized such that their sum is 1. The weight values are represented in grayscale in Figure 8 and listed in Table 7.

3.3.4. Nine-Point Configuration with Gauss–Hermite Quadrature

In the cases presented above, the multipoint weights are products of the integration weights (based on the selected integration rule), τ_{ik} , and the PDF $p(M_i, C_{L_k})$. The integration points are selected manually




Figure 6: Integration points for Case 9mu.

Mach	C_L	w_j
0.84	0.475	1/9
0.85	0.475	1/9
0.86	0.475	1/9
0.84	0.500	1/9
0.85	0.500	1/9
0.86	0.500	1/9
0.84	0.525	1/9
0.85	0.525	1/9
0.86	0.525	1/9

Table 5: Integration points and weights for Case 9mu

and might not be the best possible points. In this section, we propose a Gauss–Hermite quadrature to automatically determine the integration points and corresponding weights.

Gauss–Hermite quadrature can be used to approximate the expectation integral via a change in variable when the integration variable is assumed to have a normal distribution [42]. The Gauss–Hermite quadrature is given by

$$\int_{-\infty}^{\infty} f(x) e^{-x^2} dx \approx \sum_{i=1}^n w_i f(\zeta_i), \quad (7)$$

where w_i and ζ_i denote each integration weight and point, respectively. The weighting function in the integration, e^{-x^2} , does not appear on the right-hand side. Instead, this term is captured by the w_i coefficients.

The expectation of a normally distributed variable, $\mathbb{E}[h(y)]$ with $y \sim \mathcal{N}(\mu, \sigma^2)$ (where μ and σ^2 refer to the normal mean and variance), can be computed as

$$\mathbb{E}[h(y)] = \int_{-\infty}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} h(y) \exp\left[-\frac{(y-\mu)^2}{2\sigma^2}\right] dy. \quad (8)$$

To use the Gauss–Hermite quadrature in this expectation integral approximation, we apply the variable transformation, $y = \sigma\sqrt{2}x + \mu$. With this transformation, Equation (8) becomes

$$\begin{aligned} \mathbb{E}[h(y)] &= \int_{-\infty}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} h(\sqrt{2}\sigma x + \mu) \exp(-x^2) \sigma\sqrt{2} dx \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{\pi}} h(\sqrt{2}\sigma x + \mu) \exp(-x^2) dx. \end{aligned} \quad (9)$$




Figure 7: Integration points for Case 9mh.

Mach	C_L	w_j
0.84	0.475	0.045
0.85	0.475	0.113
0.86	0.475	0.034
0.84	0.500	0.131
0.85	0.500	0.280
0.86	0.500	0.082
0.84	0.525	0.077
0.85	0.525	0.181
0.86	0.525	0.057

Table 6: Integration points and weights for Case 9mh

After finding n Gauss–Hermite quadrature weights and nodes using a numerical algorithm, we can compute $\mathbb{E}[h(y)]$ using the approximation

$$\mathbb{E}[h(y)] \approx \sum_{i=1}^n \frac{1}{\sqrt{\pi}} w_i^{\text{GH}} h(\sqrt{2}\sigma\zeta_i^{\text{GH}} + \mu), \quad (10)$$

where the superscript GH is used to denote “Gauss–Hermite.” We can decompose multivariate weighted Gaussian quadrature with zero correlation between the variables into a sequence of nested one-dimensional quadratures.

To use the Gauss–Hermite quadrature points and weights, we first need to determine a normal distribution $\mathcal{N}(\mu, \sigma^2)$ that approximates the flight-condition histogram. Assuming zero correlation between Mach and C_L , we computed the corresponding μ and σ^2 of the normal distributions as

$$\mathcal{N}_M(0.85, 0.007^2) \quad \text{and} \quad \mathcal{N}_{C_L}(0.50, 0.02^2). \quad (11)$$

This leads to the two-dimensional multivariate normal distribution shown in Figure 9. The integration points for the 3×3 Gauss–Hermite quadrature stencil are shown as white circles, sized by their weights. Although the selection of the integration points is automated in this approach, the Gauss–Hermite integration rule is suitable only for cases where the PDFs are well approximated by normal distributions.

We now analyze the convergence of the Gauss–Hermite quadrature scheme. We conduct the convergence study using the baseline configuration, and thus no optimization is involved here. For this two-dimensional




Figure 8: Integration points for Case 5h.

Mach	C_L	w_j
0.85	0.475	0.144
0.84	0.500	0.166
0.85	0.500	0.356
0.86	0.500	0.104
0.85	0.525	0.230

Table 7: Integration points and weights for Case 5h

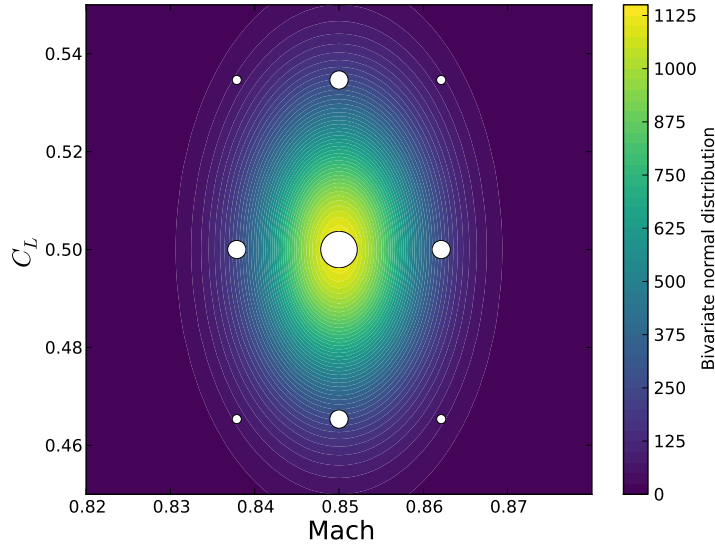


Figure 9: Normal distribution approximation of the flight-condition histogram and Gauss-Hermite quadrature points sized by the corresponding weights.

quadrature, the integration points are distributed in a uniform rectangular grid. The number of integration
 265 points ranges from 4 (2×2) to 2500 (50×50). Figure 10 shows the convergence of the relative error in

$\mathbb{E}[C_D]$ as the number of integration points is increased. The relative error is computed using the $\mathbb{E}[C_D]$ for the 2500-point integration as the reference value. The error rapidly reduces below $\mathcal{O}(10^{-4})$ and there is no significant variation beyond 500 points. With nine integration points, the error is $\mathcal{O}(10^{-2})$ (in drag counts), which is deemed sufficiently accurate for the multipoint optimization problem presented in this paper.




Figure 10: Error in approximating $\mathbb{E}[C_D]$ for different numbers of Gauss-Hermite quadrature points, in drag counts ($C_D \times 10^4$).

4. Results

In this section, we present the results for the optimization cases considered. We seek to quantify the difference in aircraft performance as well as how the different multipoint formulations affect the approximations of the expected C_D values. We also want to analyze the optimization history and quantify the computational cost for each case.

4.1. Optimization Results

The key features of the most representative results are summarized in Figures 11 to 14. For the sake of conciseness, only four of the seven cases are presented here: cases 1, 5h, 9mh, and 9tu-ADODG. In these plots all the C_D values are reported as drag counts ($C_{D_{\text{counts}}} = C_D \times 10^4$). The red and blue lines indicate the baseline and optimized configurations, respectively. On the top right we show the planform view of the wing with the C_p contours of the top surface. The left wing corresponds to the baseline design, which

is the design that the optimization starts from, and the right wing corresponds to the optimized design. A front view with the shock surface visualization and the physical thickness distribution of the wing is shown below the planform view. The next three plots on the left show the spanwise lift (with the reference elliptical distribution shown in gray), twist, and thickness-to-chord ratio (t/c) distributions. The lift distributions are
285 normalized such that the area under the curve equals 1. On the right, we show the airfoils at six spanwise locations (A through F) and the corresponding C_p distributions. For the nine-point cases (Figures 13 and 14), the C_p distributions of all nine flight conditions are shown. The thick blue lines correspond to the nominal flight condition ($M = 0.85$ and $C_L = 0.50$).

The optimized configurations exhibit some common trends. The shock surfaces are visibly reduced in
290 all the optimized results and almost completely eliminated in the single-point case. All the optimized lift distributions are closer to the reference elliptical distribution than the baseline one, which results in lower induced drag for the optimized configurations. The twist values close to the wing root and tip are similar to the values for the baseline wing, but the optimized wings exhibit a short plateau in the mid-span section.

The optimized t/c distributions show a significant increase near the root and a reduction closer to the tip
295 when compared to the baseline design, especially at sections C, D, E, and F. This change in the t/c distribution is due to a trade-off in the spanwise distribution of the viscous pressure drag. Since the optimization is allowed to reduce the thickness to 25% of the baseline, it reduces the outboard airfoil thickness to decrease the viscous pressure drag, while increasing the inboard thickness to satisfy the volume constraint, where the chord is much larger. This design trade-off is more pronounced in the single-point case and is consistent
300 with previous results for similar cases [31, 35]. The change in the spanwise viscous pressure drag distribution has been verified by Dumont and Meheut [10], who used a drag decomposition technique that is able to separate the contributions from wave drag, induced drag, and viscous pressure drag. The leading edge of the outboard wing also exhibits a smaller radius, and thus these designs would have a poor low-speed $C_{L_{\max}}$ performance, which is not considered in this study. This trend is less pronounced in the multipoint cases,
305 but the radius is still small enough that it would hinder low-speed performance.

These results suggest that both structural constraints and low-speed performance considerations are required. Thus, the optimized wings obtained herein have an unfair advantage over the baseline CRM wing, which was designed with these considerations in mind. However, our main goal is not to show that our optimized wings are better than the baseline but to study the effect of the cruise flight conditions and weights.
310 Therefore, we focus on comparing the optimized designs to determine which multipoint formulation should be used.

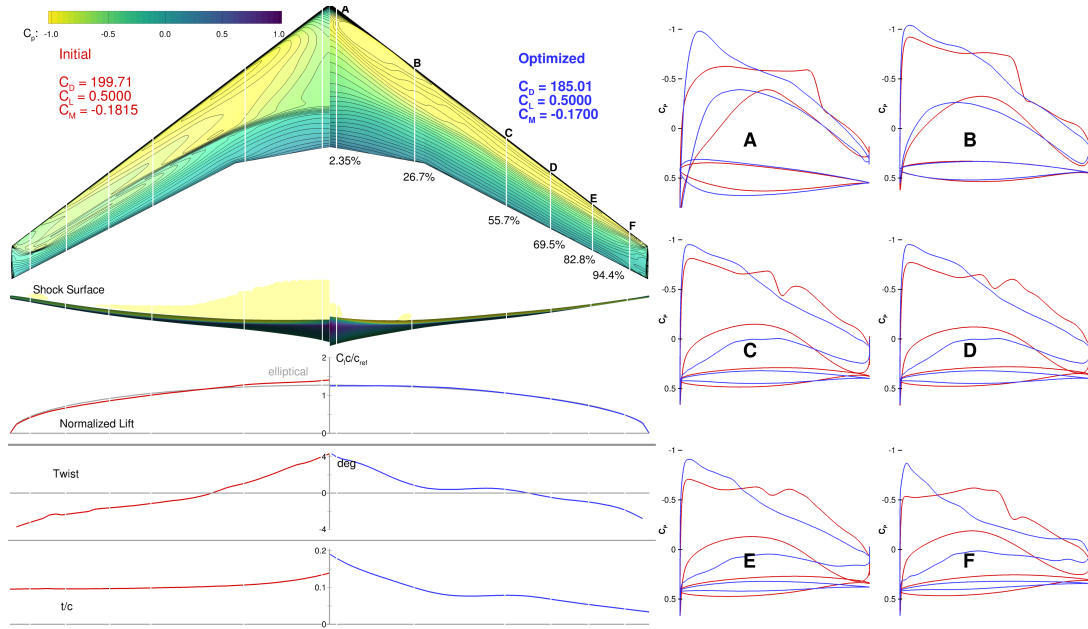


Figure 11: Case 1 optimized wing (blue) compared to baseline wing (red).

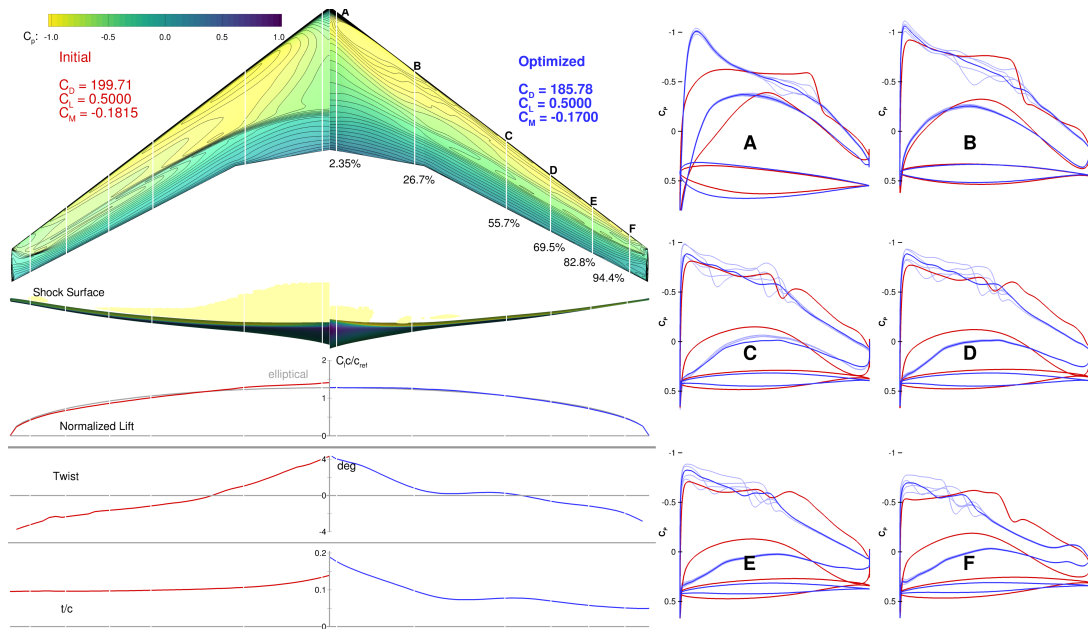


Figure 12: Case 5h optimized wing (blue) compared to baseline wing (red).

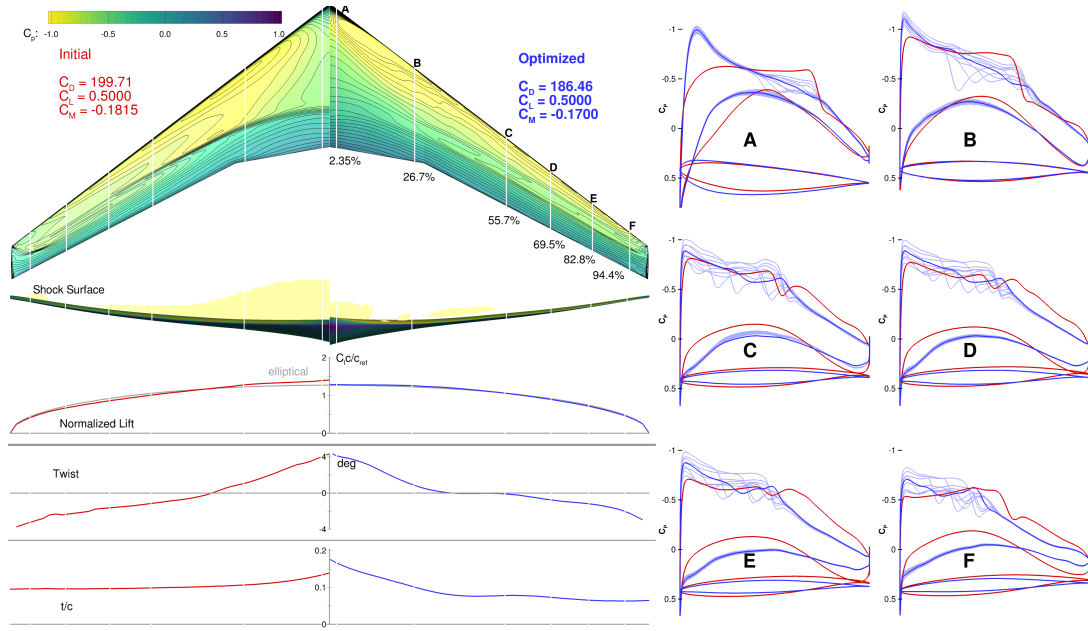


Figure 13: Case 9mh optimized wing (blue) compared to baseline wing (red).

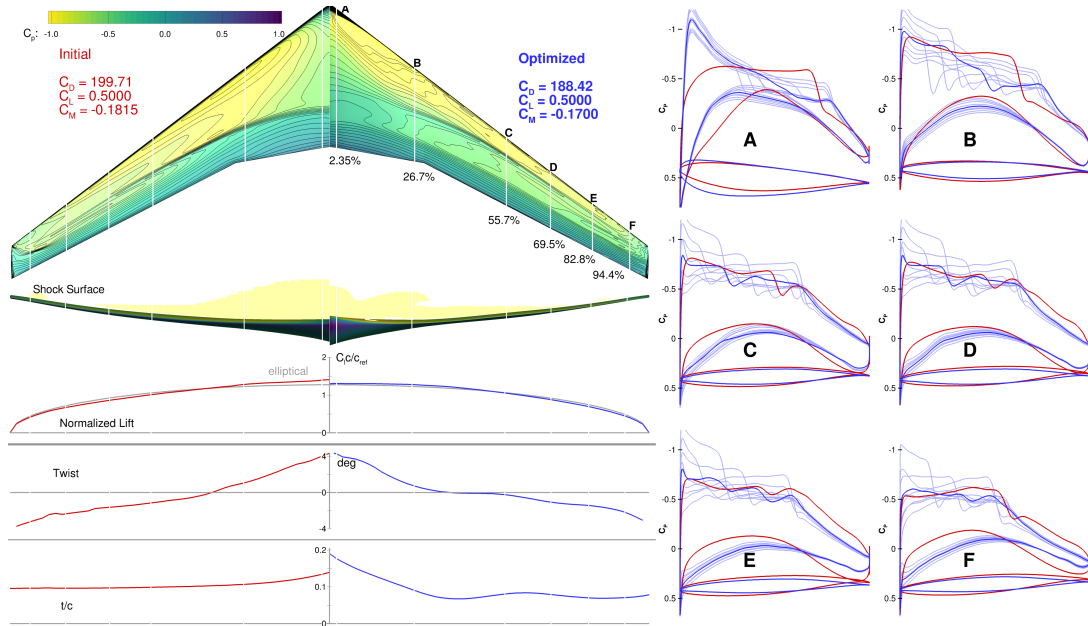


Figure 14: Case 9tu-ADODG optimized wing (blue) compared to baseline wing (red).

Figure 15 plots the drag divergence profiles for the baseline and optimization cases, for three C_L values: 0.475, 0.5, and 0.55. These plots highlight the more consistent drag reduction achieved when including more points in the formulation. The optimizer tries to reduce drag for all the flight conditions considered in the formulation. Therefore, the larger the flight-condition range covered by the integration points, the flatter the drag divergence curve. As shown in Figure 15, the flattest drag divergence profile is observed in Case 9tu-ADODG, where the integration points cover the largest range (as shown in Figure 4). The drag reduction for Case 9tu-ADODG also extends to higher Mach values. The drag reduction in the single-point optimization (Case 1), on the other hand, occurs only near the nominal condition, with higher drag values observed at conditions away from the nominal.

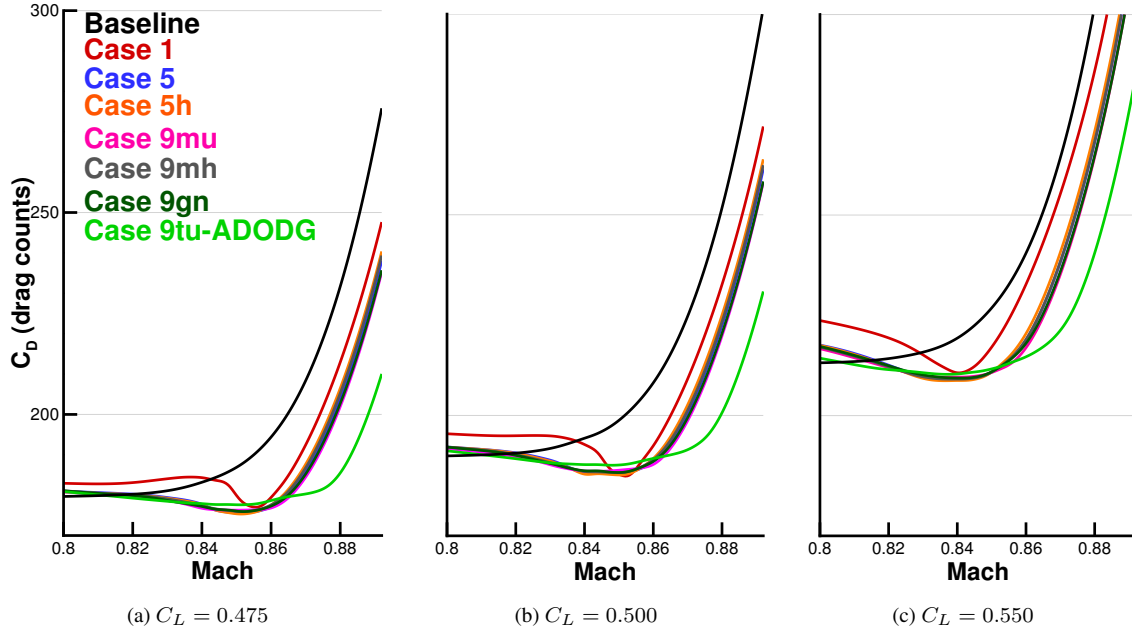


Figure 15: Comparison of drag divergence for three different lift coefficients.

Although the multipoint optimization results are more robust, as demonstrated by the more consistent drag reduction in the flight-condition space, they sacrifice some performance at the nominal condition. Table 8 summarizes the optimized C_D values at the nominal flight condition and their relative reductions (computed with respect to the baseline C_D of 199.71 counts). The reductions range from 5.65% to 7.36%. As previously discussed (in Section 2.3), allowing for thinner wings accounts for these large reductions rel-

ative to the baseline wing. The single-point optimization achieved the largest drag reduction (7.36%), which is expected because it does not need to compromise across multiple flight conditions. For the multipoint formulations with the same range in the flight conditions (cases 5 through 9gn) the difference in the drag reduction was within 0.43%. This suggests that different multipoint formulations do not significantly affect drag reductions provided the points span the desired range in cruise flight conditions.

Case	C_D	% Reduction
Baseline	199.71	—
Case 1	185.01	7.36%
Case 5	186.32	6.70%
Case 5h	185.78	6.98%
Case 9mu	186.46	6.63%
Case 9mh	186.63	6.55%
Case 9gn	186.43	6.65%
Case 9tu-ADODG	188.42	5.65%

Table 8: Drag coefficients and their relative reductions for the optimized configurations (with respect to the baseline configuration) at the nominal condition (in drag counts, $C_D \times 10^4$).

To further demonstrate the merit of multipoint versus single-point optimization, we compare the \sqrt{ML}/D contours, which can be interpreted as a two-dimensional extension of the drag-divergence curve. The \sqrt{ML}/D metric can be used to assess the range performance of a transonic commercial aircraft [19], which is based on the Breguet range equation

$$R = \frac{L}{D} \frac{V}{c_T} \ln \left(\frac{W_1}{W_2} \right), \quad (12)$$

where L/D is the lift-to-drag ratio, V is the velocity, c_T is the TSFC, W_1 is the initial cruise weight, and W_2 is the final cruise weight. Assuming that the TSFC is proportional to \sqrt{M} (which is reasonable for high-bypass ratio turbofans operating in the troposphere [28]), and that the speed of sound and the W_1/W_2 ratio are constant, the range is proportional to \sqrt{ML}/D .

Figure 16 shows the \sqrt{ML}/D contours for the baseline and single-point optimized geometries. The drag values must be corrected to include the additional parasitic drag from components other than the aircraft wing. Therefore, an additional 28 drag counts are added, such that the maximum \sqrt{ML}/D of the

baseline configuration occurs at the nominal $C_L = 0.5$ [22], as shown in Figure 16a. The thick red line marks the contour where $\sqrt{ML}/D = 20.53$, which is 99% of its maximum value 20.73. Aircraft typically operate within this enclosed area, within which the fuel-burn penalty would not exceed 1%. Therefore, this 99%(\sqrt{ML}/D)_{max} area can be used as a proxy to measure the robustness of the design. For the single-point case (Figure 16b), the blue circle indicates the maximum \sqrt{ML}/D value, and the corresponding 99%(\sqrt{ML}/D)_{max} contours are shown as thick blue lines. We refer to the area enclosed by the red line as the *baseline high-performance region* and one enclosed by the blue line as the *high-performance region*.

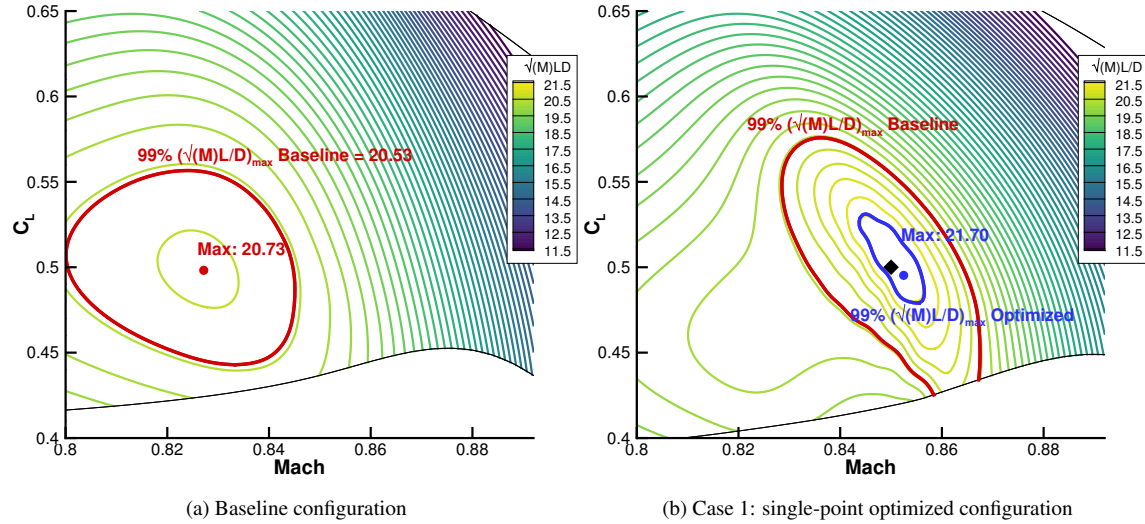


Figure 16: \sqrt{ML}/D contours for the baseline and single-point optimized configurations.

Figure 17 displays the \sqrt{ML}/D contours for all the multipoint cases. The integration points are shown as black diamonds, sized by their objective-function weight values. From these contour plots, we observe that the maximum \sqrt{ML}/D (i.e., the best-performance point) occurs at a flight condition close to the nominal point. Recall that the integration-point configurations for all the cases are centered at the nominal flight condition.

Figure 18 shows the superposition in one plot of the high-performance regions for the baseline and the optimized configurations. The points corresponding to the maximum \sqrt{ML}/D of all the optimization cases are displayed, and the values are shown inside the brackets. From this comparison, we identify two main factors that affect the shape and size of the high-performance region of each optimization case. The first

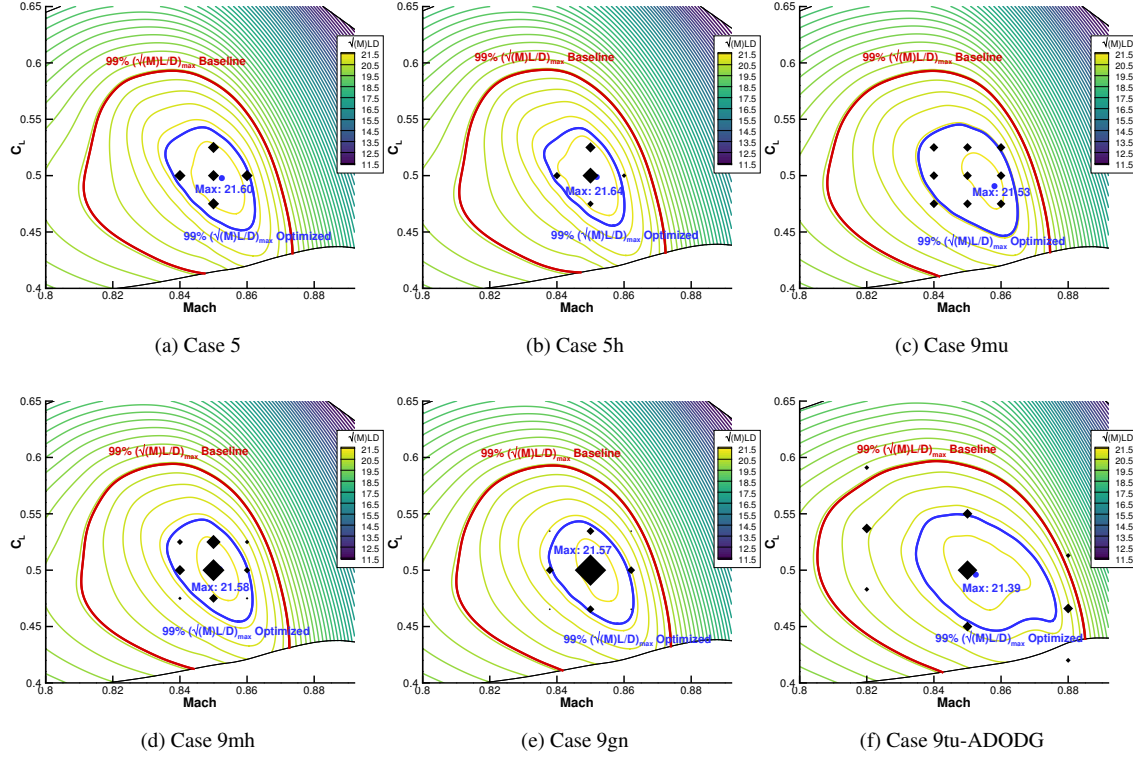


Figure 17: \sqrt{ML}/D contours for the multipoint-optimized designs.

factor is the flight-condition range covered by the integration points, which affects the flatness of the drag-divergence curves, as previously discussed. The second factor has to do with the weight values assigned at the corner points. The higher these weights (e.g., when assigning equal weights for all points), the larger the high-performance region. This is because the corner points are assigned the same importance as the nominal condition, and thus the optimizer trades off the performance at the nominal condition with the corner points. This plot shows that all the optimized configurations have higher operating Mach numbers than the baseline configuration, which means that the aircraft can fly faster without incurring a drag penalty.

Figure 18 expands on the drag divergence curve comparison of Figure 15. All the optimized configurations demonstrate improved performance compared to the baseline design, and we also observe the improved robustness of the multipoint results compared with the single-point result. The improved performance is demonstrated by the higher operating Mach number as well as the higher maximum \sqrt{ML}/D (i.e., the

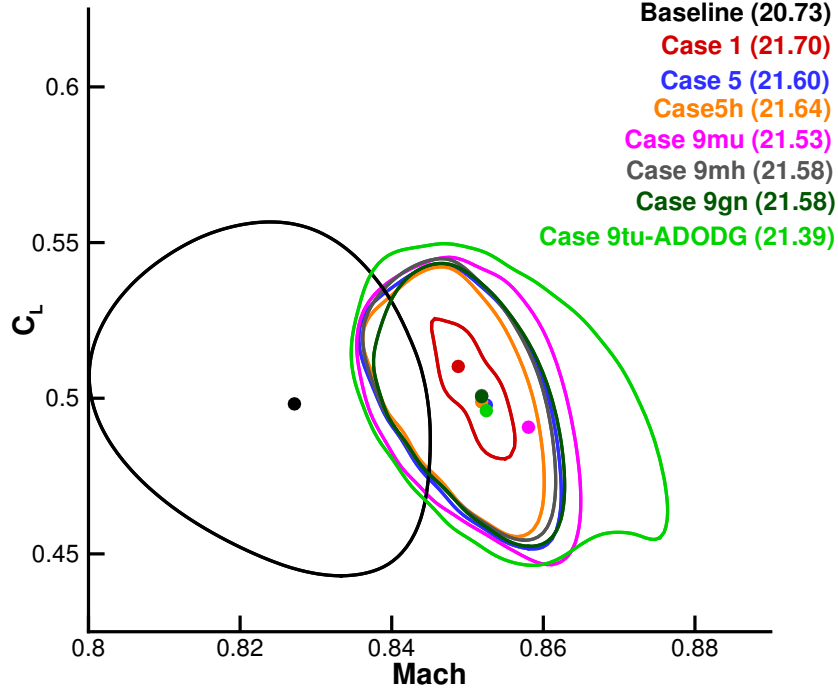


Figure 18: High-performance regions for the baseline and optimized configurations.

best-performance point). The improved robustness is shown by the larger high-performance regions of the multipoint results. In other words, the multipoint optimizations have “flatter” high-performance regions, whereas the single-point optimization has a sharper maximum. All the high-performance regions for the optimized configurations are smaller than that of the baseline configuration. This outcome shows that the optimizer tries to improve the performance around the flight conditions that truly matter, based on the information derived from the flight operation data.

4.2. Assessment of the Expected Drag Approximation

In this section, we assess the objective function accuracy in estimating the expected value of the drag coefficient, $\mathbb{E}[C_D]$. We denote the objective function as \bar{C}_D to represent the weighted average C_D . Since the actual $\mathbb{E}[C_D]$ is unknown, we use the flight-condition histogram shown in Figure 3 to compute the reference value, which is denoted $\mathbb{E}[C_D]_{\text{ref}}$. A midpoint integration rule is employed, where the flight condition at the center of each cell is used as the integration point. The flight-condition frequency in each cell (mapped

to the colors in the histogram) is used as the PDF in the expectation integral computation. The difference between $\mathbb{E}[C_D]_{\text{ref}}$ and \bar{C}_D is denoted ΔC_D ; all the quantities are reported in drag counts.

380 Table 9 summarizes the assessment results for the baseline and optimized configurations. For the baseline configuration, $\mathbb{E}[C_D]_{\text{ref}}$ is the same for every case (201.56), while the \bar{C}_D computations are different because of the different integration points and weights. The computed C_D at the nominal condition for each configuration is also shown in Table 9.

Case	Baseline configuration		Optimized configuration			
	\bar{C}_D	ΔC_D	Nominal	$\mathbb{E}[C_D]_{\text{ref}}$	\bar{C}_D	ΔC_D
Baseline	—	—	199.71	—	—	—
1	198.77	−2.79	185.01	191.18	185.01	−6.17
5	200.05	−1.51	186.32	188.12	186.43	−1.69
5h	200.37	−1.19	185.78	188.07	186.82	−1.24
9mu	200.80	−0.75	187.14	188.34	187.26	−1.08
9mh	201.14	−0.41	186.46	188.37	188.13	−0.25
9gn	200.37	−1.19	186.43	188.26	187.13	−1.13
9tu-ADODG	210.27	8.71	188.42	189.47	192.23	2.76

Table 9: Comparison of $\mathbb{E}[C_D]_{\text{ref}}$ and \bar{C}_D for the baseline and optimized configurations. For the baseline configuration, the same $\mathbb{E}[C_D]_{\text{ref}}$ reference value (201.56) is used for all the cases. All the values are in drag counts ($C_D \times 10^4$).

From Table 9, we see that Case 9mh yields the smallest ΔC_D for both the baseline and optimized configurations. This result is expected since this case derived the PDF directly from the flight-condition histogram. In general, including more points in the expectation integral approximation improves the accuracy. For the optimized configurations, the single-point approximation (Case 1) is the least accurate, since a single point is not sufficient to represent the entire integration domain. Another case with poor approximation accuracy is Case 9tu-ADODG, since it includes points outside the flight-condition histogram (Figure 3). This emphasizes the importance of ensuring that the selection of the integration points and corresponding weights reflects how the aircraft actually operates.

In addition to assessing the accuracy of the expectation integral approximation, we compare the overall expected C_D for the different optimization cases, i.e., the computed values of $\mathbb{E}[C_D]_{\text{ref}}$. Table 9 shows that

all the multipoint cases yield similar $\mathbb{E}[C_D]_{\text{ref}}$ values. To investigate this further, we analyze the C_D variation in the flight-condition space for each optimized configuration. For easier comparison, the drag values are compared to the most accurate integration, which is Case 9mh. Figure 19 shows these comparisons for Case 1 and Case 9tu-ADODG, which have the biggest discrepancies.




Figure 19: C_D deviations from the reference case (Case 9mh) for each flight condition (in drag counts).

For the other four cases, the differences are much smaller and thus a much narrower value range is used in the plots shown in Figure 20. Figures 19 and 20 show that the C_D variations for the optimized configurations fall within a narrow range. The C_D values for the two extreme cases, Case 1 and Case 9tu-ADODG, differ from the reference (Case 9mh) by -15 to 15 drag counts, whereas in the other cases the ranges are between -2 and 2 drag counts.

The above observations suggest that each case yields different configurations with different performance, but the differences are small. Moreover, the differences at the perimeter of the flight-condition histogram become even more insignificant in calculating the expected drag because they are less frequent flight conditions. Therefore, the resulting $\mathbb{E}[C_D]_{\text{ref}}$ values are within 1 drag count of each other, as shown in Table 9.

In Figures 11–14 we show that the various multipoint cases result in similar geometries. In this comparison, we again use Case 9mh as the reference case, and compute the root-mean-square deviation (RMSD) between any airfoil geometry and this reference airfoil. This deviation is based on the Euclidean distance between each point of a given airfoil and the corresponding point on the reference airfoil, i.e.,

$$d_i = \sqrt{(x_i - x_i^{\text{ref}})^2 + (z_i - z_i^{\text{ref}})^2}, \quad i = 1, 2, \dots, N_{\text{points}}. \quad (13)$$




Figure 20: C_D deviations from the reference case (9mh) for each flight condition (in drag counts).

The RMSD is

$$\text{RMSD} = \sqrt{\frac{\sum_{i=1}^N d_i^2}{N}}, \quad (14)$$

where N is the number of points defining each airfoil. In addition to computing the RMSD for each airfoil, we also compute the overall RMSD for each case using the geometry coordinates of six airfoils. The airfoil geometry comparison is summarized in Table 10.

Case	A	B	C	D	E	F	Overall
Case 1	2.8822	4.0880	3.1423	3.4589	3.6993	3.8524	3.5440
Case 5	0.0204	0.0134	0.0242	0.0125	0.0163	0.0230	0.0189
Case 5h	0.0287	0.0140	0.0313	0.0182	0.0269	0.0381	0.0274
Case 9mu	0.0203	0.0155	0.0300	0.0112	0.0109	0.0185	0.0189
Case 9gn	0.0214	0.0149	0.0164	0.0171	0.0131	0.0140	0.0164
Case 9tu-ADODG	0.0394	0.0365	0.0181	0.0284	0.0220	0.0260	0.0294

Table 10: Quantification of the airfoil geometry differences, using Case 9mh as the reference geometry (RMSD in meters)

The results in Table 10 show that Case 1 has the highest discrepancy with the reference case, which is consistent with the $\mathbb{E}[C_D]_{\text{ref}}$ comparison results presented in Table 9. The overall deviation for this case is around 3.5 m. This deviation is illustrated in Figure 21 for six airfoil slices; it is lowest near the wing root and increases toward the wing tip. For the other cases, the deviations are less than 3 cm. Compared to the overall wing dimensions (the wing mean aerodynamic chord is 7 m), these deviations are small, which explains the similarities in the $\mathbb{E}[C_D]_{\text{ref}}$ values.

4.3. Optimization History and Computational Cost

Figure 22 shows the evolution of the merit function and the optimality of four representative cases. These plots are representative of all the cases, since they all exhibit similar trends. In SNOPT, the merit function is used in the augmented Lagrangian formulation, and optimality refers to how closely the current point satisfies the first-order Karush–Kuhn–Tucker conditions [13]. For all the cases, the optimality and feasibility tolerances are set to 10^{-4} and 10^{-5} , respectively. In Figure 22, we see that the nine-point cases require 100 to 110 iterations to converge, the five-point case requires about 140 iterations, and the single-point case requires over 160 iterations. Based on Figures 11–14 and the C_D reductions summarized in Table 8, we see

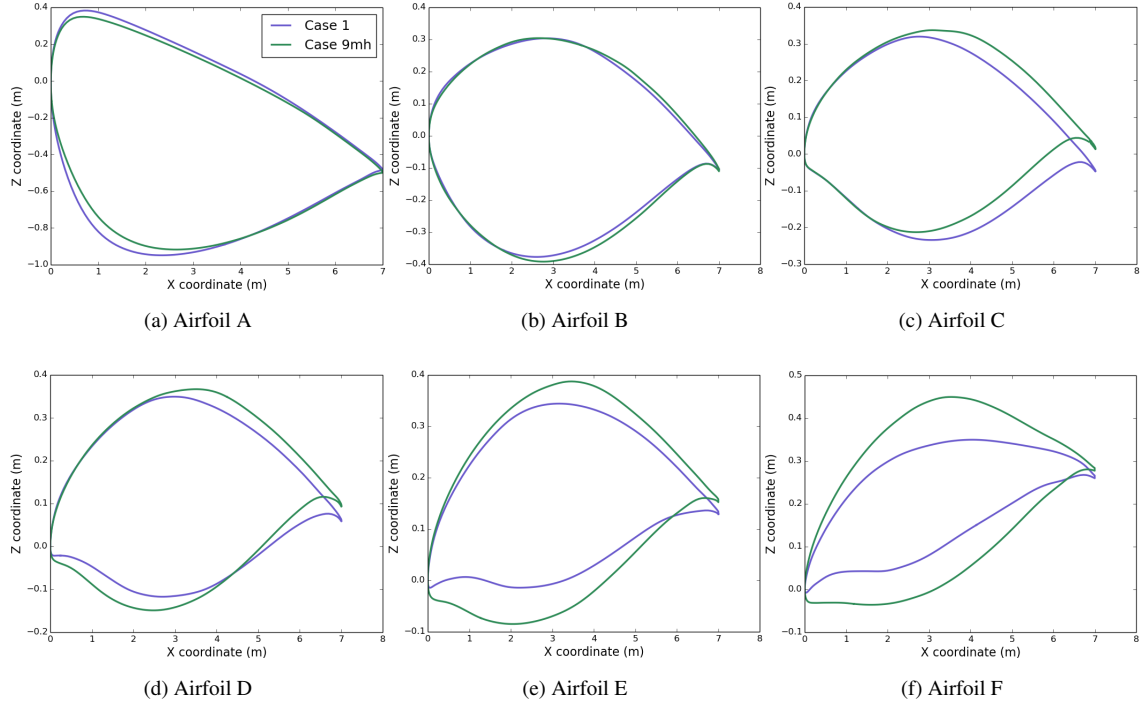


Figure 21: Airfoil geometry comparison for Cases 1 and 9mh.

that the single-point optimization results in the largest change relative to the baseline configuration in both drag and shape, which explains the increased number of iterations.

425 All the optimized designs satisfy the lift, moment, volume, and thickness constraints within the specified tolerance. All the cases are run in parallel [27]: the L2 cases are performed with 32 processors, while the L1 cases use 480 processors. Table 11 summarizes the total number of major iterations and computation time (in hours) for all the cases. The coarse optimizations, where most of the gains are achieved, range from 31 to 59 hours of wall time for the multipoint cases using 32 processors, so they could be completed overnight
430 using 128 processors.

5. Conclusion

Most multipoint aerodynamic shape optimization investigations to date have relied on heuristics to choose the flight conditions and the corresponding weights. In this work, we investigate the effects of

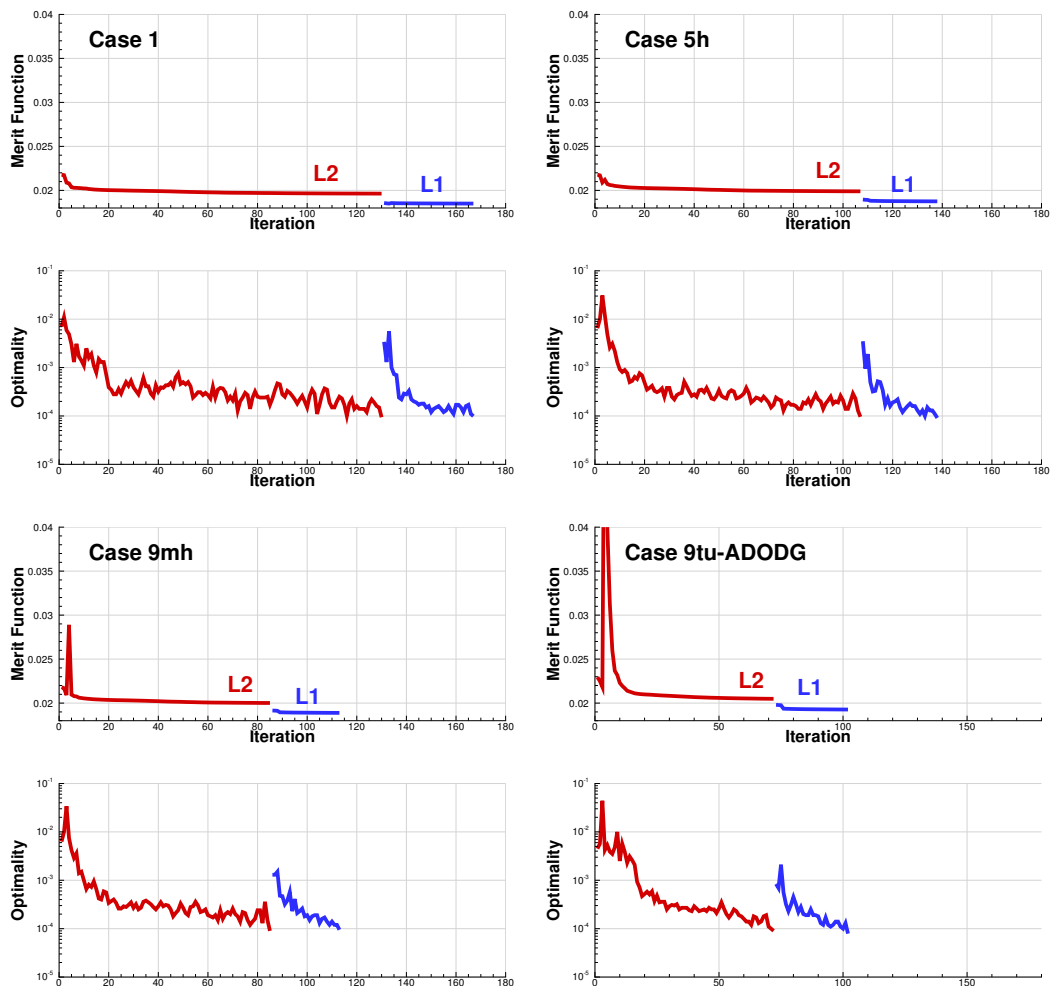


Figure 22: Objective function and optimization history for four cases.

Case	Coarse optimization (32 procs)		Fine optimization (480 procs)	
	Major iterations	Time (h)	Major iterations	Time (h)
Case 1	129	9	36	7
Case 5	94	31	23	9
Case 5h	106	34	30	12
Case 9mu	81	52	23	15
Case 9mh	83	52	27	18
Case 9gn	97	59	29	18
Case 9tu-ADODG	71	48	29	19

Table 11: Number of major iterations and total computational time for the coarse optimizations (with 32 processors) and fine optimizations (with 480 processors)

different multipoint optimization formulations on the expected drag minimization. In particular, we aim to
435 reduce drag over all the flight conditions that the aircraft normally operates at during cruise. To this end,
we use actual mission operational data to generate a PDF of the flight conditions, from which a series of
multipoint objective functions is formulated. We also perform a single-point optimization for comparison
purposes. All the optimizations are for the CRM wing, subject to lift, moment, volume, and thickness con-
straints. There are 768 design variables and 752 constraints. The wing planform shape is fixed, and the
440 aerodynamic analysis consists in solving the RANS equations.

We formulate the expected drag minimization problem as a multipoint problem by computing the ex-
pectation integral numerically, employing midpoint and Gauss–Hermite quadrature rules. We perform the
integration in the two-dimensional space of Mach and C_L . We obtain the flight-condition distribution by
performing detailed mission analyses based on actual flight data. We then use this distribution to select
445 the integration points and weights for all the multipoint cases except one, which is a previously defined
benchmark case (9tu-ADODG). In addition, we approximate this distribution via a normal distribution to
automatically generate integration points and weights using the Gauss–Hermite quadrature.

All the multipoint cases result in more consistent performance gains than the single-point case, especially
around the high-frequency region of the flight-condition distribution. This results in flatter drag divergence
450 curves and larger high-performance regions in the \sqrt{ML}/D plots. The single-point optimization, on the

other hand, resulted in localized performance gain around the nominal flight condition with the highest drag reduction at this nominal point.

Comparing the various multipoint results suggests that it is important to choose the right points, but the determination of the integration weights is secondary if the focus is on minimizing the overall expected drag. This conclusion is based on the fact that all the multipoint cases (except 9tu-ADODG) yield similar designs. Recall that Case 9tu-ADODG is the only multipoint case that did not use the derived flight-condition histogram to determine the flight conditions and weights. The cases that are based on the flight-condition distribution (5, 5h, 9mu, 9mh, 9gn) have similar integration point stencils; the only difference between the five-point and nine-point stencils is the absence of corner points. Each case uses a different rationale for determining the flight-condition weights.

Our results also highlight the importance of distinguishing between flight conditions where we want to optimize the performance, and conditions where we want only to ensure adequate handling for safety. Case 9tu-ADODG, for example, includes flight conditions at which the aircraft would not normally operate during cruise. These conditions are included to emulate the off-design conditions to provide a margin between cruise conditions and buffet onset. This approach to addressing buffet onset reduced the performance at the flight conditions relevant for normal operation. Thus, a better approach would be to consider only the relevant flight conditions in the objective and to include the off-design conditions as constraints [23].

Although most of the multipoint designs are similar, the expected C_D values exhibit different accuracies when compared to the actual drag integral. The most accurate is Case 9mh, where the weights are chosen to integrate the flight histogram, with a 0.25 drag count discrepancy. For our aerodynamic shape optimizations, an accurate $\mathbb{E}[C_D]$ approximation at each iteration is not critical, since we achieve similar results. However, for aerostructural optimization [21], objective function accuracy is critical because we need to accurately quantify the trade-off between aerodynamics and structural disciplines at each optimization iteration.

We conclude that a five-point configuration with equal weights is sufficient in the multipoint formulation if the focus is on the final optimized shape and only aerodynamics is considered. On the other hand, if the expectation integral accuracy is important at each iteration, e.g., when doing trade-off analysis in the loop, then the multipoint formulation should be based on the actual flight-condition distribution. While this conclusion might not be applicable to flight-condition distributions that are different from those considered here, our approach can still be used to minimize the expected drag or another objective, such as fuel burn.

480 Acknowledgments

The first author is grateful for the funding provided by the Vanier Canada Graduate Scholarships. The computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; the Ontario Research Fund—Research Excellence; and the University of Toronto. The authors would
485 like to recognize the other members of their research group, especially Zhoujie (Peter) Lyu, for their contributions to the solvers and framework.

References

- [1] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 9th printing*. Denver, New York, 1972.
- 490 [2] W. Kyle Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers and Fluids*, 28(4):443–480, 1999. doi:10.1016/S0045-7930(98)00041-3.
- [3] Howard P. Buckley and David W. Zingg. Approach to Aerodynamic Design Through Numerical Optimization. *AIAA Journal*, 51(8):1972–1981, August 2013. doi:10.2514/1.J052268.
- 495 [4] Susan E. Cliff, James J. Reuther, David A. Saunders, and Raymond M. Hicks. Single-Point and Multipoint Aerodynamic Shape Optimization of High-Speed Civil Transport. *Journal of Aircraft*, 38(6):997–1005, 2001. doi:10.2514/2.2886.
- [5] G. B. Constantino and T. L. Holst. Numerical Optimization Design of Advanced Transonic Wing Configurations. Technical report, NASA TM 85950, 1984.
- 500 [6] Germund Dahlquist and Åke Björck. *Numerical Methods in Scientific Computing: Volume 1, pages 521–604*. SIAM e-books. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 2008.
- [7] Germund Dahlquist and Åke Björck. *Numerical Methods, pages 290–304*. Dover Books on Mathematics. Dover Publications, 2012.

- 505 [8] Philip J. Davis and Philip Rabinowitz. *Methods of Numerical Integration*. Dover Books on Mathematics Series. Dover Publications, 2007. URL <http://books.google.ca/books?id=gGCKdqka0HAC>.
- [9] Mark Drela. Pros and Cons of Airfoil Optimization. In D. A. Caughey and M. M. Hafez, editors, *Frontiers of CFD 1998*, pages 363–381. World Scientific, 1998.
- 510 [10] Antoine Dumont and Michael Meheut. Gradient-based optimization of CRM wing-alone and wing-body-tail configurations by RANS adjoint technique. In *54th AIAA Aerospace Sciences Meeting*. American Institute of Aeronautics and Astronautics (AIAA), January 2016. doi:10.2514/6.2016-1293.
- [11] Jonathan Elliott and Jaume Peraire. Constrained, Multipoint Shape Optimization for Complex 3D Configurations. *Aeronautical Journal*, 102(1017):365–376, 1998.
- 515 [12] François Gallard, Matthieu Meaux, Marc Montagnac, and Bijan Mohammadi. Aerodynamic aircraft design for mission performance by multipoint optimization. In *21st AIAA Computational Fluid Dynamics Conference*. American Institute of Aeronautics and Astronautics, Jun 2013. doi:10.2514/6.2013-2582.
- [13] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP Algorithm for Large Scale
520 Constrained Optimization. *SIAM Review*, 47(1):99–131, 2005. doi:10.1137/S0036144504446096.
- [14] R. M. Hicks and P. A. Henne. Wing Design by Numerical Optimization. *Journal of Aircraft*, 15: 407–412, July 1978.
- [15] R. M. Hicks, E. M. Murman, and G. N. Vanderplaats. An assessment of airfoil design by numerical optimization. Technical report, NASA TM X-3092, July 1974.
- 525 [16] A. Jameson. Aerodynamic Design via Control Theory. *Journal of Scientific Computing*, 3(3):233–260, 1988.
- [17] A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes. *AIAA Paper* 81-1259, 1981.
- [18] Antony Jameson. Automatic Design of Transonic Airfoils to Reduce the Shock Induced Pressure
530 Drag. In *Proceedings of the 31st Israel Annual Conference on Aviation and Aeronautics*, pages 5–17, Technion-Israel, Haifa, Israel, 1990.

- [19] Antony Jameson, John C. Vassberg, and Sriram Shankaran. Aerodynamic-Structural Design Studies of Low-Sweep Transonic Wings. In *46th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, 7–10 January 2008. AIAA 2008-145.
- 535 [20] Gaetan K. Kenway and Joaquim R. R. A. Martins. Aerodynamic shape optimization of the CRM configuration including buffet-onset conditions. In *54th AIAA Aerospace Sciences Meeting*, San Diego, CA, January 2016. American Institute of Aeronautics and Astronautics. doi:10.2514/6.2016-1294.
- [21] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multi-point High-fidelity Aerostructural Optimization of a Transport Aircraft Configuration. *Journal of Aircraft*, 51:144–160, 2014.
540 doi:10.2514/1.C032150.
- [22] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. *AIAA Journal*, 54(1):113–128, January 2016. doi:10.2514/1.J054154.
- [23] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Buffet onset constraint formulation for aerodynamic shape optimization. *AIAA Journal*, 2017. (In press).
545
- [24] Christopher Lee, David Koo, Karla Telidetzki, Howard Buckley, Hugo Gagnon, and David W. Zingg. Aerodynamic shape optimization of benchmark problems using jetstream. In *53rd AIAA Aerospace Sciences Meeting*, Kissimmee, FL, January 2015.
- [25] Rhea P. Liem, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Multimission Aircraft Fuel
550 Burn Minimization via Multipoint Aerostructural Optimization. *AIAA Journal*, 53(1):104–122, 2015. doi:10.2514/1.J052940.
- [26] Rhea P. Liem, Charles A. Mader, and Joaquim R. R. A. Martins. Surrogate models and mixtures of experts in aerodynamic performance prediction for mission analysis. *Aerospace Science and Technology*, 43:126–151, 2015. doi:10.1016/j.ast.2015.02.019.
- 555 [27] Chris Loken, Daniel Gruner, Leslie Groer, Richard Peltier, Neil Bunn, Michael Craig, Teresa Henriques, Jillian Dempsey, Ching-Hsing Yu, Joseph Chen, L. Jonathan Dursi, Jason Chong, Scott Northrup, Jaime Pinto, Neil Knecht, and Ramses Van Zon. SciNet: Lessons learned from building a power-efficient top-20 system and data centre. *Journal of Physics: Conference Series*, 256(1):012026, 2010. doi:10.1088/1742-6596/256/1/012026.

- 560 [28] Zhoujie Lyu. *High-Fidelity Aerodynamic Design Optimization of Aircraft Configurations*. PhD thesis, University of Michigan, 2014.
- [29] Zhoujie Lyu, Gaetan K. W. Kenway, Cody Paige, and Joaquim R. R. A. Martins. Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, Jul 2013.
- 565 [30] Zhoujie Lyu, Zelu Xu, and Joaquim R. R. A. Martins. Benchmarking optimization algorithms for wing aerodynamic design optimization. In *Proceedings of the 8th International Conference on Computational Fluid Dynamics*, Chengdu, Sichuan, China, July 2014. ICCFD8-2014-0203.
- [31] Zhoujie Lyu, Gaetan K. Kenway, and Joaquim R. R. A. Martins. Aerodynamic shape optimization studies on the Common Research Model wing benchmark. *AIAA Journal*, 53(4):968–985, April 2015. doi:10.2514/1.J053318.
- 570 [32] Michael Méheut, Daniel Destarac, Gérald Carrier, George Anderson, Siva Nadarajah, Daniel Poole, John Vassberg, and David W. Zingg. Gradient-Based Single and Multi-point Aerodynamic Optimizations with the *elsA* Software. In *Proceedings of the 53rd AIAA Aerospace Sciences Meeting*, Kissimmee, FL, January 2015.
- 575 [33] Siva Nadarajah and Antony Jameson. A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization. In *Proceedings of the 38th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2000. AIAA 2000-0667.
- [34] M. Nemec, D. W. Zingg, and T. H. Pulliam. Multipoint and Multi-Objective Aerodynamic Shape Optimization. *AIAA Journal*, 42(6):1057–1065, June 2004.
- 580 [35] Lana Maria Osusky. *A Novel Numerical Tool for Aerodynamic Shape Optimization in Turbulent Flow*. PhD thesis, University of Toronto, 2013.
- [36] Ruben E. Perez, Peter W. Jansen, and Joaquim R. R. A. Martins. pyOpt: a Python-based object-oriented framework for nonlinear constrained optimization. *Structural and Multidisciplinary Optimization*, 45(1):101–118, 2012. doi:10.1007/s00158-011-0666-3.
- 585 [37] J. Reuther, A. Jameson, J. Farmer, L. Martinelli, and D. Saunders. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. In *Proceedings of the 34th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, January 1996. AIAA 1996-0094.

- [38] James J. Reuther, Antony Jameson, Juan J. Alonso, Mark J. Rimlinger, and David Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 2. *Journal of Aircraft*, 36(1):61–74, 1999.
- [39] J. A. Samareh. Aerodynamic Shape Optimization Based on Free-form Deformation. In *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, volume 6, pages 3672–3683, 2004. doi:10.2514/6.2004-4630.
- [40] G. R. Shubin. Obtaining Cheap Optimization Gradients from Computational Aerodynamics Codes. International Paper AMS-TR-164, Boeing Computer Services, June 1991.
- [41] G. R. Shubin and P. D. Frank. A Comparison of the Implicit Gradient Approach and the Variational Approach to Aerodynamic Design Optimization. International Paper AMS-TR-164, Boeing Computer Services, June 1991.
- [42] Benjamin S. Skrainka and Kenneth L. Judd. High Performance Quadrature Rules: How Numerical Integration Affects a Popular Model of Product Differentiation. Technical Report CEMMAP Working Paper CWP03/1, Centre for Microdata Methods and Practice, London, 2012.
- [43] Karla Telidetzki, Lana Osusky, and David W. Zingg. Application of jetstream to a suite of aerodynamic shape optimization problems. In *52nd Aerospace Sciences Meeting*, Feb. 2014. doi:10.2514/6.2014-0571.
- [44] E. van der Weide, G. Kalitzin, J. Schluter, and J. J. Alonso. Unsteady Turbomachinery Computations Using Massively Parallel Platforms. In *4th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, January 2006. doi:10.2514/6.2006-421.
- [45] D. W. Zingg, M. Nemec, and T. H. Pulliam. A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization. *European Journal of Computational Mechanics*, 17(1–2):103–126, January 2008. doi:10.3166/remn.17.103-126.