

Cross-Validation of High-Fidelity Aerodynamic Shape Optimization Methodologies for Aircraft Wing-Body Optimization

Thomas A. Reist,^{*} David Koo,[†] and David W. Zingg[‡]
Institute for Aerospace Studies, University of Toronto, Toronto, Ontario, M3H 5T6, Canada

Pascal Bochud,[§] Patrice Castonguay,[¶] and David Leblond^{||}
Bombardier Aerospace, Saint-Laurent, Québec, H4S 2A9, Canada

This paper presents cross-validation of two aerodynamic shape optimization methodologies in order to validate, characterize, and compare the two methodologies. Both methodologies use gradient-based optimization based on the Reynolds-averaged Navier-Stokes equations driven by the discrete adjoint method. The first methodology uses a B-spline surface representation of the geometry, coupled with free-form deformation and axial curve-based geometric control, and an efficient linear elasticity-based mesh deformation method. The second methodology uses class-shape-transforms for airfoil geometry control and either in-CAD or out-of-CAD surface modelling. Two benchmark geometries are used for comparison, a wing-body configuration representative of a large transport aircraft, and a wing-body configuration with aft-mounted nacelles typical of a small business jet. Both single and multipoint optimizations are performed. Cross-validation is conducted such that flow evaluations are performed on the geometries produced by the two methodologies using a single flow solver and consistent meshes. It is found that both methodologies provide similar geometries and performance improvements. The primary differences which exist are a result of the distinct geometry parameterization and control methods, which, in some cases, allow for more localized geometric control. The degree of agreement is indicative of the maturity that has been achieved by modern, state-of-the-art aerodynamic shape optimization methodologies.

Nomenclature

α = Angle of attack

^{*}Research Associate, tom.reist@utoronto.ca

[†]Research Associate, tai.koo@utoronto.ca

[‡]University of Toronto Distinguished Professor of Computational Aerodynamics and Sustainable Aviation, Centre for Research in Sustainable Aviation, and Associate Fellow AIAA, dwz@utias.utoronto.ca

[§]Engineering Specialist, Advanced Aerodynamics, pascal.bochud@aero.bombardier.com

[¶]Engineering Specialist, Advanced Aerodynamics, patrice.castonguay@aero.bombardier.com

^{||}Senior Engineering Specialist, Advanced Aerodynamics, david.leblond@aero.bombardier.com

- b = Wing semi-span
- c = Local chord length
- C_D = Drag coefficient
- C_L = Lift coefficient
- C_M = Pitching moment coefficient
- C_p = Pressure coefficient
- J = Objective function
- M = Mach number
- t/c = Thickness-to-chord ratio
- w_i = Objective function weights

I. Introduction

The need for more fuel-efficient aircraft, for both environmental and economic reasons, necessitates the use of state-of-the-art design tools in order to obtain the highest possible performance from a design. Aerodynamic shape optimization, which couples high-fidelity computational fluid dynamics (CFD) solvers with numerical optimization methods, is one such tool. With the information provided by the CFD solver, the optimizer can find the optimal geometry, potentially more quickly than a human designer, particularly when the design space is large and the optimum is unintuitive.

Aerodynamic shape optimization was pioneered by Hicks and Henne[1] in the 1970s, with significant advances made by Jameson[2] the following decade. Since then, aerodynamic shape optimization methods have become a focus of research and are widely used in the aerospace industry. Due to the increasing use of aerodynamic shape optimization, cross-validation of methodologies is important so that designers can have confidence that their optimization results are not overly dependent on the methodology used. Until recently, cross-validation of different aerodynamic shape optimization algorithms has been limited. Epstein et al.[3] performed optimization of a three-dimensional wing in turbulent flow using three different aerodynamic shape optimization methodologies. The drag reductions achieved by each optimization methodology were found to be in relatively good agreement when computed with an independent solver; however some geometric and performance differences were evident. This paper was presented as a step toward the establishment of benchmark problems for cross validation of aerodynamic shape optimization methodologies.

In 2013 the American Institute for Aeronautics and Astronautics launched the Aerodynamic Design Optimization Discussion Group, which established a set of benchmark aerodynamic shape optimization problems to be used for cross-validation of aerodynamic shape optimization methodologies[4–8]. Using a two-dimensional benchmark problem specified by the Aerodynamic Design Optimization Discussion Group, Destarac et al.[9] performed cross-validation of

seven optimized airfoils produced by the various discussion group participants. Each of these airfoils was produced using the same optimization problem formulation, i.e. objective and constraints, but with different optimization methodologies and numbers of design variables. With all optimized airfoils analyzed on a consistent set of grids and with the same solver, the authors showed discrepancies in computed drag when compared to the original optimization results. The paper then investigated the non-uniqueness of inviscid solutions and hysteresis effects, and the pitfalls that these imply in the use of aerodynamic shape optimization.

Using the Common Research Model (CRM) wing optimization problem specified by the Aerodynamic Design Optimization Discussion Group, LeDoux et al.[10] investigated the use of two optimization methodologies, one based on a potential flow analysis with a boundary layer solver and the second using a full Reynolds-averaged Navier-Stokes solver. This study examined the impact of the two different fidelity levels in the underlying flow solver, as well as the impact of the number of section shape design variables. The results highlighted the impact of flow solver fidelity on the optimized shape and performance, and how model differences, particularly in viscous modelling, can be exploited by the optimizer. At the main design point the impact of the number of shape design variables was found to be limited, but to play a large role in off-design performance.

This paper presents cross-validation of two state-of-the-art high-fidelity gradient-based aerodynamic shape optimization algorithms, one developed at the University of Toronto Institute for Aerospace Studies, the other at Bombardier Aerospace. The paper follows the spirit of the Aerodynamic Design Optimization Discussion Group, with an emphasis on industrially-relevant optimization problems, and with cross-validation results presented so as to enable direct comparisons independent of the final analysis solvers. Cross-validation is presented on two three-dimensional aircraft optimization problems for both single and multipoint optimizations. Comparisons will be made between the optimized shapes, pressure and lift distributions, and aerodynamic performance. This cross-validation exercise will establish whether these two state-of-the-art methodologies, developed independently and using different underlying algorithms for the flow solver, geometry control, mesh movement, and optimization, produce similar optimal shapes and aerodynamic performance.

The paper is structured as follows. Section II describes the two methodologies used, Section III details the optimization problem definition, including design variables, constraints, and operating conditions, and Section IV compares the optimal shapes and performance of the two methods. Section V explores the impact of the number of section shape design variables in one of the optimization methodologies.

II. Methodologies

This section describes the two methodologies and their underlying components. The first, called Jetstream, has been developed at the University of Toronto Institute for Aerospace Studies, and the second, called BOOST, at Bombardier Aerospace.




Fig. 1 Illustration of B-spline geometry parameterization and FFD control.

A. Jetstream

Jetstream comprises four main components: 1) a Newton-Krylov-Schur solver for the Reynolds-averaged Navier-Stokes equations coupled with the Spalart-Allmaras (SA-neg) turbulence model[11, 12], 2) an integrated geometry parameterization and linear elasticity-based mesh movement scheme[11], 3) a free-form and axial curve deformation technique for geometry control[13], and 4) the sparse nonlinear optimization package SNOPT[14], where gradients are computed via the discrete adjoint method which is tightly coupled to the above three components[11, 15].

The flow solver is a parallel implicit solver that uses summation-by-parts operators for spatial discretization and simultaneous approximation terms for the imposition of boundary conditions and block interface conditions. It uses multiblock structured grids, and has demonstrated good scaling behaviour to thousands of processors. The Krylov subspace method Generalized Minimum Residual (GMRES) is used with approximate Schur preconditioning in an inexact Newton method for the solution of the discrete equations. Details of the flow solver can be found in Hicken and Zingg[16] and Osusky and Zingg[12]. The flow solver has been validated through participation in the 5th AIAA Drag Prediction Workshop[17].

At each optimization iteration for which a geometric shape change occurs, the computational grid must be moved to reflect this change. To accomplish this, each block of the computational grid is fitted with a B-spline volume; this leads to an approximation of the initial geometry. The B-spline parameterization on the surface is embedded within a free-form deformation volume that can be controlled through free-form deformation (FFD) volumes and ‘axial curves’, as described by Gagnon and Zingg[13]. This is illustrated in Figure 1. The B-spline parameterization also facilitates mesh deformation. As the B-spline control points on the aerodynamic surface are moved, each B-spline volume block is treated as a linear elastic solid, for which a finite-element solution is obtained to define the new shape of the B-spline volume. The computational grid is then recovered from this new B-spline volume. This method has been found to be very robust for large shape changes while being relatively inexpensive. Details can be found in Hicken and Zingg[11].

Due to the high cost of evaluating the flow equations, a gradient-based optimizer is used for optimization, as

gradient-based optimizers typically require fewer function evaluations than genetic algorithms[18]. The gradients of the objective and flow-dependent constraints are evaluated using the discrete adjoint method. The number of adjoint solutions required is proportional to the number of flow-dependent functions, i.e. objective and constraints. Since this can require significant computational cost for practical problems, an efficient method of solving the linear system of the adjoint problem is required. For this, a modified, flexible version of the Generalized Conjugate Residual with Orthogonalization and Truncation (GCROT) algorithm is used[19]. The majority of the gradient terms required in the adjoint equations are calculated analytically, with some particularly complex terms evaluated using the complex step method. The flow solver, mesh movement, and geometry control are all tightly coupled in the development of the adjoint equations. The gradient-based optimizer SNOPT is used, as it allows for the solution of large-scale constrained problems. Details of the adjoint method and its integration with the flow solver and mesh movement are given by Hicken and Zingg[11], while the details of SNOPT are described by Gill et al.[14]

This optimization framework has been used extensively for the aerodynamic optimization of wings[5], conventional aircraft[15], and unconventional aircraft configurations[20–22].

B. BOOST

The BOOST adjoint-based optimization framework used at Bombardier consists of three main components: 1) the FANSC (Full Aircraft Navier-Stokes Code) flow solver; 2) a geometric modeller which can either be in-CAD or out-of-CAD, complemented by a block-structured mesh generator and a mesh-mover scheme; and 3) the efficient gradient-based interior point optimizer technique implemented in the IPOPT package.

FANSC is a finite-volume cell-centred three-dimensional multi-block Navier-Stokes flow solver. It features a second-order accurate spatial discretization using matrix dissipation or Roe's upwind scheme. The discretized system of equations is solved implicitly with the symmetric Gauss-Seidel scheme, or the generalized minimal residual (GMRES) approach, combined with acceleration techniques such as Full Approximate Storage Multigrid. Available turbulence models are the widely used Spalart-Allmaras and $k\text{-}\omega$ -SST equations, with the standard Spalart-Allmaras model used for this work. The code includes a farfield drag prediction method for drag computation analysis. Details on the flow solver can be found in Mohamed et al.[23]

During the optimization, the computational mesh needs to be moved according to the design variables which control the geometry. The Class-Shape-Transform (CST) methodology[24] is used to parameterize the wing section shapes. To obtain the surface mesh, two options are available: in-CAD or out-of-CAD. In the in-CAD approach, the aircraft surfaces are modelled and updated using the CATIA V5 modeller. The surface mesh is regenerated using the MBGRID mesh generator.[25] In the out-of-CAD approach, B-Spline surfaces are generated through the CST airfoil shapes and discretized in sets of Radial Basis Functions (RBF) control points. The displacement between an initial set and new set of RBF control points determines the movement of the surface mesh points. Once the surface mesh is obtained, either




Fig. 2 The two aircraft configurations under consideration, with control stations illustrated.

in-CAD or out-of-CAD, the volume mesh point displacements are calculated through an inverse distance interpolation scheme. Both of these geometry control methods are used in this work, as will be described later.

The discrete adjoint method is implemented in the FANSC flow solver to evaluate the gradients of the objective function and constraints. The mesh sensitivities are obtained by finite differencing for both the out-of-CAD and in-CAD approaches. The function and gradient information is then fed into the gradient-based optimizer IPOPT, which is an open source software well suited for large scale non-linear optimization that uses the interior point method.[26]

III. Optimization Problem Definition

A. Geometries

Two three-dimensional geometries are considered. The first is NASA's CRM wing-body geometry[27]. A variant of this geometry has been used as one of the Aerodynamic Design Optimization Discussion Group benchmark cases. The second case is representative of a small business aircraft, and is referred to as the Bombardier Research Aircraft Configuration (BRAC). Both of these geometries are shown in Figure 2.

B. Objective and Operating Conditions

Single and multipoint optimizations are considered for both the CRM and BRAC geometries. For the CRM, the main design point is at $M = 0.85$ and $C_L = 0.50$, where M is the Mach number and C_L is the lift coefficient of the aircraft. For the multipoint optimization, a nine point stencil in Mach number- C_L space is considered, with $M = [0.80, 0.85, 0.90]$, and $C_L = [0.45, 0.50, 0.55]$. While the high Mach number- C_L pairs are not a practical design condition, they provide a challenging condition for the cross-validation.

For the BRAC, the main design point is at $M = 0.80$ and $C_L = 0.45$. For the multipoint optimization, two operating points are considered, the first at $M = 0.80$ and $C_L = 0.45$, and the second at $M = 0.83$ and $C_L = 0.35$.

The objective for each optimization is to minimize total drag. In the case of the multipoint optimizations, the

Table 1 Summary of optimization cases

Geometry	M	C_L	Description
CRM	0.85	0.50	Single-point lift-constrained drag minimization
CRM	0.80, 0.85, 0.90	0.45, 0.50, 0.55	Multipoint lift-constrained drag minimization
BRAC	0.80,	0.45	Single-point lift-constrained drag minimization
BRAC	0.80, 0.83	0.45, 0.35	Multipoint lift-constrained drag minimization
BRAC	0.80, 0.83	0.45, 0.35	Multipoint trim-constrained drag minimization

objective is to minimize the sum of the weighted drag at each operating point, i.e.

$$J = \sum_{i=1}^N w_i C_{D_i}. \quad (1)$$

For the CRM case, the weights, w_i , are selected so that the sum forms a trapezoidal quadrature approximation of the integral of drag over the Mach number- C_L space, i.e. $w = [1/16, 1/8, 1/16, 1/8, 1/4, 1/8, 1/16, 1/8, 1/16]$. For the BRAC, the weights are $w = [2/3, 1/3]$. These weights prioritize drag reduction at the main design point, while incorporating a higher speed cruise condition.

Two multipoint cases are considered for the BRAC. For the first, only the lift is constrained to equal the targets listed above. For the second, a constraint is imposed on the trim, with the lift constraint adjusted accordingly. No horizontal stabilizer is included in the BRAC geometry, so the trimming capability of the tail is accounted for through an analytical trimming model. The adjusted trim lift to be constrained is given by

$$C_{L_{\text{trimmed}}} = C_{L_{\text{tail-off}}} + C_{M_{\text{tail-off}}} \frac{l_{\text{ref}}}{\cos(\alpha)(x_{\text{MAC, tail}} - x_{\text{CG}}) + \sin(\alpha)(z_{\text{MAC, tail}} - z_{\text{CG}})} \quad (2)$$

which accounts for a tail load required to trim the aircraft. The drag incurred by the tail is not accounted for in the objective. The tail and CG locations are fixed, and l_{ref} is a reference length, here taken to be the MAC. The angle of attack, α , is a design variable at each operating point and is bounded by $\alpha \leq 3.00$ deg for both the CRM and BRAC. The five cases considered in this paper are summarized in Table 1.

C. Geometric Variables and Constraints

For both the CRM and BRAC optimizations the planform is fixed and the geometric design variables include the wing twist and section shapes at a number of spanwise locations - 11 for the CRM and nine for the BRAC.* These stations are shown in Figure 2. At each of these stations, the section shape is free to vary using the methods described in Section II, as is the twist. In Jetstream, the section shapes are controlled by 6 chordwise FFD pairs of control points which are uniformly distributed along the chord, as illustrated in Figure 1, for both the CRM and BRAC. Therefore,

*For the BRAC optimization in Jetstream, four additional FFD control stations are required at the winglet root to properly enclose the geometry.

there are 143 and 195 geometric design variables for the CRM and BRAC, respectively. An additional study into the impact of the number of chordwise FFD design variables is presented in Section V, where 11 chordwise FFD control point pairs are used to control the section shape. In BOOST, the CRM and BRAC geometry are controlled with CST functions. At each control station the section shape is controlled by 12 CST shape variables, 6 for each of the lower and upper surfaces which are uniformly distributed along the chord, plus one twist variable. Therefore, there are 143 and 117 geometric design variables for the CRM and BRAC, respectively. For the CRM optimization, BOOST uses the out-of-CAD geometry control approach described in Section II.B, while the BRAC optimization uses the in-CAD approach.

Special treatment is required to handle geometric changes at the wing-body junction. For the CRM, the wing section at the fuselage intersection ($y/b = 10.4\%$) is fixed, and the wing surface between this fixed wing-fuselage intersection and the first control station (located at $y/b = 16\%$) varies based on the interpolation provided by either the FFD in Jetstream, or the B-spline surface in BOOST. For the BRAC, changes to the wing-fuselage intersection ($y/b = 9.7\%$) are permitted; however, the intersection is handled differently between Jetstream and BOOST. In Jetstream, changes at the wing root are propagated onto the fuselage using an algebraic relationship applied to the B-spline control points. The propagation function is such that the changes to the fuselage shape are limited to a small region near the wing-fuselage intersection. This results in small changes to the fuselage and belly fairing shape. In BOOST, the fuselage is invariant, with a new intersection of the wing and fuselage calculated in-CAD for each new wing geometry. The flow at the wing-body junction is important, with complexities such as side-of-body separation. Since designing the wing-body junction to eliminate these features is not the purpose of this work, the reduced geometric freedom in this region is acceptable so long as both optimizers have similar levels of control. For example, the known side-of-body separation present on the CRM is not removed in the optimizations presented here.

Since multidisciplinary requirements are not explicitly accounted for, geometric constraints are introduced to capture their effect and prevent unrealistic geometries. To account for structural and internal space requirements, a number of thicknesses are constrained. At each control station the thickness at 15%, 90%, and 100% chord is constrained to be greater than or equal to the original value. The maximum thickness, the location of which is free to move, is constrained to be greater than or equal to the initial value. To account for low-speed performance requirements, the leading-edge radius is constrained to be greater than or equal to its initial value. These multidisciplinary requirements can also be addressed through other constraints, such as the use of volume or sectional area constraints to reflect structural requirements, or explicit $C_{L_{\max}}$ constraints to account for low speed requirements[28].

IV. Results

This section presents the results of the five optimization cases as obtained by each of the aerodynamic shape optimization methodologies, and compares and contrasts the results. To facilitate cross comparison of the two

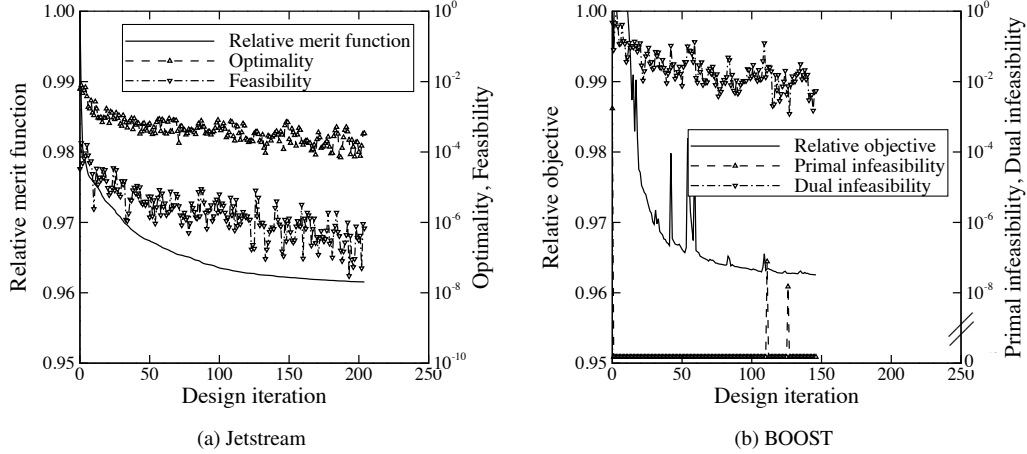


Fig. 3 Optimization convergence histories for the single-point CRM optimization with both optimizers.

methodologies and remove a source of discrepancy, all functionals, pressure distributions, and lift distributions are computed with BOOST unless otherwise stated.

A. NASA Common Research Model (CRM)

BOOST optimizations are performed on the medium level of the AIAA Drag Prediction Workshop (DPW) grids, with 5.1 million nodes. Due to fitting requirements within Jetstream, the DPW grids cannot be used for optimization; thus in-house grids are created with properties which mimic the DPW grids. These grids have 8.0 million nodes. Grids used for optimization must be sufficiently accurate and properly resolve the relevant flow features. For example, shocks and the wing-body separation region present on the CRM geometry must be captured, as they are a significant source of drag which the optimizer must be able to minimize or remove. The CRM Drag Prediction Workshop grids are known to be of good quality and to resolve these features. The in-house grids created for use with Jetstream closely mimic the Drag Prediction Workshop grids and capture these features. Determination of grid-independent results is important for performance comparison, so grid convergence studies are conducted on all optimized designs when reporting the performance of optimized geometries. The grid-converged drags reported in this paper are obtained via Richardson extrapolation based on the three finest grid levels. Grid-convergence studies conducted using both sets of grids on the baseline CRM geometry give grid-converged drag values within 0.3 drag counts when computed by both solvers.

1. Single-Point Optimization

The optimization convergence histories for the two methodologies are shown in Figure 3. The use of two different optimization packages, SNOPT and IPOPT, necessitates reporting of different optimization statistics. The relative merit function reported by SNOPT is analogous to the relative objective in the IPOPT convergence history. These two quantities are equal when the constraints are satisfied in the SNOPT optimization. In SNOPT, optimality is a measure of the accuracy of the dual variables and reflects the degree of convergence of the problem. Feasibility is the norm




Fig. 4 The optimized CRM shape produced by Jetstream, showing nonlinear tip twist feature.

of the constraint violations. In IPOPT, primal infeasibility is a measure of the constraint violation and is analogous to feasibility in SNOPT. The dual infeasibility is the maximum of the infeasibility of the IPOPT dual variables, and is analogous to optimality in SNOPT. At each function call, Jetstream solves the flow and adjoint problems to a final residual tolerance of at most 10^{-8} . BOOST uses a flow solver convergence criterion based on C_L of at most $\Delta C_L = 10^{-5}$, while the adjoint problem is solved to a residual tolerance of at most 10^{-5} . While these convergence tolerances affect the time to optimization convergence, they are sufficiently low that the final optimized design is independent of their values.

Figure 3 shows design, or major, iterations on the abscissa. However, multiple flow solves may be required during the line search procedure or backtracking, thus requiring more computation cost than is suggested in the figure. For the CRM optimization, Jetstream performed 213 flow solves for the 204 design iterations, while BOOST performed 264 flow solves for the 147 design iterations.

Both optimizations are run until all constraints are satisfied and there is minimal change in the objective between iteration. For both optimizations, the majority of the drag improvement is achieved in the early iterations. For example, Jetstream has reduced the drag by 2.8%, 3.3% and 3.7% by iteration 25, 50, and 100, respectively. In comparison, BOOST reduces the drag by 2.7%, 3.4%, and 3.7% at the same iterations. The total reduction achieved by Jetstream and BOOST is 3.8% and 3.7%, respectively.

For the Jetstream optimization, at the converged solution no geometric variables are at their bounds, and the angle-of-attack is 2.72° . For both optimizations all of the leading-edge radius constraints are active, with the exception of the root station in the Jetstream optimization. The minimum t/c_{\max} constraints are active, the t/c constraints at 15% and 90% chord are active at all but the two tip control stations. The lift constraint is active in both the Jetstream and BOOST optimizations. In the BOOST optimization, three of the section shape variables are at their bounds, and the angle-of-attack is 2.57° .

The initial CRM optimization conducted with Jetstream produces the shape shown in Figure 4. The optimizer utilizes the twist freedom near the tip to create a highly nonlinear twist variation localized near the tip. This feature is not seen in the BOOST optimization results, although the freedom to obtain it is present in the design space. This

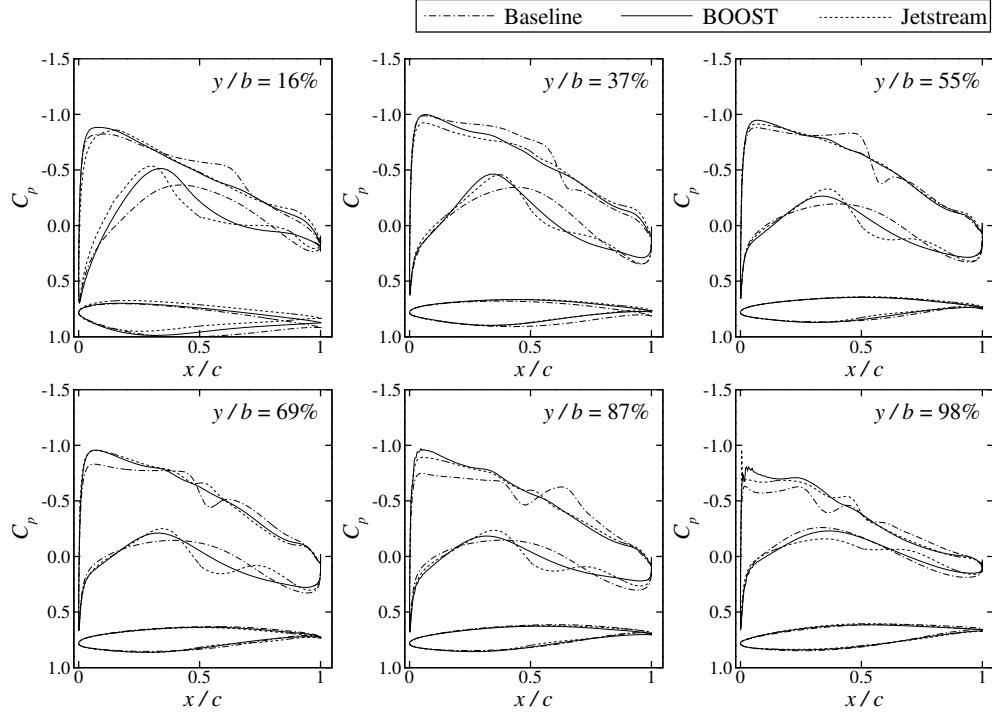


Fig. 5 Pressure distributions on the baseline and single-point optimized CRM designs.

feature is the largest discrepancy between the BOOST and Jetstream optimization results. To determine its performance benefit and to identify how similar the BOOST and Jetstream results would be if such a feature was not present, the ability to select a nonlinear twist distribution over the four outboard-most control sections was removed by constraining these control sections to have a linearly varying application of twist design variables. With this nonlinear tip feature removed, a comparison is made between the BOOST and Jetstream results. Figure 5 shows the sectional shapes and pressure distributions at six spanwise slices. The slice at 16% semispan corresponds with the location of the first control station. Pressure distributions are computed on the optimization level grids using BOOST’s flow solver, FANSC. The Jetstream section shapes in this figure are for the geometry with the nonlinear tip feature removed through the additional constraints. Over the span, the suction side of the wing has a very similar pressure distribution between the two optimized shapes. On the pressure side of the wing, the Jetstream result shows more curvature variation, with more lift carried just aft of the mid-chord and near the trailing edge. Figure 6 details differences between the section shapes generated by BOOST and Jetstream. The largest geometric differences are on the lower surface on the aft half of the chord. This variation is a consequence of the FFD geometry control and the presence of the t/c constraint at 90% chord. The combination of this constraint and the coarse chord-wise FFD control distribution causes the optimizer to thicken the section aft of approximately 70% chord to satisfy the constraint, while minimizing the thickness upstream of 70% to minimize drag. The impact of the thickness constraints and the number of section shape design variables will be discussed further in Section V. The section/twist differences near the tip between the two Jetstream cases - those with

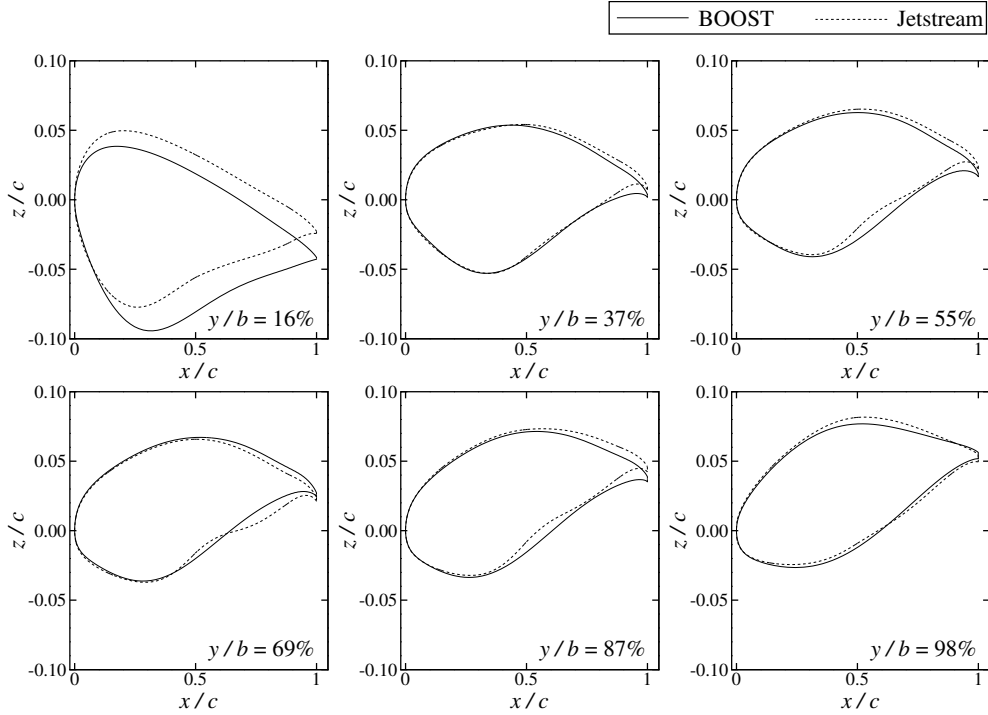


Fig. 6 Section shapes of the single-point optimized CRM designs.




Fig. 7 Variation in section shape and pressure distribution near the wing tip for designs optimized with and without nonlinear tip twist.

and without the nonlinear tip feature - are illustrated in Figure 7. Inboard of 87% chord, the section/twist variation between these two cases is minimal.

The spanwise lift distributions for the three optimized shapes are shown in Figure 8. There is little discernible difference between the three lift distributions, and both are close to an elliptical distribution, while carrying slightly less lift inboard. The Jetstream-optimized design with the nonlinear tip twist features carries slightly more lift near the tip.

The grid convergence behaviour for each of the optimized designs, in addition to the baseline CRM, is computed on grids with up to 113 million nodes and plotted in Figure 9. To remove the influence of CFD solver differences, the computations for these grid convergence studies are all done with BOOST's CFD solver, FANSC. However, similar

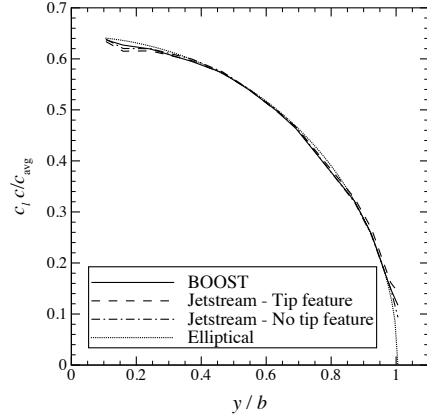


Fig. 8 Spanwise lift distribution on the single-point optimized CRM designs.

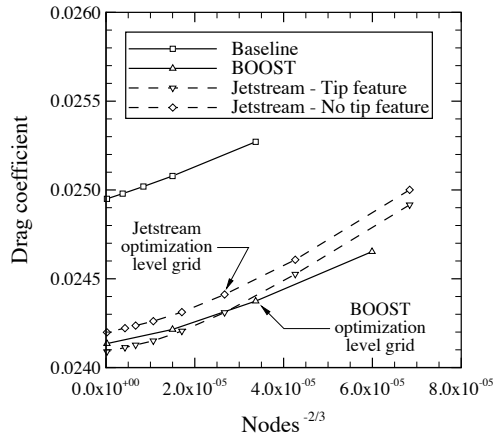


Fig. 9 Grid convergence behaviour of the baseline and single-point optimized CRM Designs. All values computed with BOOST.

grid-converged ΔC_D s are achieved when analyzed with Jetstream. For example, the Jetstream-optimized case with nonlinear tip twist has a grid-converged ΔC_D of 8.8 counts when analyzed with Jetstream, as compared to the 8.6 counts when analyzed with BOOST. Extrapolated grid-converged values are also shown and are tabulated in Table 2. The drag difference between the BOOST and Jetstream-optimized shapes is -0.7 and $+0.4$ counts for the linear and nonlinear tip twist cases, respectively. The nonlinear tip twist found in the Jetstream result provides a 1.1 count benefit compared to the Jetstream optimized design with this highly nonlinear tip twist prevented by additional constraints. The drag of the BOOST-optimized shape falls between these two results. These two optimization frameworks, using different flow solvers, optimization packages, and geometry control methodologies produce final designs with similar performance through shapes which are similar, but feature some notable differences.

2. Multipoint Optimization

For both the Jetstream and BOOST optimizations, the same constraints are active at the converged solution. The angle-of-attack goes to its upper bound of 3° at the $M = 0.80$, $C_L = 0.55$ operating point for both optimizations. As

Table 2 Grid-converged drag values for the baseline and single-point optimized CRM.

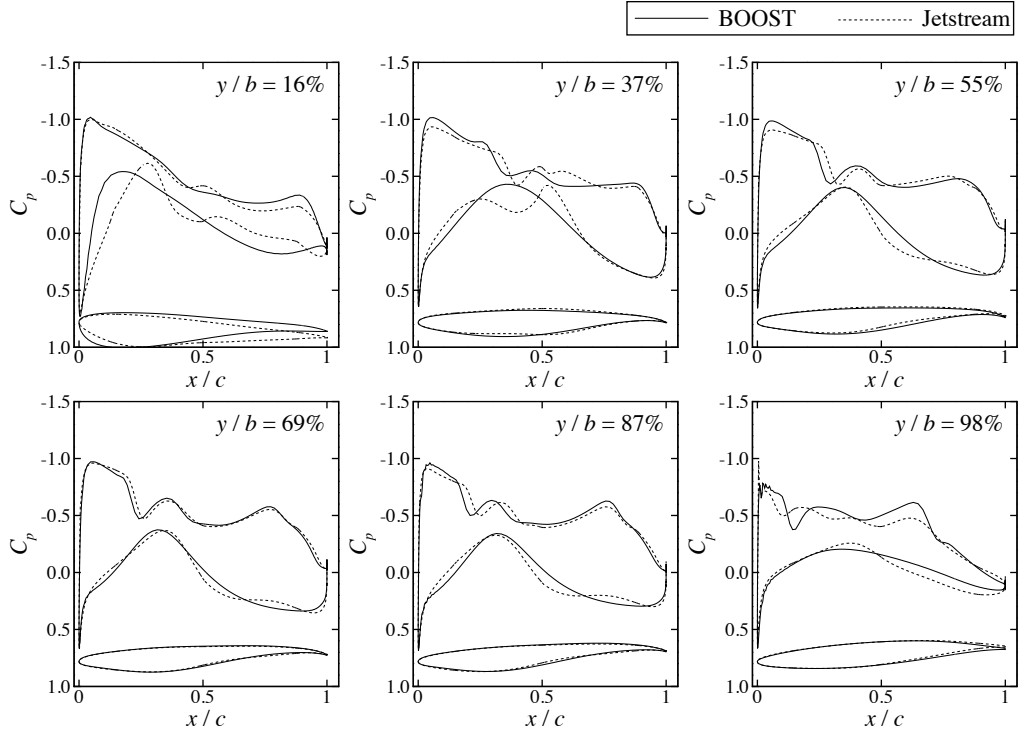
	C_D [counts]	ΔC_D [counts]
Baseline	249.5	–
BOOST	241.3	–8.2
Jetstream - nonlinear tip twist	240.9	–8.6
Jetstream - linear tip twist	242.0	–7.5

with the single-point optimizations, Jetstream produces a shape with a highly nonlinear twist variation near the tip. A second optimization is conducted with this feature prevented. Sectional shapes and pressure plots are shown in Figure 10 for the BOOST geometry and the Jetstream geometry with the linear tip twist. Results are shown for the main operating point, $M = 0.85$ and $C_L = 0.50$, and for the most challenging point of $M = 0.90$ and $C_L = 0.55$. The two optimizations produce section shapes and pressure distributions which are quite similar for the outboard half of the span. Inboard of 50% span, the section shapes, and hence pressure distributions, show larger discrepancies between the two optimizations, with significantly different section shapes at the wing root. Neither optimizer now produces a shock free design at the main design point, and strong shocks remain at the high speed - high lift operating point. As with the single-point optimization, the largest geometric differences between the two codes are produced on the lower surface, as seen in Figure 11.

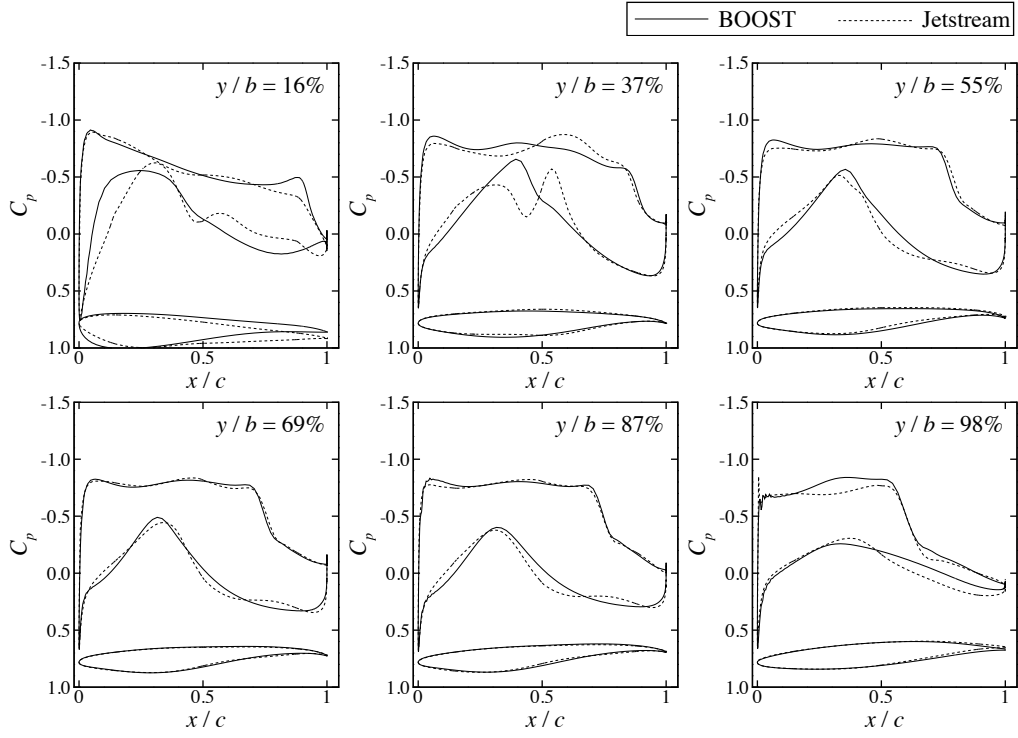
Grid convergence studies were conducted for each operating point using FANSC as was done for the single-point optimization. The grid-converged drag at each operating point for the optimized designs is shown in Figure 12. The figure shows the variation in drag with Mach number and lift coefficient, as well as the differences in performance between the three optimized designs. The drag difference between the two Jetstream-optimized shapes and the BOOST-optimized shape is shown as ΔC_D , where $\Delta C_D = C_{D_{\text{Jetstream}}} - C_{D_{\text{BOOST}}}$. At the two lower Mach numbers, $M = 0.80$ and $M = 0.85$, Jetstream produces designs with up to 4.7 counts lower drag. At the high Mach number, $M = 0.90$, BOOST produces lower drag, with up to a 4.6 drag count benefit at $C_L = 0.55$. Due to the strong shock present on the baseline design at $M = 0.90$, the majority of the drag reduction achieved by all optimizers is obtained at the high Mach number and lift coefficient operating points, with a reduction of up to almost 170 counts at $M = 0.90$ and $C_L = 0.55$. The weighted drag reduction, i.e. objective function, across all operating points achieved with BOOST and Jetstream is shown in Table 3, and is within 1.8 counts, with a 0.6 count difference between BOOST and Jetstream for the geometry with the enforced linear tip twist. As with the single-point result, the nonlinear tip twist yields a drag benefit, 1.2 counts in this case.

B. Bombardier Research Aircraft Configuration (BRAC)

The BRAC optimizations conducted with Jetstream and BOOST use two different grids. Those used by BOOST have 15.2 million nodes, those used by Jetstream have 15 million nodes. Both grids are constructed with similar resolution at



(a) $M = 0.85, C_L = 0.50$



(b) $M = 0.90, C_L = 0.55$

Fig. 10 Pressure distributions on the multipoint optimized CRM designs.

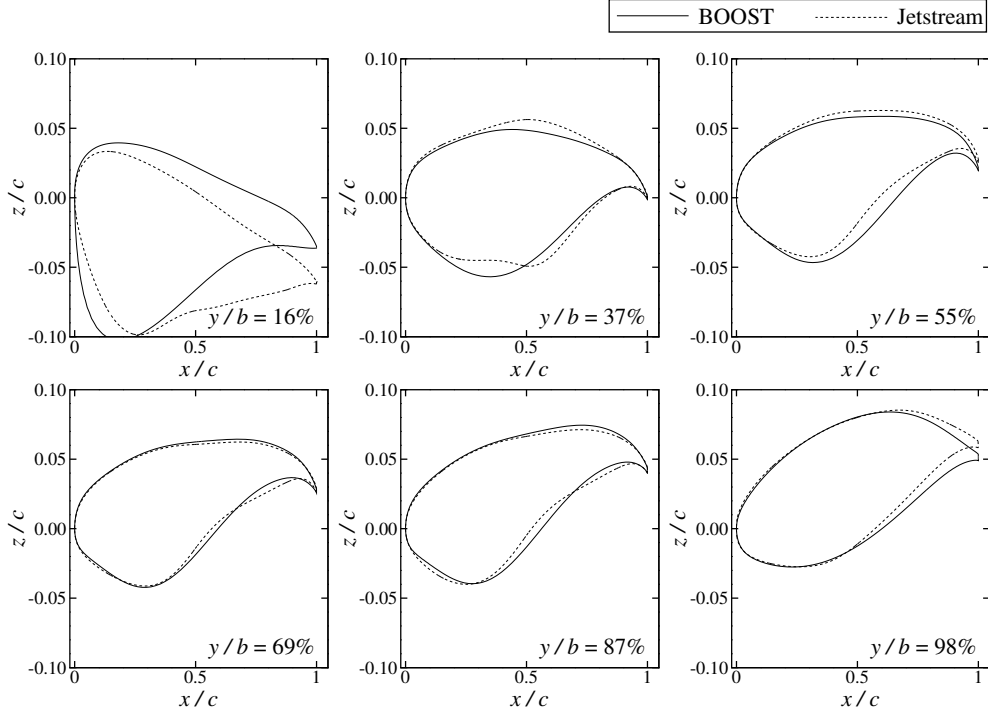


Fig. 11 Section shapes of the multipoint optimized CRM designs.

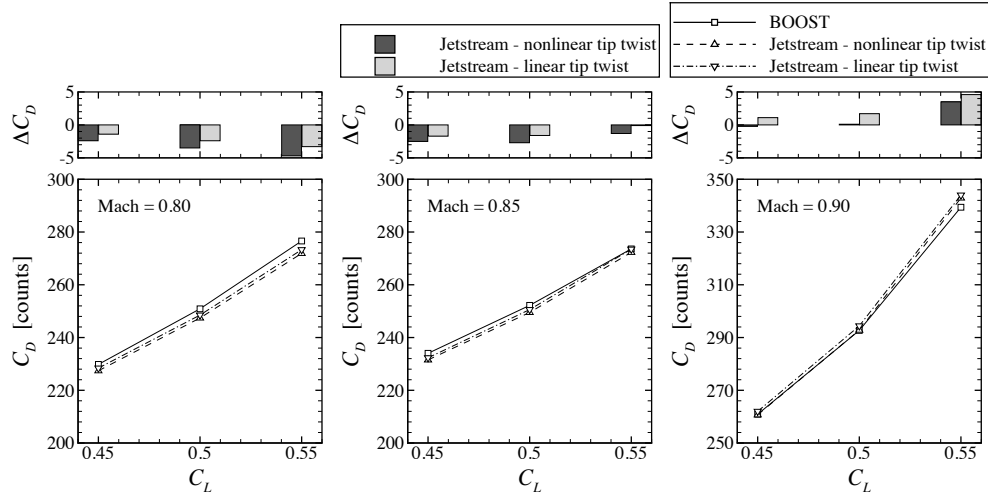


Fig. 12 Grid-converged drag at each operating point for the multipoint-optimized CRM designs, where $\Delta C_D = C_{D_{\text{Jetstream}}} - C_{D_{\text{BOOST}}}$

critical locations, e.g. off-wall spacing, leading-edge and trailing-edge spacings, etc.

1. Single-Point Optimization

The optimization convergence histories for both optimizers are presented in Figure 13. The same convergence criteria are used as for the CRM optimization. Jetstream performs 119 flow evaluations for the 112 design iterations, while BOOST performs 96 flow solves for the 66 design iterations. As with the CRM, the majority of the drag reduction

Table 3 Grid-converged objective function values for the baseline and multipoint optimized CRM.

	Objective, J [counts]	ΔJ [counts]
Baseline	294.5	–
BOOST	263.6	–30.9
Jetstream - nonlinear tip twist	261.8	–32.7
Jetstream - linear tip twist	263.0	–31.5

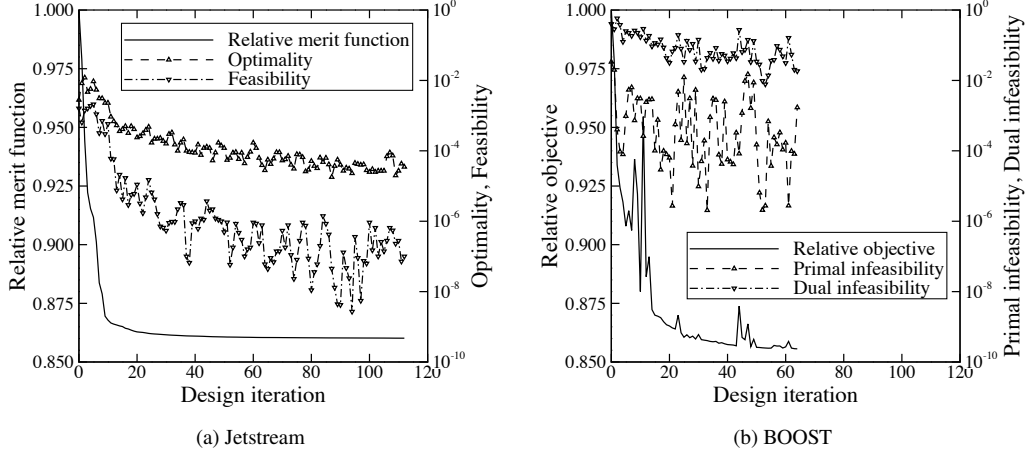


Fig. 13 Optimization convergence histories for the single-point BRAC optimization with both optimizers.

is achieved in the early stages of the optimization. The Jetstream optimization was run longer than necessary, as almost all of the drag improvement is obtained in the first 40 iterations.

At the converged solution for the Jetstream optimization, none of the geometric variables are at their bounds, with the exception of the root control station. To restrict geometric changes at the wing-fuselage junction, the section shape variable bounds are reduced at the first control station; the design variables which control the section shape are limited to scaling FFD shape variables by $\pm 10\%$. In the BOOST optimization, three of the twist design variables reach their upper bound of 2.5° . For both optimizers, the angle-of-attack goes to its upper bound of 3° . For both optimizers, the leading-edge radius and t/c_{\max} constraints are active, as are the majority of the t/c constraints at 15% and 90% chord. The lift constraints are active in all cases.

To perform the post-optimization analysis, the optimal shape produced by Jetstream is remeshed and analyzed with BOOST, in order to eliminate any differences arising from the flow solver. The section shapes and optimal pressure distributions at six wing stations are shown in Figure 14. Sections shown with $y/b > 1$ are on the winglet. The station with $y/b = 101\%$ is at the winglet root, and that with $y/b = 108\%$ is near the winglet tip. There are some differences in pressure on the upper surface, primarily near the root, where the Jetstream-optimized shape produces a smaller suction peak. On the winglet, Jetstream produces a shock-free design while BOOST retains a shock on the winglet and near the wing tip; the drag penalty from this shock is minor however, as evidenced by the 0.5 drag count difference between the

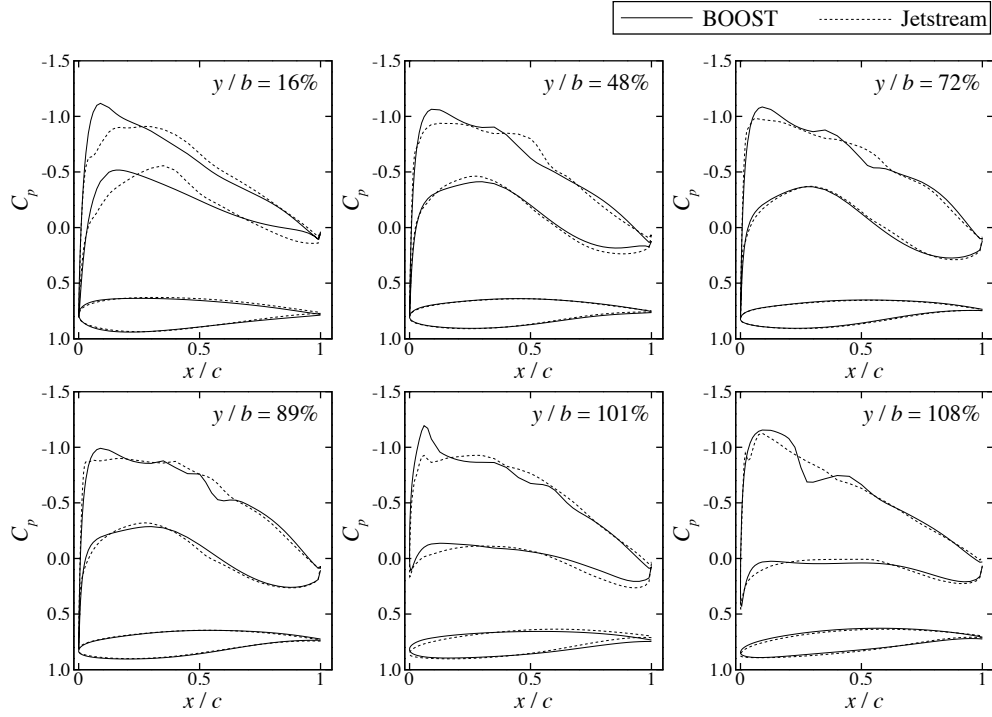


Fig. 14 Pressure distributions on the baseline and single-point optimized BRAC designs.

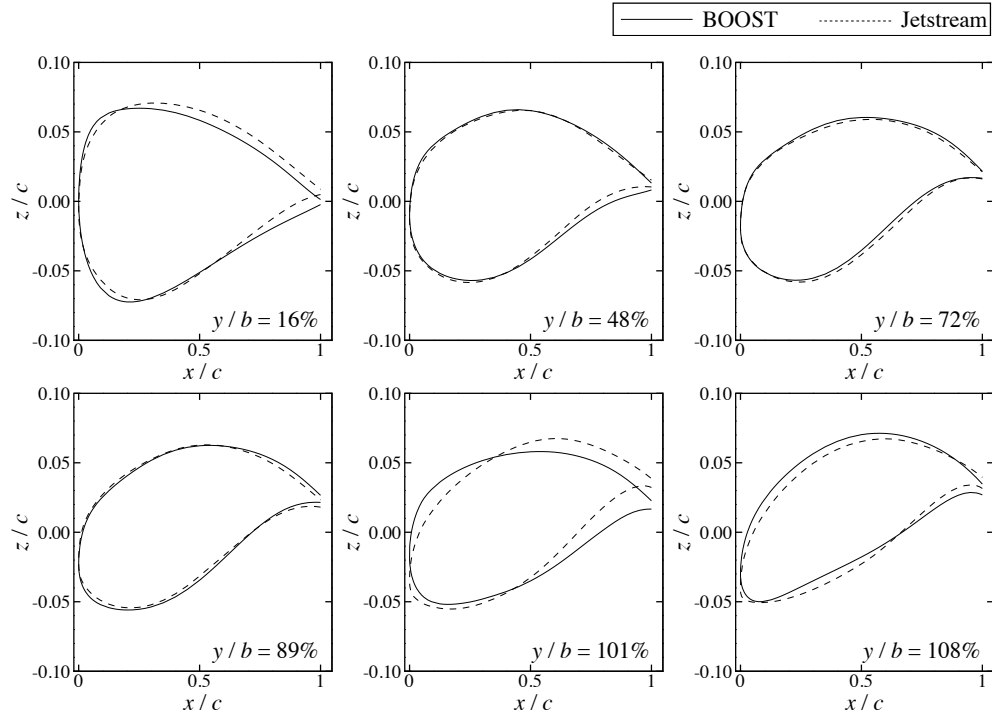


Fig. 15 Section shapes of the single-point optimized BRAC designs.

performance of the BOOST and Jetstream optimized shapes. Details of the section shapes are shown in Figure 15, where the close similarity of the two optimized shapes can be seen, with some differences near the root and on the winglet.




Fig. 16 Spanwise normal force distribution on the single-point optimized BRAC designs.

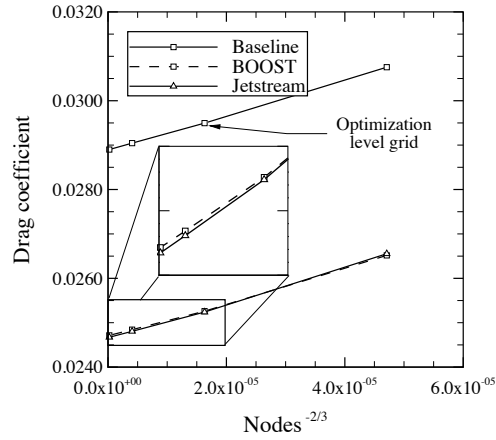


Fig. 17 Grid convergence behaviour of the baseline and single-point optimized BRAC Designs. All values computed with BOOST.

Table 4 Grid-converged drag values for the baseline and single-point optimized BRAC.

	C_D [counts]	ΔC_D [counts]
Baseline	289.0	–
BOOST	247.2	–41.8
Jetstream	246.7	–42.3

The spanwise distribution of the wing-plane-normal force are shown in Figure 16. Both optimizations produce very similar distributions which deviate from an elliptical distribution due to the winglet.

Grid convergence studies on the baseline and two optimized shapes are conducted using FANSC with up to 121 million nodes; the behaviour of drag with grid density is shown in Figure 17. Drag values reported in Table 4 are grid-converged estimates based on these grid convergence studies. Both optimizations produce very similar optimized performance, within 0.5 counts.

Table 5 Grid-converged drag values for the baseline and multipoint optimized BRAC.

	Untrimmed ¹			Trimmed ²		
	M	C_L		M	C_L	
	0.80	0.83	–	0.80	0.83	–
	0.45	0.35	–	0.45	0.35	–
	C_D		J	C_D		J
	[counts]		[counts]	[counts]		[counts]
Baseline	289.0	309.3	295.8	308.2	328.2	314.9
BOOST	249.3	225.5	241.4	259.3	236.9	251.8
Jetstream	248.7	229.4	242.3	258.1	238.8	251.7

¹ $C_L = C_{L_{\text{tail-off}}}$

² $C_L = C_{L_{\text{trimmed}}}$

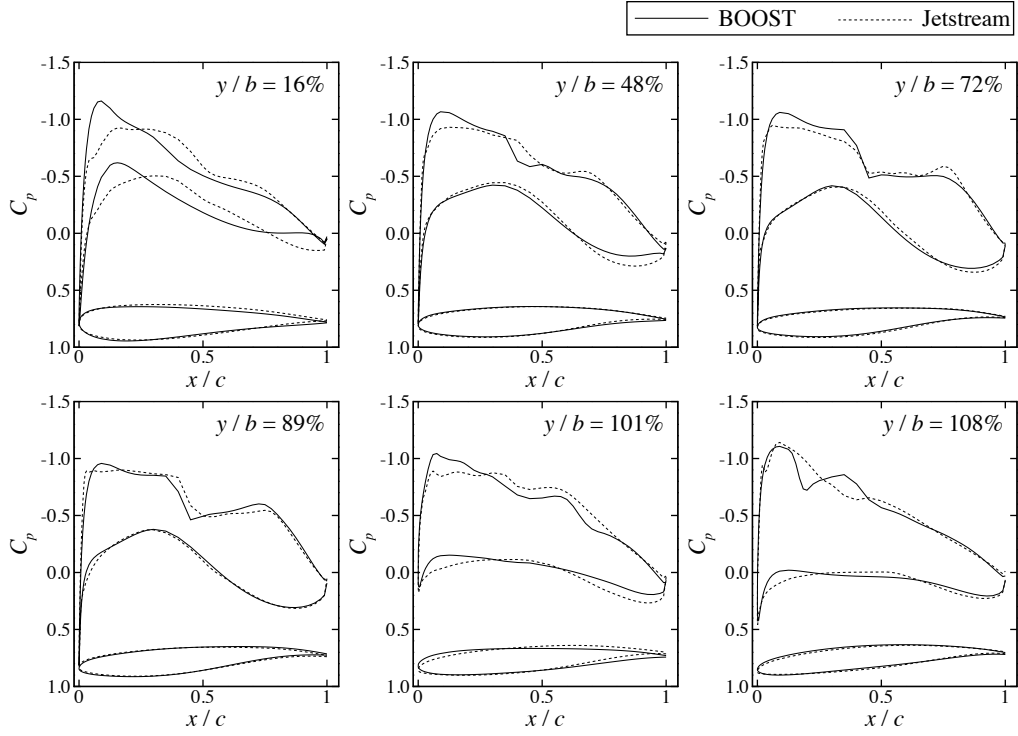
2. Multipoint Optimization

Two multipoint optimizations are presented for the BRAC. The first is similar to the single-point case, where the wing-body lift is constrained, while in the second a ‘trimmed lift’ is constrained, as described in Section III. Results for both multipoint cases are presented here. Optimization convergence histories are similar to the single-point case.

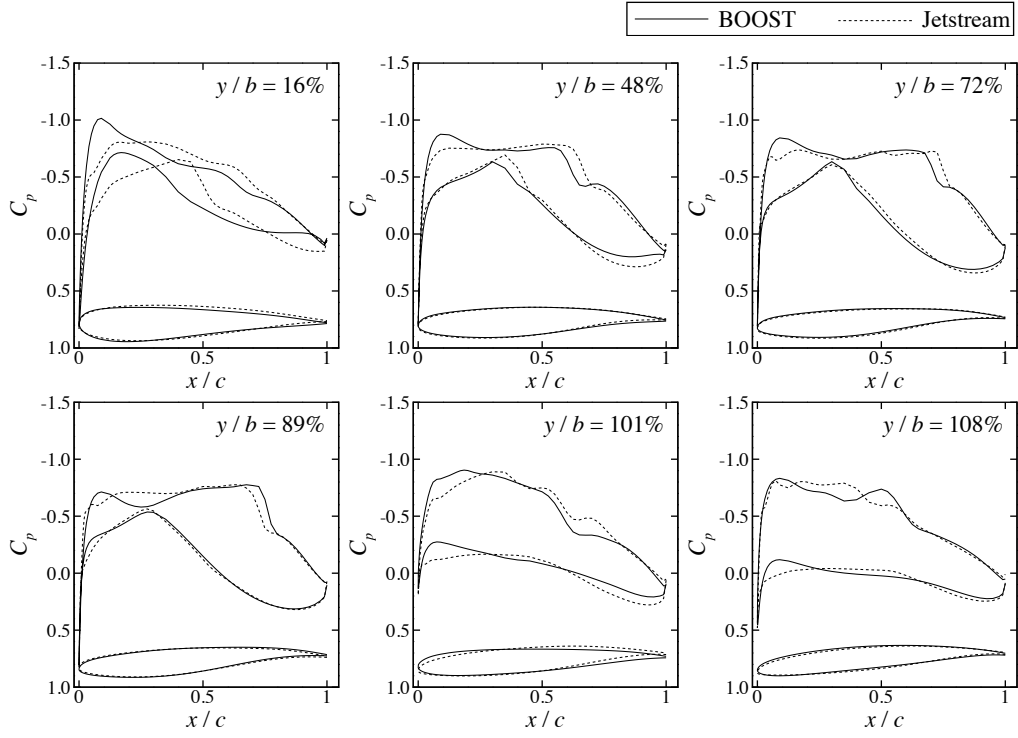
At the converged solution for all of the multipoint cases, the same geometric constraints are active as for the single-point case. The angle-of-attack goes to its upper bound at the $M = 0.80$, $C_L = 0.45$ operating point for both cases and both optimizers. The lift constraints, untrimmed and trimmed, respectively are active.

Grid convergence studies are conducted as for the single-point case. Table 5 shows the grid-converged drag estimates for both multipoint cases as found by both optimization methodologies, as well as the objective function value. Optimal drag at the main design point ($M = 0.80$, $C_L = 0.45$) is within 0.6 counts for the untrimmed case, and 1.2 counts for the trimmed case. For the high speed case ($M = 0.83$, $C_L = 0.35$) a more significant drag difference exists, particularly for the untrimmed case, of up to 3.9 counts. For both multipoint cases, BOOST achieves a lower drag at the high speed operating point, while Jetstream achieves lower drag at the main design point. The weighted objective function, also tabulated in Table 5, is within 0.9 counts between the two optimization methodologies for the untrimmed case and within 0.1 counts for the trimmed case. Figure 18 shows the pressure distributions for the two operating points of the untrimmed case. The magnitude of the differences between the section shapes found by the two optimization methodologies is similar between the single and multipoint optimizations.

For both optimization methodologies the trim requirement incurs a drag penalty. At the main design point, the penalty is 10.0 counts for the BOOST optimization and 9.4 counts for the Jetstream optimization. The penalty at the $M = 0.83$ point is similar in magnitude. Figure 19 shows select pressure distributions at the main design point for the untrimmed and trimmed designs as designed by both optimizers. Both optimizers decrease aft loading when the trim requirement is imposed, resulting in a slight increase in the shock strength. Jetstream decreases the aft loading on the inboard portion of the wing, while BOOST has a larger aft loading reduction on the outboard wing. The increased shock



(a) $M = 0.80, C_L = 0.45$



(b) $M = 0.83, C_L = 0.35$

Fig. 18 Pressure distributions on the multipoint optimized BRAC designs.

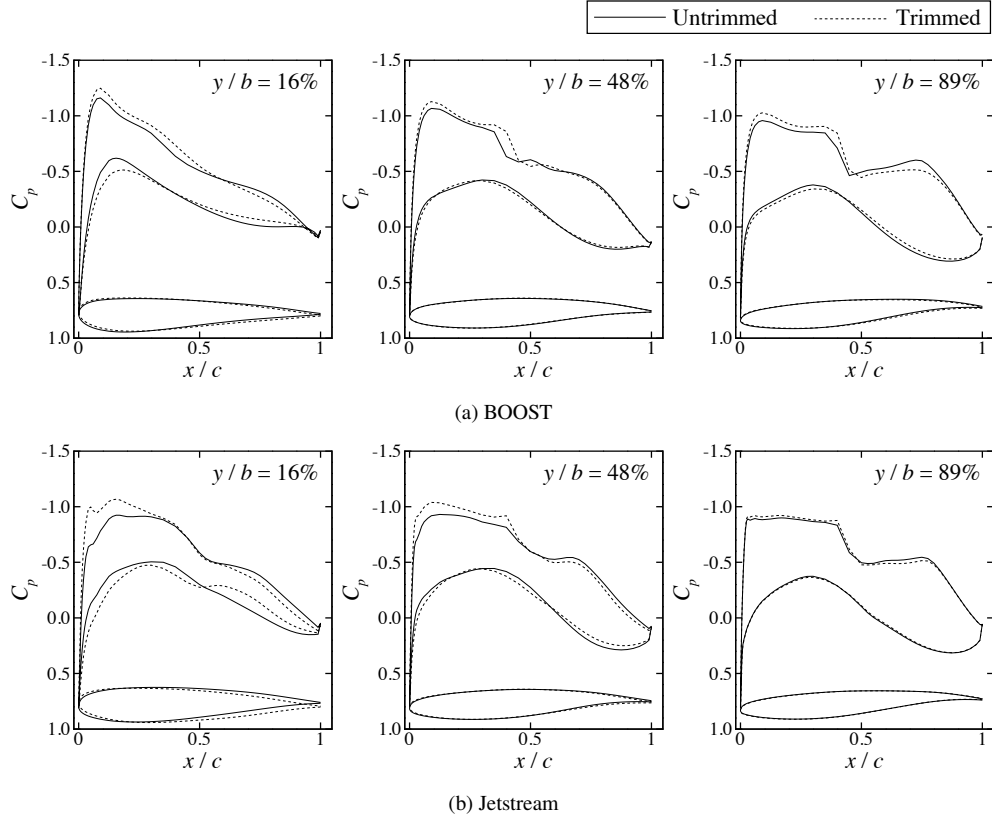


Fig. 19 Pressure distributions on the untrimmed and trimmed designs as optimized with the two optimization methodologies.

strength combined with the increased $C_{L_{\text{tail-off}}}$ results in the increased drag from the trim requirement.

V. Impact of the Number of Design Variables

To quantify the impact of the number of section shape variables on optimized shape and performance, Jetstream optimizations are conducted on both the CRM and BRAC cases with twice the number of section shape FFD design variables. Due to computational limitations, a similar study is not conducted with BOOST. In the previous optimizations, Jetstream used 6 chordwise pairs of FFD control points to control the section shape, for a total of 12 section shape design variables per section. In this study, the optimizations are repeated with 11 chordwise pairs for a total of 22 section shape variables per section. This study is conducted on the single and multipoint CRM cases, with the nonlinear twist near the wing tip permitted, and on the single and multipoint BRAC cases. Of the two BRAC multipoint cases, only the untrimmed case is used for this study. With the exception of the number of design variables, the optimization problem definition for each case remains the same as that presented previously.




Fig. 20 Section shapes and pressure distributions for the single-point Jetstream-optimized CRM, when optimized with 12 and 22 section shape variables.

A. NASA Common Research Model (CRM)

The optimized section shapes and pressure distributions for the single-point Jetstream-optimized CRM are shown in Figure 20 for the cases with 12 and 22 design variables (DV) per section. The upper surface pressure distribution is similar for both cases, with noticeable differences in shape and pressure on the lower surface. In particular, the design optimized with 22 section shape design variables per section carries more lift on the aft half of the airfoil. This is the result of increased camber and reduced thickness of the airfoil from approximately 50-90% of the chord. With 6 FFD control point pairs controlling the section shape, the thickness constraint at 90% chord prevents thinning of the section immediately upstream of 90% chord due to the coarseness of the control parameterization. With 11 FFD control point pairs over the chord, there is sufficient control to allow the optimizer to satisfy the 90% chord thickness constraint while still thinning the section between 50-90% chord. The influence of the thickness constraint at 90% chord is evident by the thickening of the section at this location and the resulting sharp suction spike on the lower surface in this region. This increased freedom allows for a drag reduction of 1 drag count compared to the 12 variables-per-section case.

For the multipoint CRM case, the benefit of the increased freedom becomes more pronounced. The increased freedom allows the optimizer to better tailor the section shapes at each operating point. The resulting performance at each operating point is shown in Figure 21. The multipoint BOOST-optimized results from Section IV.A.2 are included for reference, and $\Delta C_D = C_{D_{\text{Jetstream}}} - C_{D_{\text{BOOST}}}$. With the increased geometric control, greater drag reductions are seen at each operating point, with the largest benefit of 12.3 drag counts at $M = 0.90$ and $C_L = 0.50$, relative to the




Fig. 21 Grid-converged drag at each operating point for the multipoint-optimized CRM designs, including Jetstream optimization with varying numbers of section shape design variables, where $\Delta C_D = C_{D, \text{Jetstream}} - C_{D, \text{BOOST}}$

Table 6 Grid-converged drag values for the baseline and multipoint optimized BRAC, including Jetstream optimization with varying numbers of section shape design variables.

	M	0.80	0.83	–
	C_L	0.45	0.35	–
		C_D		J
		[counts]		[counts]
Baseline		289.0	309.3	295.8
BOOST		249.3	225.5	241.4
Jetstream - 12 DVs / sec.		248.7	229.4	242.3
Jetstream - 22 DVs / sec.		247.6	226.1	240.4

12 variables-per-section case. The weighted drag objective function value is reduced by 5.2 counts relative to the 12 variables-per-section case, with most of this reduction coming from the high Mach number operating points.

B. Bombardier Research Aircraft Configuration (BRAC)

For the BRAC, the single-point optimization benefit is similar to that of the CRM case, at 0.9 drag counts. The interaction of the thickness constraint at 90% chord with the increased geometry control is again seen, but to a lesser extent. The larger impact in this case is a reduction in washout. The tip washout in the 12 variables-per-section case is 4.0° , while that of the 22 variable case is 0.2° . This is compensated for with increased camber to maintain an equivalent sectional lift, and hence similar spanload.

As with the CRM, the multipoint BRAC case experiences a greater benefit from the increased freedom. The weighted drag objective function is reduced by 1.9 counts due to the increased freedom. The majority of this reduction, is achieved at the higher Mach number operating point. The drag reduction for each case is given in Table 6.

As expected, increasing the geometric freedom available to the optimizer to control the section shapes results in designs with lower drag. Part of the reason for this in the cases presented here is the interaction between the geometry control and the thickness constraints, particularly the thickness constraint at 90% chord. It may be possible to formulate an alternative constraint, which aims to prevent an excessively thin trailing edge, that does not exhibit this interaction with the geometry control method so as to achieve improved performance even with fewer design variables. The increased level of geometric control also leads to less smooth geometries, again, partly due to the thickness constraint, which may be undesirable. This trade-off between improved performance and geometries which may be deemed undesirable should be considered and potentially addressed through the addition of constraints to remove the undesirable geometric feature rather than omitting them from the design space by reducing the geometric flexibility.

VI. Conclusions

This paper presents cross-validation of two state-of-the-art aerodynamic shape optimization methodologies, one developed in academia, the other in industry. The methodologies use different geometry control, mesh movement, CFD solvers, and optimization algorithms. The purpose is to investigate the degree of agreement achieved with respect to the optimized geometries and their performance. Cross-validation is conducted using grid convergence studies and a single solver so as to remove grid and CFD solver dependencies.

Two aircraft configurations were considered, with single and multipoint optimizations being conducted. The geometries resulting from the single-point optimizations are in good agreement for both configurations, with some differences stemming from the interaction between the geometry control and thickness constraint formulation. Despite the geometric differences, the optimized drag differs by no more than 0.7 counts. The most notable geometric difference was the introduction of highly nonlinear twist near the wing tip for the CRM case by the Jetstream methodology. The multipoint-optimized CRM showed larger section shape differences at the wing root than the single-point case, with a corresponding performance difference of up to 4.6 drag counts at the high Mach number-high lift coefficient operating point, and of 1.6-2.7 drag counts at the main design point. For the multipoint BRAC cases, the two optimization methodologies produced very similar section shapes, with a maximum performance difference of up to 3.9 drag counts at the high speed design point, and of 0.6-1.2 drag counts at the main design point.

When more section shape design variables are used in the Jetstream optimizations, performance is improved, particularly for the multipoint optimizations. However, part of this improvement is again due to the interaction of the geometry control and a thickness constraint, and revised geometric constraints may be needed with the increased geometric flexibility in order to avoid geometries that are too thin in places.

Both optimization methodologies are found to exhibit very similar optimization convergence behaviour, such that the cost, as characterized by the number of function evaluations to achieve a given objective improvement, is similar.

It is thus concluded that these different optimization methodologies using different flow solvers (based on the same

underlying equations), geometry control methodologies (with similar levels of geometric control), and optimization techniques are capable of finding optimal aerodynamic shapes which produce equivalent performance. This lends confidence to the use of these tools, as it gives confidence in the methodology-independence of aerodynamic shape optimization results obtained by different but consistent codes.

Acknowledgments

Financial support is provided in part by Bombardier Aerospace and the National Science and Engineering Research Council of Canada. Computations were performed on the Niagara 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; Ontario Research Fund - Research Excellence; and the University of Toronto.

References

- [1] Hicks, R. M. and Henne, P. A., “Wing Design by Numerical Optimization,” *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412. doi:10.2514/3.58379.
- [2] Jameson, A., “Aerodynamic Design via Control Theory,” *Journal of Scientific Computing*, Vol. 3, No. 3, 1988, pp. 233–260. doi:10.1007/BF01061285.
- [3] Epstein, B., Jameson, A., Peigin, S., Roman, D., Vassberg, J., and Harrison, N., “Comparative Study of 3D Wing Drag Minimization by Different Optimization Techniques,” *46th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA-2008-0326, 2008. doi:10.2514/6.2008-326.
- [4] Lee, C., Koo, D., Telidetzki, K., Buckley, H., Gagnon, H., and Zingg, D. W., “Aerodynamic Shape Optimization of Benchmark Problems Using Jetstream,” *53rd AIAA Aerospace Sciences Meeting*, AIAA-2015-0262, Kissimmee, FL, January 2015. doi:10.2514/6.2015-0262.
- [5] Koo, D. and Zingg, D. W., “Investigation into Aerodynamic Shape Optimization of Planar and Nonplanar Wings,” *AIAA Journal*, Vol. 56, No. 1, 2018, pp. 250–263. doi:10.2514/1.J055978.
- [6] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark,” *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985. doi:10.2514/1.J053318.
- [7] Poole, D. J., Allen, C. B., and Rendall, T., “Control Point-Based Aerodynamic Shape Optimization Applied to AIAA ADODG Test Cases,” *53rd AIAA Aerospace Sciences Meeting*, AIAA-2015-1947, 2015. doi:10.2514/6.2015-1947.
- [8] Dumont, A., Méheut, M., Baumgärtner, D., and Bletzinger, K.-U., “Aerodynamic Shape Optimization Progress on ADODG Benchmark Problems Using the elsA Software,” *35th AIAA Applied Aerodynamics Conference*, AIAA-2015-4081, 2015. doi:10.2514/6.2017-4081.

- [9] Destarac, D., Carrier, G., Anderson, G. R., Nadarajah, S., Poole, D. J., Vassberg, J. C., and Zingg, D. W., "Example of a Pitfall in Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 56, No. 4, 2018, pp. 1532–1540. doi:10.2514/1.J056128.
- [10] LeDoux, S. T., Young, D. P., Fugal, S., Elliot, J., Huffman, W. P., and Melvin, R. G., "An Updated Study for the AIAA Aerodynamic Design Optimization Discussion Group Test Case-4," *53rd Aerospace Sciences Meeting*, AIAA-2015-1717, 2015. doi:10.2514/6.2015-1717.
- [11] Hicken, J. E. and Zingg, D. W., "Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement," *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 401–413. doi:10.2514/1.44033.
- [12] Osusky, M. and Zingg, D. W., "A Parallel Newton-Krylov-Schur Flow Solver for the Navier-Stokes Equations Discretized Using Summation-By-Parts Operators," *AIAA Journal*, Vol. 51, No. 12, 2013, pp. 2833–2851. doi:10.2514/1.J052487.
- [13] Gagnon, H. and Zingg, D. W., "Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 53, No. 7, 2015, pp. 2015–2026. doi:10.2514/1.J053575.
- [14] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *Society for Industrial Applied Mathematics Review*, Vol. 47, No. 1, 2005, pp. 99–131. doi:10.1137/S0036144504446096.
- [15] Osusky, L., Buckley, H., Reist, T., and Zingg, D. W., "Drag Minimization Based on the Navier-Stokes Equations Using a Newton-Krylov Approach," *AIAA Journal*, Vol. 53, No. 6, 2015, pp. 1555–1577. doi:10.2514/1.J053457.
- [16] Hicken, J. E. and Zingg, D. W., "A Parallel Newton-Krylov Solver for the Euler Equations Discretized Using Simultaneous Approximation Terms," *AIAA Journal*, Vol. 46, No. 11, 2008, pp. 2773–2786. doi:10.2514/1.34810.
- [17] Osusky, M., Boom, P. D., and Zingg, D. W., "Results from the Fifth AIAA Drag Prediction Workshop obtained with a parallel Newton-Krylov-Schur flow solver discretized using summation-by-parts operators," *31st AIAA Applied Aerodynamics Conference*, AIAA-2013-2511, San Diego, CA, June 2013. doi:10.2514/6.2013-2511.
- [18] Zingg, D. W., Nemec, M., and Pulliam, T. H., "A Comparative Evaluation of Genetic and Gradient-Based Algorithms Applied to Aerodynamic Optimization," *European Journal of Computational Mechanics*, Vol. 17, No. 1, 2008, pp. 103–126. doi:10.3166/remn.17.103-126.
- [19] Hicken, J. E. and Zingg, D. W., "A Simplified and Flexible Variant of GCROT for Solving Nonsymmetric Linear Systems," *SIAM Journal of Scientific Computing*, Vol. 32, No. 3, 2010, pp. 1672–1694. doi:10.1137/090754674.
- [20] Gagnon, H. and Zingg, D. W., "Euler-Equation-Based Drag Minimization of Unconventional Aircraft Configurations," *Journal of Aircraft*, Vol. 53, No. 5, 2016, pp. 1361–1371. doi:10.2514/1.C033591.
- [21] Reist, T. A. and Zingg, D. W., "High-Fidelity Aerodynamic Shape Optimization of a Lifting-Fuselage Concept for Regional Aircraft," *Journal of Aircraft*, Vol. 54, No. 3, 2017, pp. 1085–1097. doi:10.2514/1.C033798.
- [22] Reist, T. A., Zingg, D. W., Rakowitz, M., Potter, G., and Banerjee, S., "Multifidelity Optimization of Hybrid Wing-Body Aircraft with Stability and Control Requirements," *Journal of Aircraft*, Vol. 56, No. 2, 2019, pp. 442–456. doi:10.2514/1.C034703.

- [23] Mohamed, K., Laurendeau, E., and Nadarajah, S., "Implementation and Validation of Detached Eddy Simulation in the Bombardier Aerospace Navier-Stokes Flow Solver," *56th Canadian Air and Space Institute Aeronautics Conference*, Ottawa, ON, Canada, 2009.
- [24] Kulfan, B. M., "A Universal Parametric Geometry Representation Method — "CST"," *The 45th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA-2007-0062, Reno, NV, January 2007. doi:10.2514/6.2007-62.
- [25] Piperni, P. and Boudreau, J., "The Evolution of Structured Grid Generation at Bombardier Aerospace," *50th Canadian Air and Space Institute Annual General Meeting and Conference*, Montreal, QC, Canada, 2003.
- [26] Wächter, A. and Biegler, L. T., "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, Vol. 106, No. 1, 2006, pp. 25–57. doi:10.1007/s10107-004-0559-y.
- [27] Rivers, M. B., "NASA Common Research Model: A History and Future Plans," *AIAA Aviation 2019 Forum*, AIAA-2019-3725, 2019. doi:10.2514/6.2019-3725.
- [28] Buckley, H. P. and Zingg, D. W., "An Approach to Aerodynamic Design Through Numerical Optimization," *AIAA Journal*, Vol. 51, No. 8, 2013, pp. 1972–1981. doi:10.2514/1.J052268.