Aerodynamic Shape Optimization of an S-Duct Intake for a Boundary-Layer Ingesting Engine

Christopher Chiang*, David Koo†, and David W. Zingg‡
Institute for Aerospace Studies, University of Toronto, Toronto, Ontario, M3H 5T6, Canada

A high-fidelity aerodynamic shape optimization framework based on the Reynolds-averaged Navier-Stokes equations is applied to the optimization of a boundary-layer ingesting S-duct designed for embedded engines on a high-subsonic, unmanned flight vehicle. The optimizations initially target a cruise operating condition and are further extended to single-point and multipoint optimizations considering descent and climb. Two different composite objective functions are used. The first combines distortion and swirl at the fan interface plane as well as total pressure recovery, with user-defined weights for each objective, while the second involves pressure recovery, fan blade load variation, and fan blade incidence variation. Pareto fronts show the trade-offs between objectives. The results indicate that compared to the baseline geometry, a simultaneous improvement in all objectives contained in the composite objective function can be obtained, depending on the priorities of each objective pre-assigned by the user. It was also found that if swirl can be ignored, then fan-face distortion can be greatly reduced while simultaneously reducing total pressure loss in the S-duct. Similarly, fan blade load variation and fan blade incidence variation can be significantly reduced while reducing total pressure loss. Finally, the multipoint optimization results show that a single S-duct geometry can perform well during cruise, climb, and descent conditions.

I. Nomenclature

A_c = cross-sectional area

A_j = outlet circumferential ring area

A_{rings} = outlet total ring area

$\boldsymbol{b}^{(m)}$ = deformed B-spline control points

C_p = static pressure coefficient

*MASc Candidate, AIAA Student Member, chris.chiang@mail.utoronto.ca
†Fluid Dynamics Analyst, Pratt and Whitney Canada, dkoo15@gmail.com
‡Professor, University of Toronto Distinguished Professor of Computational Aerodynamics and Sustainable Aviation and Associate Fellow AIAA, dwz@utias.utoronto.ca

This work is associated with conference paper 2021-2468 presented virtually on August 2 at the 2021 AIAA AVIATION Forum.

dA_i = differential area element within ring

d_{H} = hydraulic diameter

D = total pressure distortion

D_0 = initial total pressure distortion

G = objective function gradient

$\mathcal{J}$ = objective function

$L_{\mathrm{avg,j}}$ = average blade loading within ring

L_{i} = local blade loading at element i

L_{var} = circumferential blade load variation

$L_{\mathrm{var,0}}$ = initial circumferential blade load variation

$\mathcal{M}$ = mesh residuals

M_{inlet} = inlet Mach number

M_∞ = freestream Mach number

$\mathbf{q}$ = non-dimensional flow variables

$p_{\mathrm{s,out}}$ = outlet static pressure

$p_{\mathrm{t,avg,j}}$ = average outlet total pressure within ring

$p_{\mathrm{t,i}}$ = local total pressure at element i

$p_{\mathrm{t,in}}$ = inlet total pressure

$p_{\mathrm{t,out}}$ = outlet total pressure

$p_{\mathrm{t,out,0}}$ = initial outlet total pressure

$p_{\mathrm{t,60^\circ}}$ = outlet total pressure in 60° sector

P_{w} = wetted perimeter

PR = pressure recovery

$\mathcal{R}$ = flow residuals

S = swirl distortion

S_0 = initial swirl distortion

U_θ = circumferential component of the local velocity vector

U_x = axial component of the local velocity vector

V_i = flow velocity relative to the fan blade at element i

$\boldsymbol{v}$ = geometric design variables

Greek

α, β = generic weights

α_{inc} = true fan blade angle of attack

α_{s} = swirl angle

ρ_i = local density at element i

$\phi_{\text{avg,j}}$ = average blade incidence within ring

ϕ_i = local blade incidence angle at element i

ϕ_{var} = circumferential blade incidence variation

$\phi_{\text{var,0}}$ = initial circumferential blade incidence variation

II. Introduction

Boundary-layer ingestion (BLI) is an aero-propulsive concept in which the propulsor ingests the incoming boundary layer developing over the airframe to improve overall system efficiency. By ingesting the lower momentum boundary layer, the propulsion system can potentially expend less power than that required by a conventional podded engine. Furthermore, re-energizing the aircraft wake enables less wasted kinetic energy, and thus lower losses, than a conventional propulsor [1]. The concept has long been known amongst the research community for its fuel saving potential [2, 3], where applications for marine propulsion in ships and torpedoes have been extensively studied [4]. Its resurgence in recent years is due to renewed interest in applying the concept to highly-integrated aeropropulsive systems for more fuel-efficient aircraft [5–7] and urban air mobility vehicles [8].

Embedded engines allow for the inlet to remain flush with the airframe to take advantage of BLI while also providing some degree of noise suppression. Intakes for embedded engines often take the shape of S-ducts. To avoid any ambiguity, the term "intake" refers to the system that delivers air to the engine while the term "inlet" refers to the entrance of the intake. The curvature in S-ducts diffuses the air faster than in conventional straight ducts, which leads to more compact designs and potential weight savings. This is also favourable for unmanned aerial and urban air mobility vehicles, as their total size is often determined by the propulsion system [9]. Moreover, the offset lowers the radar and infrared signatures for missiles, which reduces the possibility of detection and increases survivability [10]. With varying S-duct cross-sectional area and shape, the cross-sections can transition from rectangular, oval, semi-elliptical etc. at the entrance to the circular engine face. The centerline curvature of the duct can vary with different offsets and turning angles. S-duct geometries are typically described by the offset and length ratios, based on the engine face diameter, and the diffusion ratio, which is the ratio of the outlet to inlet areas. An additional aspect ratio is used for rectangular inlets to define its width to height ratio.

The main objective of air breathing intakes is to deliver air at conditions suitable for the engine fan or compressor, supplied at the rate of mass flow required. The fan-face conditions are often at a mismatch from freestream flow

conditions, and the efficiency of this transformation process drives the design. An efficient intake recovers a large portion of the freestream total pressure at the engine interface by minimizing total pressure losses. Total pressure recovery is generally considered to be directly proportional to engine thrust in subsonic flow conditions [11]. However, maximizing total pressure recovery alone is not typically sufficient for a good design. The total pressure field at the fan face not only needs to be high in an average sense, but it also needs to be uniform. Although a lower average in total pressure recovery can lead to thrust and efficiency penalties, distortions in the total pressure field can significantly reduce engine life in two ways. The first is through fan blade fatigue due to cyclical stresses caused by fan vibrations. Noise is often increased as a result of the vibrations. Second, distortion can reduce the surge margin of the compressor in highly loaded sections, which compromises engine stability [12]. Since the 1950s, it has been known that intakes with very high average pressure recoveries were not robust. In fact, they could be sensitive to off-design conditions that yield higher distortion values over a range of operation than a lower recovery intake [13]. For these reasons, total pressure distortion is introduced as another objective into the design problem.

Several documented instances [14–17] suggest that the bi-objective design problem of total pressure recovery and distortion may not be adequate in capturing the accuracy of intake performance over the entire range of operating conditions. Swirl is a secondary flow phenomenon that describes the flow angularity in the intake. Wind tunnel measurements of swirl in an S-duct conducted by Guo and Seddon [18] showed that at high incidence, a large vortex forms around an area of flow separation after the first bend and a pair of counter-rotating vortices in the flow appear after the second bend. Historically, the use of inlet guide vanes placed upstream of the fan interface plane avoided many operability problems that did not explicitly account for swirl distortion [19]. Nevertheless, a large amount of swirl can significantly impinge on the incidence angle of inlet guide vanes and induce a cascade effect on the downstream fan and compressors [20]. Depending on the design criteria, these flow control devices are also not always practical and hence cannot be deployed. Intakes without inlet guide vanes could induce high levels of swirl distortion, leaving the engine vulnerable to compressor stalls and other operability problems [21].

Researchers have targeted the reduction of total pressure losses while counteracting total pressure and swirl distortion through active, passive, and hybrid flow control as well as shape optimization techniques. Several experimental studies [22–26] focus on the characterization and influence of inlet distortion profiles, as well as the effectiveness of flow control to mitigate distortion. Concurrently, computational work in the form of steady and unsteady numerical simulations on S-ducts has sought to achieve similar goals [27–29]. A detailed literature review of flow control techniques and optimization in S-ducts is given by Wojewodka et al. [30].

The focus here will be on shape control that deforms the duct walls. Some previous work in the literature has had varying degrees of success in this respect. A commonly selected S-duct for optimization is the Royal Aircraft Establishment M2129, introduced by researchers at British Aerospace Industries in the early 1980s [31, 32]. It has offset, length, and diffusion ratios of 0.9, 3, and 1.52, respectively. An S-duct optimization studied by Zhang et al. [33]

featured an irregular bump on the upper wall after the first bend. The intent is to introduce a separated flow region far upstream that helps reduce cross-flow downstream of the duct. The two design parameters were the height and width of the surface perturbation. The objective function was the total pressure distortion at the engine interface. The inflow Mach number was 0.6, the Reynolds number based on the inlet diameter was 2.6 million, and the BLI height was 7.3% of the inlet diameter. Using a gradient-based optimization algorithm, distortion was reduced by 63% compared to the baseline M2129 duct. Pressure recovery remained the same, as the bump behaved like a closed bubble and the separation was confined to that region.

With the same M2129 duct, D'Ambros et al. [34] and Chiereghin et al. [35] both used the Free Form Deformation (FFD) method [36] to control the S-duct geometry but do not consider BLI. The flow conditions were consistent with those of Zhang et al. [33]. Gradient-free genetic algorithms were used to search the design space. Considering pressure recovery and swirl as the two objectives, they were able to achieve considerable improvements over the baseline. Both were able to show parts of a convex Pareto curve that demonstrated the trade-off between pressure recovery and swirl. Lee and Kim [37] used an adjoint method approach to optimize for two separate single objectives at a throat Mach number near 0.85 and a Reynolds number around 7.7 million based on the throat diameter. The first distortion minimization resulted in a 25% reduction from the baseline and the total pressure recovery increased by about 1%. The second optimization was a total pressure recovery maximization, which improved by 1.5% and decreased distortion by 3-7%. They also examined the performance of the design from the second case at several off-design conditions by increasing the Mach number. The optimized duct was able to outperform the baseline in both distortion and pressure recovery at all of the off-design conditions surveyed.

Koo et al. [38] developed an aerodynamic shape optimization approach for a fixed rectangular S-duct intake with boundary-layer ingestion, intended for use on a high-subsonic, unmanned flight vehicle. The flight condition was a cruise Mach number of 0.7 and BLI that was 60% of the inlet height. The study was based on steady-state Reynolds-Averaged Navier-Stokes (RANS) simulations, and the geometry was controlled with an FFD volume. The goal of the optimization was to simultaneously maximize total pressure recovery and minimize distortion at the outlet, and they were able to achieve substantial improvements in both objectives. A convex Pareto curve was obtained showing the competing nature between total pressure recovery and distortion. However, it was observed that the swirl was increased in the optimized ducts compared to the baseline. This suggests that considering total pressure recovery and distortion may not be sufficient, and motivates the addition of swirl into the composite objective function.

The objective of this paper is to build on the work of Koo et al. [38] and adapt the existing shape optimization methodology to investigate two different composite objective functions that pertain to the aerodynamic performance of the S-duct. The first is a composite objective function composed of distortion and swirl taken at the fan interface plane and total pressure recovery. The second is a combination of pressure recovery, blade load variation, and blade incidence variation at various fan speeds.

At the cruise design point, Pareto fronts are generated with respect to each composite objective function to provide insight into the trade-offs when it comes to prioritizing objectives. We then select some Pareto optimal points from the front to further evaluate their performance and analyze the resulting internal flow field. Lastly, we extend this optimization methodology to include descent and climb conditions to investigate the robustness of the design and to evaluate the aerodynamic performance of a multipoint design compared to single-point optimized designs.

The remainder of the paper is organized as follows. Section III provides a brief overview of the aerodynamic shape optimization framework, along with the objectives and composite objective functions. Section IV describes preliminary studies that were completed to refine the optimization methodology. Section V presents the results for two separate composite objective functions, including the Pareto front characterization, single-point optimizations, and multipoint optimizations. Conclusions are drawn in Section VI.

III. Methodology

The methodology for aerodynamic shape optimization can be broken down into five main components: geometry parameterization and control, mesh deformation, flow solver, gradient computation, and the optimization algorithm. The optimization problem is formulated as

$$\begin{aligned}
\min \quad & \mathcal{J} = (\boldsymbol{v}, \boldsymbol{b}^{(m)}, \mathbf{q}, p_{\text{s,out}}), && \text{(1a)}\\
\text{w.r.t.} \quad & \boldsymbol{v}, && \text{(1b)}\\
\text{s.t.} \quad & \boldsymbol{\mathcal{M}}^{(i)}(\boldsymbol{v}, \boldsymbol{b}^{(i)}) = 0,\ i = 1, 2, ..., m, && \text{(1c)}\\
& \boldsymbol{\mathcal{R}}(\boldsymbol{v}, \boldsymbol{b}^{(m)}, \boldsymbol{q}, p_{\text{s,out}}) = 0, && \text{(1d)}\\
\text{and} \quad & M_{\text{inlet}}(\boldsymbol{v}, \boldsymbol{b}^{(m)}, \boldsymbol{q}, p_{\text{s,out}}) = M_{\infty}, && \text{(1e)}
\end{aligned}$$

where $\mathcal{J}$ is the objective function, $\boldsymbol{v}$ is the vector of design variables, $\boldsymbol{b}^{(m)}$ are the deformed B-spline control points, $\mathbf{q}$ are the flow variables, and $p_{\text{s,out}}$ is the outlet static pressure. The first constraint is the mesh deformation residual $\boldsymbol{\mathcal{M}}$ at increment i, which is a set of m linear equations that are solved to obtain the B-spline volume control points that control the mesh at the i^{th} increment. The second is the flow residual $\boldsymbol{\mathcal{R}}$, which represents the discretized steady RANS equations that must be solved in order to obtain a converged solution for the non-dimensional flow variables $\mathbf{q}$. Both residuals are theoretically driven to zero. The last constraint is to ensure the area-averaged inlet Mach number satisfies the specified operating condition. This is done iteratively by adjusting $p_{\text{s,out}}$ using the Newton method.

A. Geometry Parameterization and Control

The baseline geometry is provided by Asghar et al. [39]. The main features of the duct are its length-to-diameter ratio of 2.703, offset-to-diameter ratio of 1, diffusion ratio of 1.57, and inlet aspect ratio of 1.5, which are all held

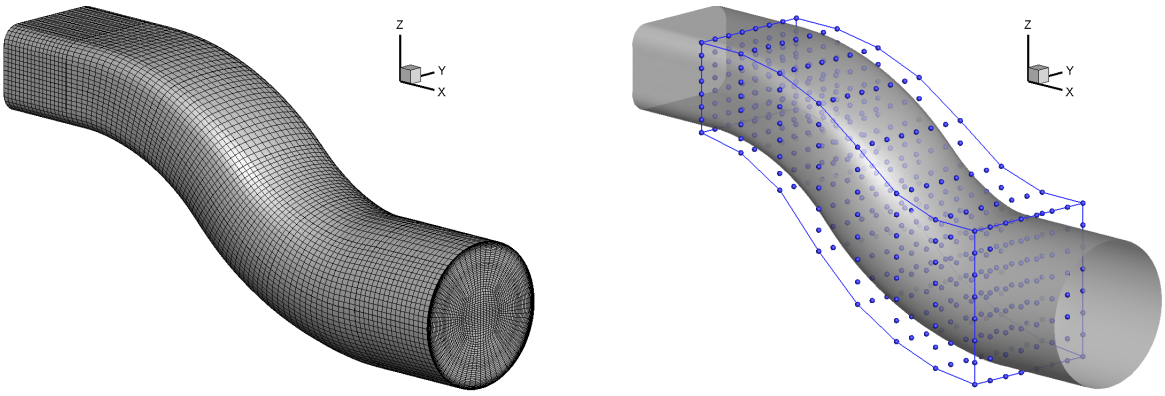

Fig. 1 Computational grid and FFD control volume for the S-duct.

constant throughout the optimization. As shown on the left of Figure 1, the computational domain is discretized into a multi-block structured, O-topology grid. To parameterize the geometry and control the grid deformation, the grid is fitted to a B-spline control volume using spatially-varying knot vectors. The B-spline control points on the duct surface are then embedded inside an FFD lattice, as shown on the right of Figure 1. These lattice points become the design variables that drive the deformation of the embedded surface while maintaining its connectivity. For this work, an FFD with 8 cross-sections, each with an 8-by-8 array of control points was found to provide adequate geometric freedom to the optimizer. This was determined from a study that is detailed in Section IV.B.1.

The hydraulic diameter is commonly used as a dimensionless quantity to describe the geometry in non-circular channel and duct flows. It is defined as

$$d_{\mathrm{H}} = \frac{4Ac}{P_{\mathrm{w}}}, \tag{2}$$

where A_{c} is the cross-sectional area of the duct and P_{w} is the wetted perimeter. Control points are restricted to stay within 10% of the hydraulic diameter at their respective cross-section. Further exploration of this constraint is discussed in Section IV.B.2. Linear constraints are set to enforce geometric symmetry about the vertical xz-plane. Details of the geometry control system are given in Gagnon and Zingg [40].

B. Mesh Deformation

Once the design variables have been manipulated by the optimizer, the embedded B-spline surface points are deformed accordingly. Changes in the FFD shape apply to the B-spline surface control points, not directly to the nodes of the computational mesh. The sensitivity of the aerodynamic surface has an exact mapping to the sensitivity with respect to the FFD control points. Surface displacements are propagated to the rest of the B-spline volume control mesh through the use of a linear-elasticity model [41, 42]. The mesh deformation equations are divided into a number of linear increments, which are solved using the preconditioned conjugate gradient method. Once the volume control mesh has been deformed, the computational grid is recomputed algebraically based on the B-spline mapping and the

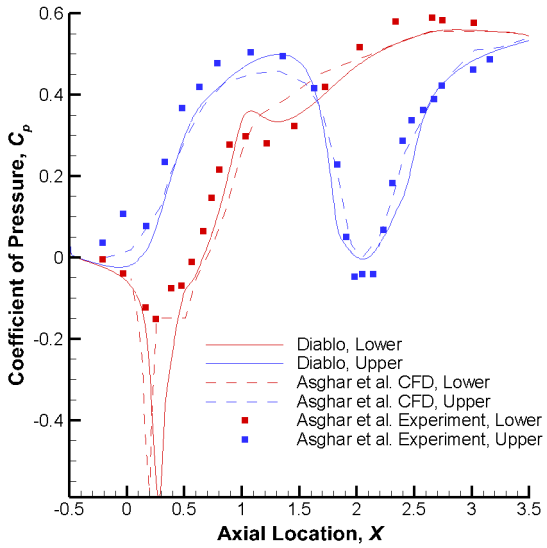

Fig. 2 Diablo results for the pressure coefficient [38] compared to numerical and experimental results from Asghar et al. [39] at an inlet Mach number of 0.8 without BLI.

parametric coordinates of the nodes. Thus, the analytical representation of the geometry is preserved.

C. Flow Solver

Diablo is a parallel, implicit, multi-block structured finite-difference solver that is used to solve the RANS equations, fully coupled with the Spalart-Allmaras turbulence model with quadratic constitutive relations [43, 44]. Second-order summation-by-parts operators are used for spatial discretization, and boundary and block interface conditions are enforced weakly through simultaneous approximation terms [45]. The steady RANS equations are solved in two stages - an approximate-Newton start-up phase for globalization, followed by an inexact-Newton phase for deep convergence. The flexible generalized minimal residual (GMRES) method with an approximate-Schur parallel preconditioner is used to solve the linear system at each Newton iteration. Additional details can be found in Osusky and Zingg [43].

For internal flows, inlet total pressure and outlet static pressure boundary conditions were implemented. In order to simulate boundary-layer ingestion, the turbulent flow over a two-dimensional flat plate is first solved for the desired boundary-layer height relative to the inlet height. The results of this simulation are then used to extract total pressure and turbulent viscosity profiles, which are applied to the nodes on the inlet face of the S-duct grid.

Figure 2 displays a plot of the pressure coefficient on the upper and lower surfaces of the S-duct, compared to both CFD and experimental results. The experimental and CFD results for the baseline geometry from Asghar et al. [39] were used to validate the flow solver for this class of problem. Note that this validation is at an inlet Mach number of 0.8 and does not consider BLI. There is reasonable agreement, indicating that the Spalart-Allmaras turbulence model is adequate for this class of flows, consistent with the conclusions of Yaras and Grosvenor [46].

D. Gradient Computation

Due to the large computational cost associated with evaluating the flow equations, a gradient-based optimizer is used, as they typically require fewer function evaluations than genetic algorithms [47]. However, aerodynamic shape optimization problems have been known to be multimodal, and the local minima found are not guaranteed to be global minima [48, 49]. The gradients of the objective and constraints are evaluated using the discrete-adjoint method. The method allows for the computational cost of each gradient evaluation to be independent of the number of design variables, which is advantageous for this class of problem with many design variables. The introduction of Lagrange multipliers leads to the flow, inlet Mach, and mesh adjoint equations respectively [50]. The adjoint equations are solved using the augmented adjoint approach used by Rashad and Zingg [51]. Additional details of the gradient computation can be found in Osusky et al. [52].

E. Optimization Framework

The Sparse Nonlinear Optimizer (SNOPT) package is used to solve the optimization problem and provide updated design variables. The mesh is then deformed, followed by the flow, adjoint, and gradient evaluations. To complete the loop, the gradient is fed back in to SNOPT for the next design iteration. Unless otherwise stated, the optimization is consistently initialized from the baseline design. The SNOPT algorithm is capable of handling linear and nonlinear constraints and will satisfy linear constraints exactly. The details of SNOPT are further described by Gill et al. [53].

F. S-Duct Performance Metrics

Intake performance is typically evaluated based on the total pressure recovery ratio that describes the total pressure loss in the duct. Total pressure recovery is defined as

$$PR = \frac{p_{\mathrm{t,out}}}{p_{\mathrm{t,in}}}, \tag{3}$$

which is an efficiency metric that should be maximized. Circumferential flow distortion is also significant since it can negatively impact engine performance. Circumferential flow distortion at the outlet is calculated as

$$D = \sum_{j}^{m} \sum_{i}^{n_j} \frac{(p_{\mathrm{t,i}} - p_{\mathrm{t,avg,j}})^2 dA_i}{A_j}. \tag{4}$$

The outlet plane is partitioned into $j = 1 \ldots m$ rings, each with n_j discrete elements. Within each ring, the area-weighted sum of squared error of total pressure is calculated. The total pressure is non-dimensionalized by the free stream dynamic pressure, $\rho_\infty (a_\infty)^2$. The reference velocity u_∞ is replaced by the reference speed of sound, since u_∞ could be zero for some flow problems.

Intakes can also generate significant flow angularity in secondary flows, otherwise known as swirl distortion. In this

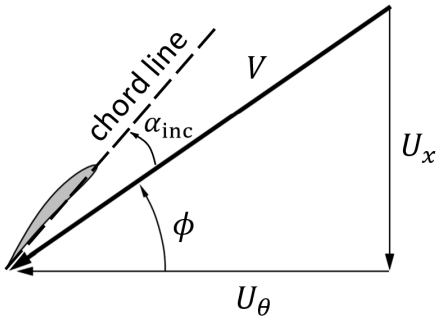

Fig. 3 Resultant velocity triangle from radial blade element discretization.

work, the swirl is defined as the angle (in radians) between the circumferential component of the local velocity vector and the axial vector normal to the outlet. The swirl angle is calculated as

$$\alpha_{\text{s}} = \arctan\left(\frac{U_\theta}{U_x}\right). \tag{5}$$

The swirl is integrated over the outlet plane, calculated as

$$S = \frac{1}{A_{\text{rings}}} \sum_{j}^{m} \sum_{i}^{n_j} \left| \alpha_{\text{s,i}} \right| dA_i, \tag{6}$$

with the same partitioning of rings and discrete elements that are used in Eq. 4. The magnitude of the swirl angle is used as opposed to the directionality since it is a value more suitable for optimization.

Alternative performance metrics are derived to describe the circumferential variation of fan blade loading (i.e. lift) and angle of incidence that a single fan blade would experience through one rotation. Minimizing blade load variation is intended to ease the circumferential loading of the fan blade in light of aerostructural concerns. Minimizing blade incidence variation enables the design of more optimal blade airfoil sections. From blade element theory, the fan blade is radially discretized into independent 2D sections. The resultant velocity triangle in Figure 3 shows the axial, tangential, and relative induced velocity components. For this objective function the fan speed must be specified. If the fan speed is variable, multi-point optimization can be applied, but we consider only a single fan speed here. Moreover, the effect of the fan on the flow in the duct is ignored in this study. Representative fan speeds for this application are modeled after the Williams F107 small turbofan engine, with an upper limit of approximately 35,000 RPM [54]. No fan blades are modeled and no source terms are used in the simulation; only the tangential velocity contribution of the blade is captured in the calculations of relative angle ϕ and relative velocity V.

The local blade load variation at each radial point includes only the portion of the lift per unit area that varies circumferentially and excludes the contribution from constants. Similarly, the local blade incidence variation considers the variation in the relative angle of attack since the blade metal angle is constant. The relative angle ϕ (in radians)

is used to describe the variation in blade incidence instead of α_{inc} as the variation of both angles is the same. These variations are formally defined at each element i as

$$L_{\text{i}} = \rho_i {V_{\text{i}}}^2 \phi_{\text{i}} \tag{7}$$

and

$$\phi_{\text{i}} = \arctan\left(\frac{U_x}{U_\theta}\right), \tag{8}$$

respectively. Similar to total pressure distortion, to form a scalar objective, the circumferential variation in blade loading is written as

$$L_{\text{var}} = \sum_j^m \sum_i^{n_j} \frac{(L_{\text{i}} - L_{\text{avg,j}})^2 dA_i}{A_j}. \tag{9}$$

The circumferential variation in blade incidence is written as

$$\phi_{\text{var}} = \sum_j^m \sum_i^{n_j} \frac{(\phi_{\text{i}} - \phi_{\text{avg,j}})^2 dA_i}{A_j}. \tag{10}$$

G. Objective Functions

For multi-objective optimization, the weighted-sum method is used due to its simplicity and ease of implementation. Two separate composite objective functions are considered in this work. A tri-objective function with total pressure recovery, distortion, and swirl is formed such that

$$\mathcal{J}_1 = \alpha\beta \frac{P_{\text{t,out,0}}}{P_{\text{t,out}}} + (1-\beta)\frac{D}{D_0} + \beta(1-\alpha)\frac{S}{S_0}, \tag{11}$$

where the relative weights between all three objectives are determined by α and β. The subscript 0 is the initial value used to normalize each objective, which is typically the baseline value. Since the inlet shape and total pressure are fixed, minimizing the first term is equivalent to minimizing the total pressure losses at the outlet. Note that with special values of α and β, several single- and bi-objective functions can be obtained. With $\alpha = 0$, pressure recovery does not factor in to the objective and the relative weighting between distortion and swirl is determined entirely by β. Similarly, an α value of 1 yields a bi-objective function with pressure recovery and distortion. Setting β equal to 1 gives a bi-objective function with pressure recovery and swirl, while a β value of 0 provides a single-objective function to minimize distortion. Setting α and β to 1 solely maximizes pressure recovery, while the combination of $\alpha = 0$ and $\beta = 1$ gives a single-objective minimization of swirl.

The second composite objective function seeks to maximize pressure recovery while simultaneously minimizing the

Table 1 Operating points.

Condition	Reynolds No.	Mach No.	Boundary-Layer Height	Fan RPM
Cruise	1×10^6	0.7	60% inlet height	35.0×10^3
Descent	1×10^6	0.6	60% inlet height	31.5×10^3
Climb	5×10^5	0.3	80% inlet height	38.5×10^3

circumferential variation of blade loading and blade incidence, written as

$$\mathcal{J}_2 = \alpha\beta\frac{P_{\text{t,out},0}}{P_{\text{t,out}}} + (1-\beta)\frac{L_{\text{var}}}{L_{\text{var},0}} + \beta(1-\alpha)\frac{\phi_{\text{var}}}{\phi_{\text{var},0}}. \tag{12}$$

A set of points on a Pareto front can be generated by varying the relative weights between 0 and 1. Pareto fronts are useful in determining the trade-off between competing objectives and provide a family of solutions that can be selected depending on which objectives are deemed more important by the designer. Although this approach can have difficulty with concave fronts, for example, this was not an issue here.

Next, objective functions for multipoint optimizations are discussed. Table 1 summarizes the three operating points considered. Two objective functions are formulated for multipoint optimization. The first is written as

$$\mathcal{J}_{\text{mp1}} = 0.5\mathcal{J}_{\text{cruise}} + 0.25\mathcal{J}_{\text{descent}} + 0.25\mathcal{J}_{\text{climb}}, \tag{13}$$

where cruise is prioritized twice as much as descent and climb. The second is written as

$$\mathcal{J}_{\text{mp2}} = 0.75\mathcal{J}_{\text{cruise}} + 0.125\mathcal{J}_{\text{descent}} + 0.125\mathcal{J}_{\text{climb}}, \tag{14}$$

where cruise is weighted six times as much as descent and climb. These weights were chosen under the assumption that, in most BLI applications, cruise is the dominant design condition. Note that each composite function with operating conditions at cruise, descent, and climb on the right-hand side of Eq. 13 and Eq. 14 has its own initial values for normalizing the objectives.

IV. Preliminary Studies

The S-duct is initially designed for the cruise operating condition. Figure 4 shows the baseline performance at cruise, which demonstrates the non-uniform distribution of high and low total pressure zones at the outlet. The non-uniform circumferential distributions of relative blade loading and relative blade incidence are also shown. As expected, the region of lower relative blade loading and lower relative blade incidence corresponds to the low total pressure region and the lower axial velocity of the incoming flow. To the right, the streamlines near the boundary layer indicate the presence

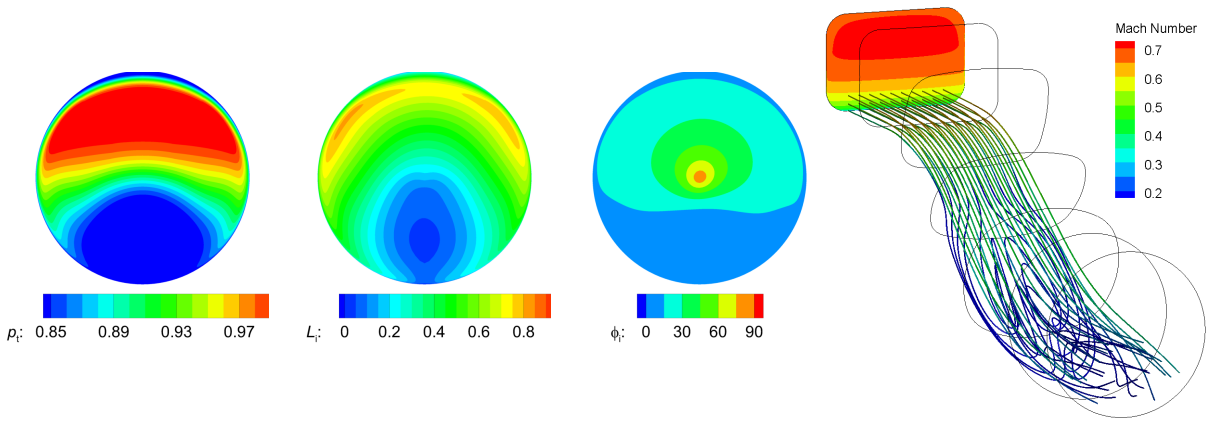

Fig. 4 Baseline S-duct performance at cruise. Total pressure (left), relative fan blade loading (middle-left), and relative fan blade incidence angle in degrees, (middle-right) at the outlet. Streamlines near the boundary layer (right).

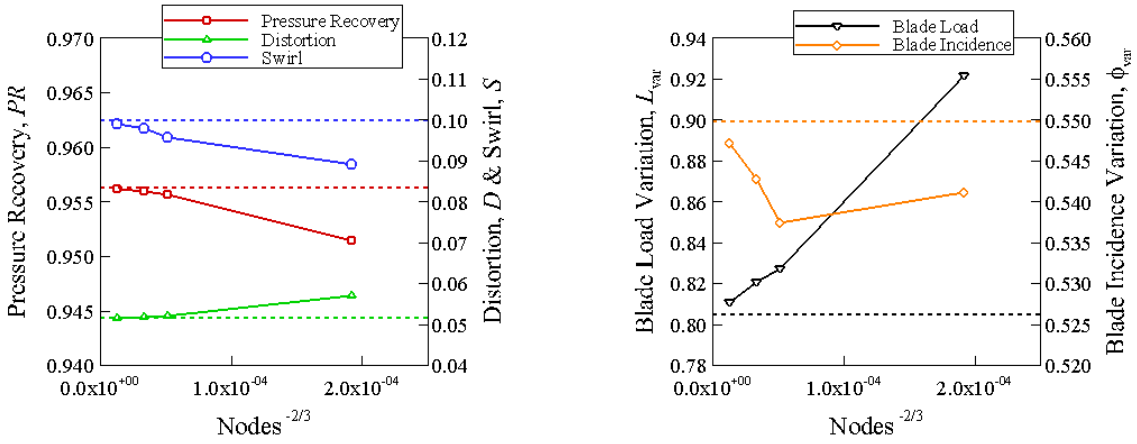

Fig. 5 Grid convergence for the baseline geometry. Dashed lines indicate the Richardson extrapolated values.

of secondary flows that generate a pair of counter-rotating vortices downstream of the S-duct at the second bend.

A. Grid Density

To evaluate the sensitivity of the S-duct performance metrics to grid refinement, a grid density study is completed for the baseline geometry using four grids that range in size from 377,000 to 21 million nodes. The results are plotted in Figure 5, along with Richardson extrapolated values. It can be noted that the blade incidence variation does not show monotonic convergence. This may be due to the inability of the coarsest grid to capture the blade incidence variation accurately due to insufficient discrete elements within a ring that could produce error cancellation from poorly resolved values. The error for pressure recovery is less than one percent on all grids, while the error for blade incidence variation is less than five percent on all grids. The errors are much larger for distortion, swirl, and blade load variation, as much as 18% using the coarsest grid. Thus, for optimizations we find that a reasonable balance between accuracy and

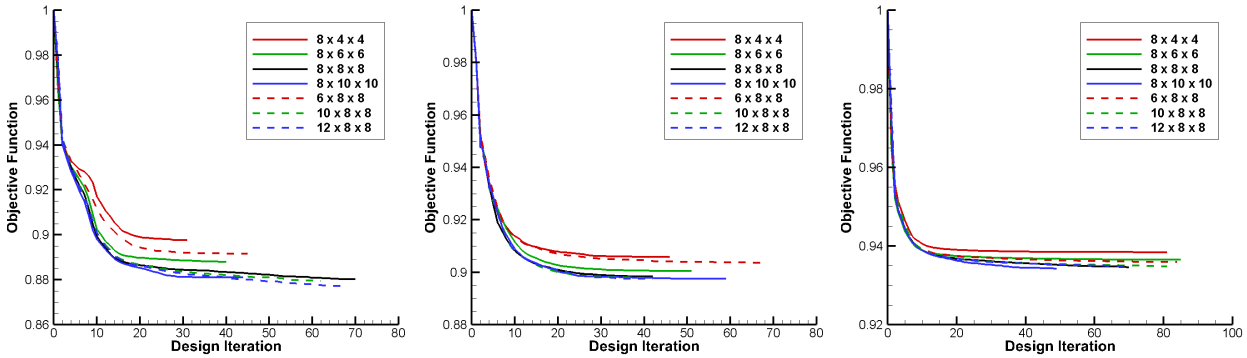

(a) $\alpha = 0.95$ and $\beta = 0.90$ (PR, D, and S).

(b) $\alpha = 1.00$ and $\beta = 0.90$ (PR and L_{var}).

(c) $\alpha = 0.95$ and $\beta = 1.00$ (PR and ϕ_{var}).

Fig. 6 Optimization convergence of FFD with m cross-sections and n-by-n control points denoted as m x n x n.

computational cost is to use the 2.7 million node grid, which is the second point from the right in Figure 5.

B. Geometric Freedom Studies

The purpose of these studies is to explore the sensitivity of the optimization to the geometric freedom allowed by the FFD control system and determine a set of preferred choices. Here, the coarsest grid was used for fast turnaround times and to identify relative trends from the optimization response to input changes to the FFD volume definition.

1. Free-Form Deformation Volume Definition

The number of cross-sections and the number of control points within each cross-section of the FFD volume are investigated to determine a set that provides an adequate amount of geometric freedom to the optimizer. Three cases are considered, with the first using the composite objective function with pressure recovery, distortion, and swirl with $\alpha = 0.95$ and $\beta = 0.90$. The next case uses the composite objective function with pressure recovery and blade load variation with $\alpha = 1.00$ and $\beta = 0.90$. The final case considers pressure recovery and blade incidence variation with $\alpha = 0.95$ and $\beta = 1.00$. For each case, the number of cross-sections varies from 6 to 12, while the number of control points at each cross-section is varied from 4-by-4 up to 10-by-10.

Figure 6 shows the results of the optimization runs for each case. Based on the three plots, there is no significant improvement gained by adding either more axial control sections or more cross-sectional control points past the 8 x 8 x 8 case. Hence, an FFD volume with 8 cross-sections and 8-by-8 control points in each cross-section was selected for all subsequent studies described in this paper.

2. Geometric Linear Box Constraints

This section investigates the geometric freedom at each control point of the FFD, which are the design variables of the optimization. The design variables are bounded by a box within their respective cross-section. The bounding boxes

Table 2 Optimized performance values for pressure recovery, distortion, and swirl with $\alpha = 0.80$ and $\beta = 0.80$ at cruise and different amounts of geometric freedom.

Bounds	J	PR	$D \times 10^2$	$S \times 10^2$
$0.05d_{\mathrm{H}}$	0.7528	0.9635	3.028	0.297
$0.1d_{\mathrm{H}}$	0.6725	0.9536	0.503	0.329
$0.2d_{\mathrm{H}}$	0.6589	0.9426	0.260	0.095

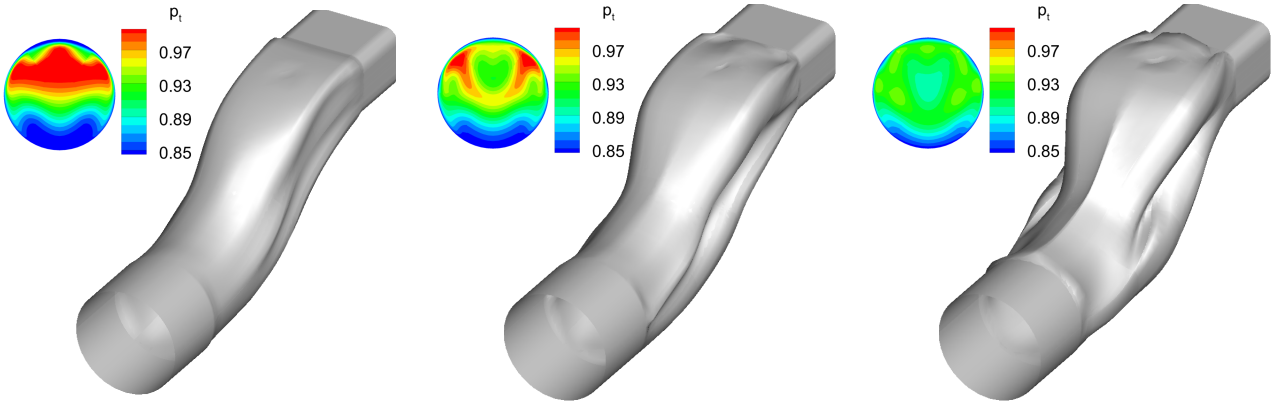

(a) Design variable bounds at $0.05d_{\mathrm{H}}$. (b) Design variable bounds at $0.1d_{\mathrm{H}}$. (c) Design variable bounds at $0.2d_{\mathrm{H}}$.

Fig. 7 Optimized geometry and outlet total pressure for pressure recovery, distortion, and swirl with $\alpha = 0.80$ and $\beta = 0.80$ at cruise and different amounts of geometric freedom.

are constrained as a percentage of the hydraulic diameter of the cross-section. Following the results of the previous FFD volume definition study, an FFD consisting of 8 cross-sections with 8-by-8 control points is used here. For this study, we consider only the composite objective function with pressure recovery, distortion, and swirl with α and β values of 0.80.

The objective function value and the optimized performance values are shown in Table 2 for the set of three design variable bounds investigated. There is further improvement in the objective function as the bounds are increased. Given this pair of α and β weights, the optimizer trades pressure recovery for lower amounts of distortion and swirl with the increased geometric freedom.

Figure 7 presents the three examples of the optimized geometries with various box constraint values. With Figure 7a, the optimizer was clearly not given enough design variable freedom. Given the limited design space, it produced a geometry that still displayed pre-existing flow features from the baseline, namely the regions of high total pressure in the top half of the outlet and the low total pressure region in the bottom half. As the design variable bounds are increased, the optimizer is able to better achieve its objective. In Figure 7c, the geometrical features are much more pronounced. From its total pressure contour, the distortion appears quite favourable compared to the other two shown. We see that the optimized geometries share some similar geometric features. The trend of the upper surface is to expand at the first bend, then form a bottleneck at the second bend. The bottom surface is also gradually expanded as much as possible.

Despite further improvements in S-duct performance with design variable bounds of 0.2 d_{H}, the increased geometric freedom leads to convergence problems with the optimization algorithm. As the design variable bounds are extended, the bounds will inevitably intersect and encroach on the domains of its neighbouring control points. This allows the possibility for the optimizer to displace a control point past its neighbour. Thus during mesh movement, the surface may fold in on itself, which triggers a failure. As a result, the optimization with 0.2 d_{H} experienced failures in the mesh movement algorithm after about 30 design iterations that hindered the optimization progress.

Furthermore, real-world manufacturability constraints could also have a considerable impact on the design variable bounds and the realizability of optimized geometries. For these reasons, a bound nominally set to 10% of the hydraulic diameter is conservatively selected, as it avoids most of the optimization convergence issues, yet is still able to achieve desirable S-duct performance with a reasonable amount of geometric deformations. Future work could look to increase the design variable bounds to 0.2 d_{H} to further investigate the capability of the optimization algorithm and whether the aerodynamic performance is robust, with consideration for the manufacturability of the optimized geometries. Alternatively, additional constraints could be implemented in the future to address these concerns and allow for greater freedom to the design variable bounds. Experimental testing of the optimized geometries should also be beneficial for validation and to verify that the optimizer is not going beyond the accuracy of the models used.

V. Results

A. Pareto Front Characterization

The Pareto front for the composite objective function with total pressure recovery, distortion, and swirl is shown in Figure 8a with boundaries at α values of 0.0 and 1.0 and a β value of 1.0. These boundaries are the three combinations of bi-objective functions resulting from Eq. 11. A rough sampling of relatively distributed points was chosen to get a general idea of the shape of the Pareto front and how the optimization process behaves with varying relative weights. The three-dimensional Pareto front shows that the surface is entirely convex and that all three objectives are competing.

The trade-off between pressure recovery and distortion is plotted in Figure 8b. The other trade-off between pressure recovery and swirl is shown in Figure 8c. Although not shown explicitly, the trade-off between distortion and swirl can be seen from the black curve. For better visual presentation, the baseline performance is plotted as a red line rather than a point in Figure 8b since its distortion far exceeds any result encountered. These plots demonstrate that with carefully chosen relative weights, the optimizer is able to simultaneously improve pressure recovery, distortion, and swirl. However, as one moves away from the elbow region of the front, small improvements in one objective lead to large penalties in the other.

Figure 9 plots the Pareto curve for the other composite objective function considering total pressure recovery, blade load variation, and blade incidence variation. Again, the baseline geometry is represented with a red line rather than

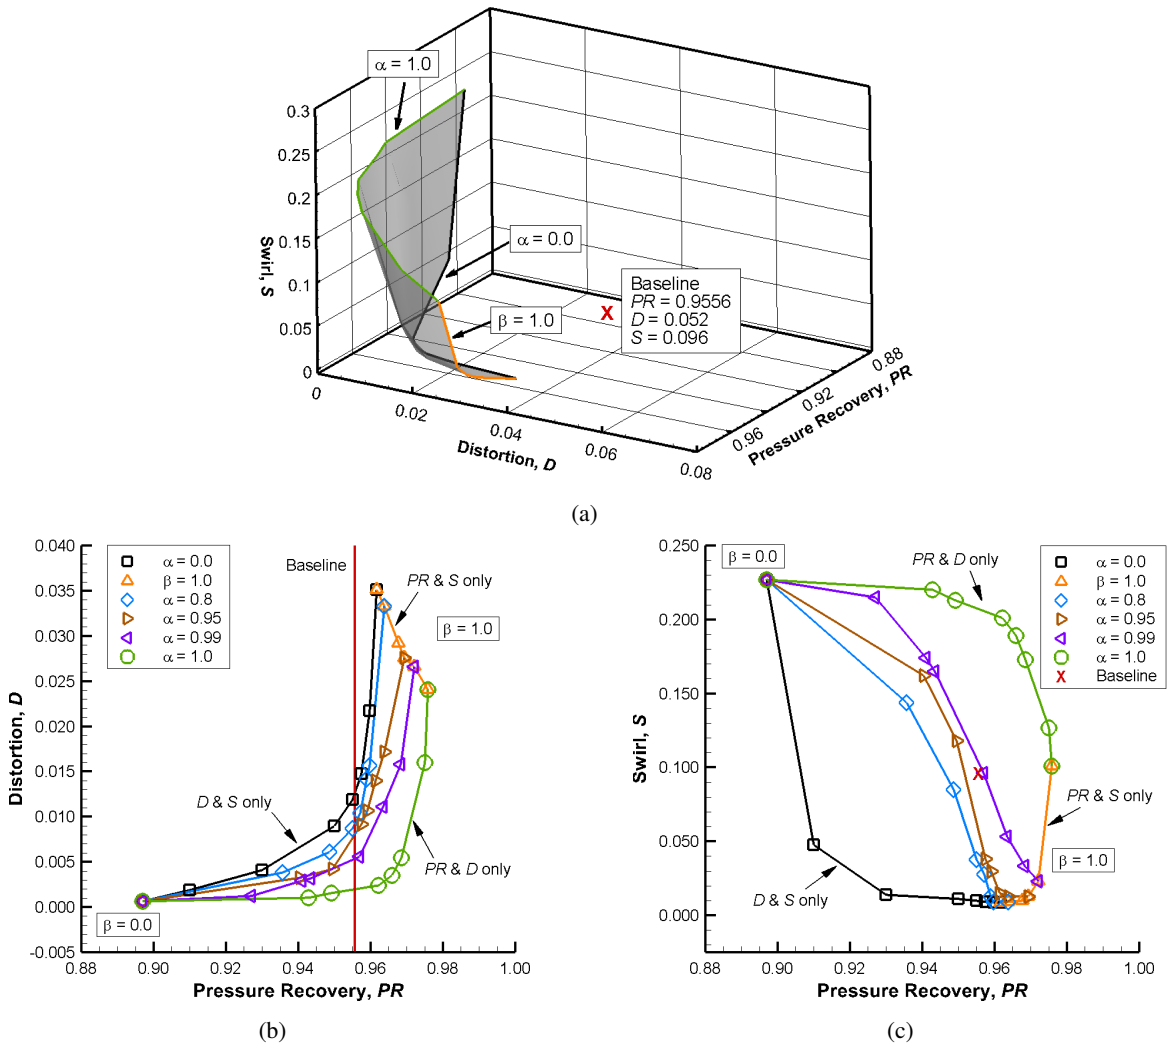

Fig. 8 Pareto front for the tri-objective function with pressure recovery, distortion, and swirl at cruise.

a point, as its variation in blade loading and blade incidence exceeds the axis range. The first thing to note is that a three-dimensional view is not necessary since the Pareto curve for pressure recovery and blade load variation, which is the black line with $\alpha = 1.00$ and varying β values, is nearly coincident to the Pareto curve for pressure recovery and blade incidence variation, which is the blue line with $\beta = 1.00$ and varying α values. Hence, these curves adequately demonstrate the convex trade-off between the objectives considered and no further combinations of relative weights are required.

In Figure 9a, the optimizer was able to considerably reduce the blade load variation across the entire range of β values. Values larger than $\beta = 0.95$ (third black circle from the right) or smaller than $\beta = 0.75$ (fourth black circle from the left) lead to diminishing returns in one objective with large penalties in the other. Thus, β values ranging from 0.75 to 0.95 are of interest since they provide substantial reduction in blade load variation while outperforming the baseline pressure recovery. Similarly, in Figure 9b the blade incidence variation is considerably reduced for all α values

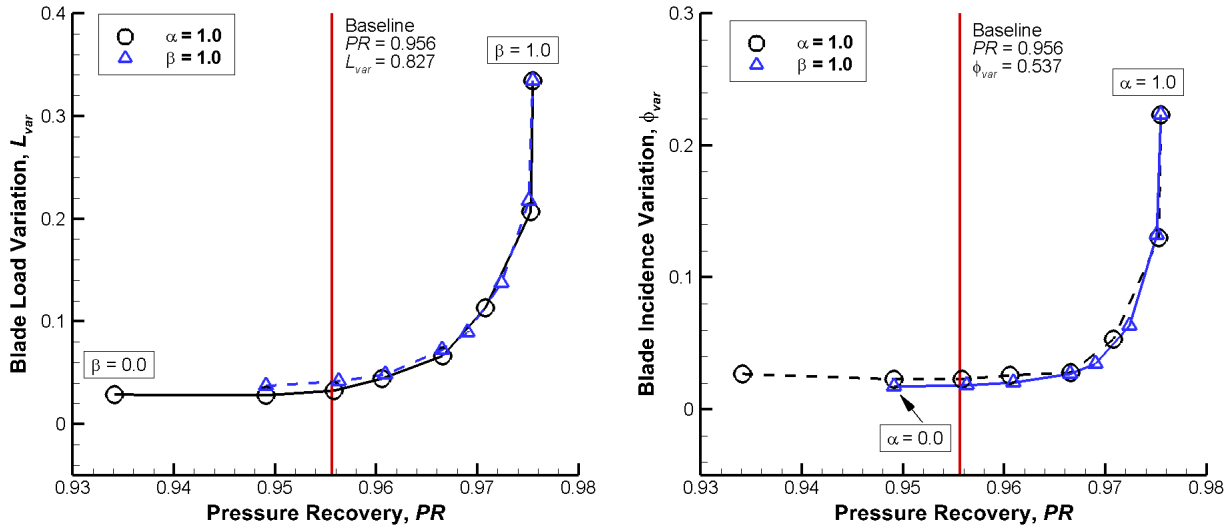

(a) (b)

Fig. 9 Pareto fronts for the tri-objective function with pressure recovery, blade load variation, and blade incidence variation at cruise and a fan speed of 35,000 RPM. Dashed lines indicate the performance of the other Pareto curve superimposed in the current domain.

Table 3 Cruise optimized performance showing multimodality with $\alpha = 0.90$, $\beta = 1.00$ (PR and ϕ_{var}).

Initial Condition	Objective Function	Pressure Recovery	Blade Incidence Variation
Optimized from baseline	0.8997	0.9666	0.0526
Optimized from $\alpha = 0.95$, $\beta = 1.00$ (PR and ϕ_{var})	0.8948	0.9690	0.0345

compared to the baseline. Values larger than $\alpha = 0.95$ (third blue triangle from the right) or smaller than $\alpha = 0.90$ (fourth blue triangle from the right) lead to diminishing returns in one objective with large penalties in the other. Thus, α values ranging from 0.90 to 0.95 are of interest since they provide substantial reduction in blade incidence variation while outperforming the baseline pressure recovery.

1. Comment on Multimodality

Evidence of multimodality with this aerodynamic shape optimization problem has been observed during the process of generating Pareto fronts for both composite objective functions. Since optimization runs are typically initiated from the baseline geometry, coupled with the fact that a gradient-based optimization algorithm is employed, in some cases this may lead to local minima that do not lie on the Pareto front. In order to increase the probability of finding a set of points that are consistent and located on the Pareto front, an optimization can be initiated from another optimized geometry that lies on the Pareto front if a multimodal result is encountered.

An example of multimodality is discussed in this section using the composite objective function with pressure recovery and blade incidence variation with $\alpha = 0.90$, $\beta = 1.00$. In Table 3, the optimizer found a solution with a better

Table 4 Optimized performance with different objective functions at cruise and a fan speed of 35,000 RPM.

Objective	Baseline	$\alpha = 1.00, \beta = 0.80$ (PR and D)	$\alpha = 0.95, \beta = 0.90$ (PR, D, and S)	$\alpha = 1.00, \beta = 0.95$ (PR and L_{var})	$\alpha = 0.90, \beta = 1.00$ (PR and ϕ_{var})
PR	0.956	0.962	0.961	0.971	0.969
$D \times 10^2$	5.205	0.239	1.360	1.015	0.697
$S \times 10^2$	9.569	20.07	1.455	14.64	16.74
L_{var}	0.827	0.061	0.259	0.113	0.089
ϕ_{var}	0.537	0.026	0.087	0.053	0.035

objective function value optimizing from another point on the Pareto curve than optimizing from the baseline. This initial point is $\alpha = 0.95$, $\beta = 1.00$, the third blue triangle from the right of Figure 9b. The fourth blue triangle from the right is $\alpha = 0.90$, $\beta = 1.00$, optimized from the third blue triangle. Furthermore, plotting the local minimum that was found by optimizing from the baseline places that point slightly above and to the left of the fourth blue triangle. Since this point is not located on the curve, it is not Pareto optimal, meaning that this point is dominated by another point on the Pareto curve that is better in both pressure recovery and blade incidence variation.

With gradient-based optimization, multimodality can be a serious concern, and it arose repeatedly during these S-duct optimization studies. One approach to addressing this concern is the gradient-based multistart approach of Chernukhin and Zingg [48, 49]. Although this approach is efficient, it is nevertheless expensive, and it was not used here. Instead, precautions were taken to avoid suboptimal local minima, as described above, and we have been careful to revisit any anomalous results obtained with a different initial geometry.

B. Single-Point Optimization

From the Pareto fronts in Figures 8 and 9, four pairs of relative weights were selected for further analysis. With these selected weights, the goal is to investigate how the optimizer manipulates the geometry and the resulting internal flow to achieve its objective at the cruise design point. The first pair is with $\alpha = 1.00$ and $\beta = 0.80$ for Eq. 11, which is a bi-objective function between pressure recovery and distortion that neglects swirl. With the same objective function formulation, the second pair is with $\alpha = 0.95$ and $\beta = 0.90$, which forms a tri-objective function that includes swirl. For the second composite objective function from Eq. 12, the third pair is with $\alpha = 1.00$ and $\beta = 0.95$, which is a bi-objective function between pressure recovery and blade load variation. The fourth pair is with $\alpha = 0.90$ and $\beta = 1.00$, which is another bi-objective function that considers pressure recovery and blade incidence variation. For a quantitative comparison of S-duct performance between the four composite objective functions, performance metrics are summarized in Table 4.

Contours of the total pressure, relative blade loading, and relative blade incidence at the outlet are compared in Figure 10. Streamlines near the boundary layer are compared in Figure 11 to visualize the extent of swirl in the flow. In

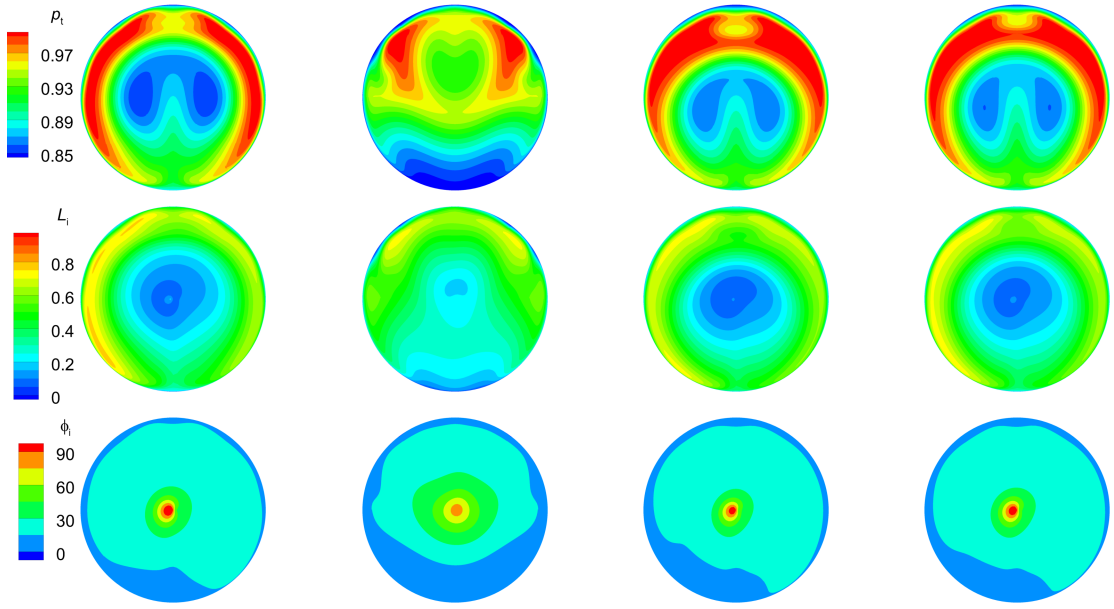

(a) $\alpha = 1.00$ and $\beta = 0.80$ (PR and D)

(b) $\alpha = 0.95$ and $\beta = 0.90$ (PR, D, and S)

(c) $\alpha = 1.00$ and $\beta = 0.95$ (PR and L_{var})

(d) $\alpha = 0.90$ and $\beta = 1.00$ (PR and ϕ_{var})

Fig. 10 Comparison of outlet contours at cruise and a fan speed of 35,000 RPM. Total pressure (top), relative blade loading (middle), and relative blade incidence in degrees (bottom).

Figure 10a, the optimizer was able to lift the low total pressure zone off the bottom surface and towards the outlet centre. The high total pressure region was dispersed circumferentially. Both of these redistributions lead to a reduction of the circumferential distortion of total pressure by 95% in relation to the baseline, while slightly improving pressure recovery by 0.6%. However, since this particular pair of $\alpha = 1.00$ and $\beta = 0.80$ relative weights removed swirl from the objective function, the optimizer was free to manipulate the secondary flows within the S-duct in order to achieve its objective.

Figure 11a illustrates how swirl was used to lift the low total pressure flow up towards the outlet centre. The increase in vortex strength to redistribute the total pressure doubles the amount of swirl in the flow compared to the baseline. Although the circumferential variations of blade loading and blade incidence are not included in this objective function, their distributions appear to be favourable. As expected, relative blade loading at the centre is low since the tangential velocity from blade rotation is small and the dominant component of the blade relative velocity comes from the incoming flow. Conversely, for the same reason the relative blade incidence angle is high at the centre and gradually reduces as the tangential velocity increases with radius. The asymmetry is a result of the addition and cancellation of tangential velocity vectors from the counter-clockwise fan rotation and the symmetric incoming flow. The swirl pair co-rotating with the fan induces a larger relative velocity which increases relative blade loading and reduces relative blade incidence, while the swirl pair counter-rotating with the fan achieves the opposite. Compared to the baseline, the blade load variation is reduced by 93%, with the blade incidence variation down by 95%.

In Figure 10b, the optimizer splits the high total pressure zone at the outlet with a lower total pressure core to create

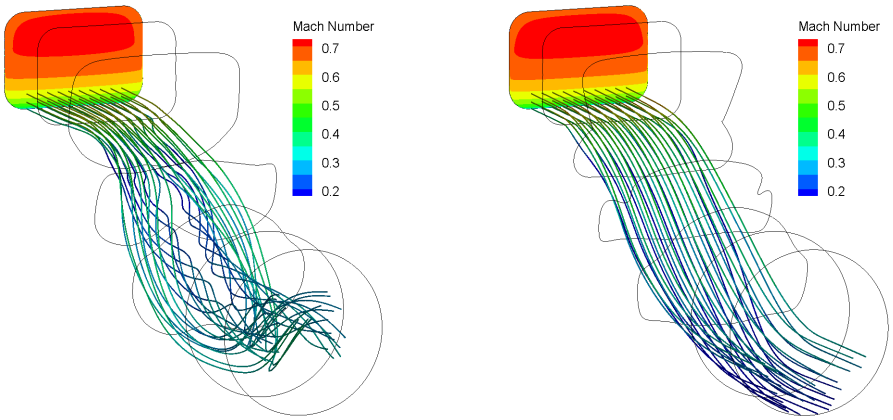

(a) α = 1.00 and β = 0.80 (PR and D)

(b) α = 0.95 and β = 0.90 (PR, D, and S)

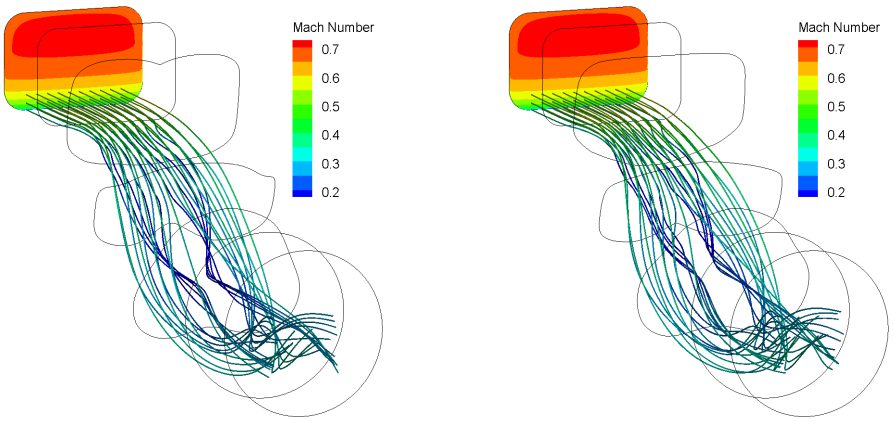

(c) α = 1.00 and β = 0.95 (PR and L_{var})

(d) α = 0.90 and β = 1.00 (PR and ϕ_{var})

Fig. 11 Comparison of streamlines at cruise and a fan speed of 35,000 RPM.

more circumferential uniformity. The bottom half of low total pressure has also been flattened in an effort to achieve the same objective of reducing distortion. Here, the distortion has been reduced by 74% compared to the baseline. The pressure recovery has maintained its slight improvement over the baseline by 0.5%. With swirl now included in the objective function, the optimizer is able to reduce swirl by 85% compared to the baseline. Referring to Figure 11b, the streamlines are shown to be relatively uniform with little flow angularity. In this case, the relative blade loading and relative blade incidence distributions are quite symmetric. This is due to the minimization of swirl, which removes the tangential component of the velocity vector for the incoming flow and leaves fan rotation as the only contributor to tangential velocity. The regions of high and low total pressure correspond well to the regions of high and low relative blade loading. With the low total pressure and relative blade load region present at the bottom surface near the tip, this may still be unfavourable for the fan to accept. This also results in a reduction of the relative blade incidence angle in the lower half of the outlet. The optimized flow indirectly reduces blade load and blade incidence variations by 69% and 84% of the baseline values, respectively.

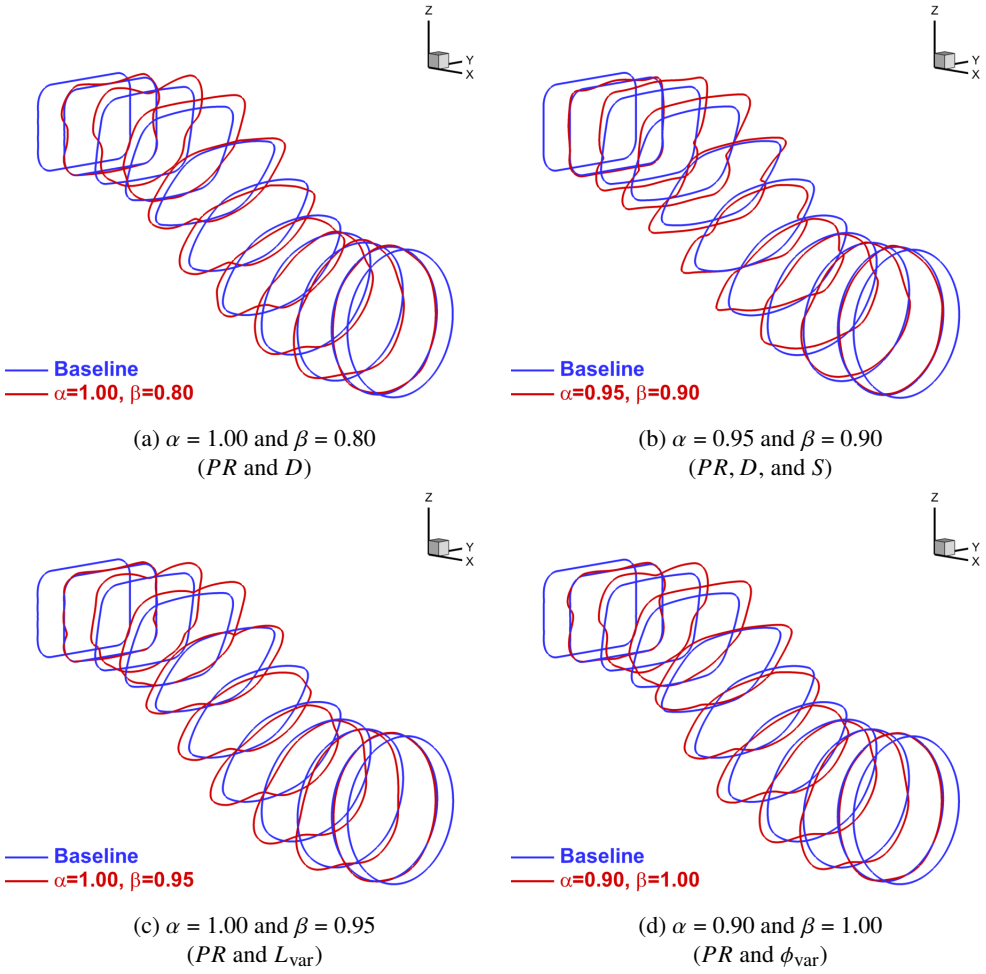

(a) α = 1.00 and β = 0.80 (PR and D)

(b) α = 0.95 and β = 0.90 (PR, D, and S)

(c) α = 1.00 and β = 0.95 (PR and L_{var})

(d) α = 0.90 and β = 1.00 (PR and ϕ_{var})

Fig. 12 Comparison of S-duct geometries at cruise and a fan speed of 35,000 RPM.

The total pressure, relative blade load, and relative blade incidence distributions in Figure 10c look similar to those in Figure 10d. In both cases, the pressure recovery is much improved, around 1.5% higher than that of the baseline. The distributions of relative blade loading and relative blade incidence are also favourable, and resemble those from Figure 10a. With α = 1.00 and β = 0.95, the objective function targets pressure recovery and blade load variation. With this set of weights, the optimizer reduces the baseline blade load variation by 87%. The blade incidence variation and distortion are reduced by 90% and 85% respectively, relative to the baseline. Using relative weights α = 0.90 and β = 1.00, the objective function focuses on pressure recovery and blade incidence variation. The blade incidence variation is reduced by 93% compared to the baseline. The blade load variation and distortion are reduced by 89% and 87% respectively, relative to the baseline. Figures 11c and 11d show that swirl is present in the flow, and the swirl value is roughly 1.5 times larger than the baseline.

The optimized shapes are compared against the baseline S-duct in Figure 12. Much like their outlet contours, the optimized geometries from Figures 12a, 12c, and 12d all look similar, indicating that the optimizer is using the same

flow features to achieve its objective. Two channels are formed on the bottom surface to accommodate the pair of vortices that shift the low total pressure region towards the centre. The upper surface near the first bend has been expanded to diffuse the flow before being accelerated by the second bend. This results in a more uniform flow velocity at the outlet plane.

In Figure 12b, the optimizer expands the upper wall at the first bend and forms a centrally-located separation bubble. The bulk flow coming from the top half bypasses this bubble and diverts to the upper corners. It then gets funnelled in a Y-pattern to split the high total pressure zone at the top half of the outlet, losing total pressure in the process, to create more circumferential uniformity. The bottom surface at the first bend has been lowered and the cross-sectional area expanded to reduce the adverse pressure gradient resulting from duct curvature, limit flow separation, and better diffuse the flow.

C. Multipoint Optimization and Assessment of a Morphing S-Duct

The current methodology is also extended to multipoint optimizations to investigate the robustness of the design. The goal is to evaluate the potential of S-duct morphing* for specific stages of flight against a static multipoint optimized design. Three cases are investigated and presented in this section. The first two cases consider the composite objective function with pressure recovery, distortion, and swirl, with relative weights $\alpha = 1.00$ and $\beta = 0.80$, and $\alpha = 0.95$ and $\beta = 0.90$ respectively. The third case uses the composite objective function with pressure recovery and blade incidence variation, with weights $\alpha = 0.90$ and $\beta = 1.00$.

Table 5 summarizes the quantitative results for single-point optimizations (i.e. different single-point designs optimized for each flight condition, which could be achieved through morphing) in the fifth column and both multipoint optimizations in the sixth and seventh columns respectively. The fourth column shows the performance of the cruise-optimized geometry and its performance at descent and climb. The comparisons are made against the baseline performance using $\alpha = 1.00$ and $\beta = 0.80$ that considers pressure recovery and distortion. All five objectives considered in this paper are presented, along with the individual components of the objective function at each operating condition. The percentage increase or decrease from the baseline is also provided in parentheses.

The general trends identified in Section V.B are consistent when extended to multiple operating conditions. With the specified weighting, the optimizer is able to maintain or slightly improve pressure recovery and significantly reduce distortion as intended. The exception here is the pressure recovery at multipoint 1, which is lower than the baseline. With swirl excluded from the objective function, the swirl generally doubles at every operating condition. This indicates that the optimizer is increasing the swirl as a method to reduce distortion and shift the low total pressure region toward the outlet centre. The single-point morphing objective function values outperform the multipoint values as expected, since the optimizer focuses on one condition specifically; however, the single-point morphing performance gain over the

*We do not consider how the morphing is achieved here, only its potential benefits with respect to the considered objectives of the S-duct design.

Table 5 Performance at multiple operating points optimizing with $\alpha = 1.00$ and $\beta = 0.80$ (PR and D).

Condition	Objective	Baseline	Cruise-Optimized	Single-point Morphing	Multipoint 1 (2:1:1)	Multipoint 2 (6:1:1)
Cruise	PR	0.956	0.962 (+0.6%)		0.951 (−0.5%)	0.960 (+0.4%)
	$D \times 10^2$	5.205	0.239 (−95%)		0.244 (−95%)	0.235 (−95%)
	$S \times 10^2$	9.570	20.07 (+110%)		21.54 (+125%)	20.71 (+116%)
	L_{var}	0.827	0.061 (−93%)		0.066 (−92%)	0.066 (−92%)
	ϕ_{var}	0.537	0.026 (−95%)		0.027 (−95%)	0.029 (−95%)
	$\mathcal{J}_{1,\text{ cruise}}$	1.000	0.804		0.814	0.806
Descent	PR	0.963	0.974 (+1.1%)	0.968 (+0.5%)	0.966 (+0.3%)	0.972 (+0.9%)
	$D \times 10^2$	3.696	0.466 (−87%)	0.167 (−95%)	0.199 (−95%)	0.351 (−91%)
	$S \times 10^2$	10.24	17.42 (+70%)	20.36 (+99%)	20.63 (+101%)	18.28 (+79%)
	L_{var}	0.620	0.078 (−87%)	0.047 (−92%)	0.054 (−91%)	0.070 (−89%)
	ϕ_{var}	0.588	0.040 (−93%)	0.031 (−95%)	0.027 (−95%)	0.034 (−94%)
	$\mathcal{J}_{1,\text{ descent}}$	1.000	0.818	0.807	0.810	0.813
Climb	PR	0.992	0.994 (+0.2%)	0.992 (0%)	0.992 (0%)	0.994 (+0.2%)
	$D \times 10^2$	0.203	0.050 (−75%)	0.014 (−93%)	0.018 (−91%)	0.043 (−79%)
	$S \times 10^2$	7.390	12.82 (+73%)	16.15 (+119%)	15.17 (+105%)	13.20 (+79%)
	L_{var}	0.178	0.035 (−80%)	0.017 (−90%)	0.021 (−88%)	0.033 (−81%)
	ϕ_{var}	0.085	0.017 (−80%)	0.004 (−95%)	0.007 (−92%)	0.015 (−82%)
	$\mathcal{J}_{1,\text{ climb}}$	1.000	0.848	0.815	0.818	0.841

multipoint optimizations is relatively small in general. In this case, between multipoint 1 and multipoint 2, placing more priority on the cruise condition with a 6:1:1 ratio does provide better cruise performance at the cost of reduced performance at descent and climb compared to multipoint 1. We also see that the cruise-optimized geometry provides better pressure recovery at descent and climb with a small trade-off in distortion compared to the single-point and multipoint cases. Although not directly accounted for in the objective function, blade load and blade incidence variation are generally much lower than the baseline across all three operating conditions.

A comparison of the outlet total pressure is shown in Figure 13. The contours are consistent at every operating condition with the region of low total pressure located near the centre and the high total pressure region being circumferentially redistributed. The low total pressure region is centred very well in the single-point column which corresponds to the lowest amount of circumferential distortion. The multipoint columns demonstrate that as less weighting is placed on descent and cruise, the low total pressure region becomes increasingly off-centred. As one would expect, this shows the trades between different flight conditions and their dependence on the multipoint weightings.

Comparing the single and multipoint columns for this particular set of weights, there is not a significant benefit for single-point morphing since the total pressure distributions are quite similar and both multi-point optimized geometries

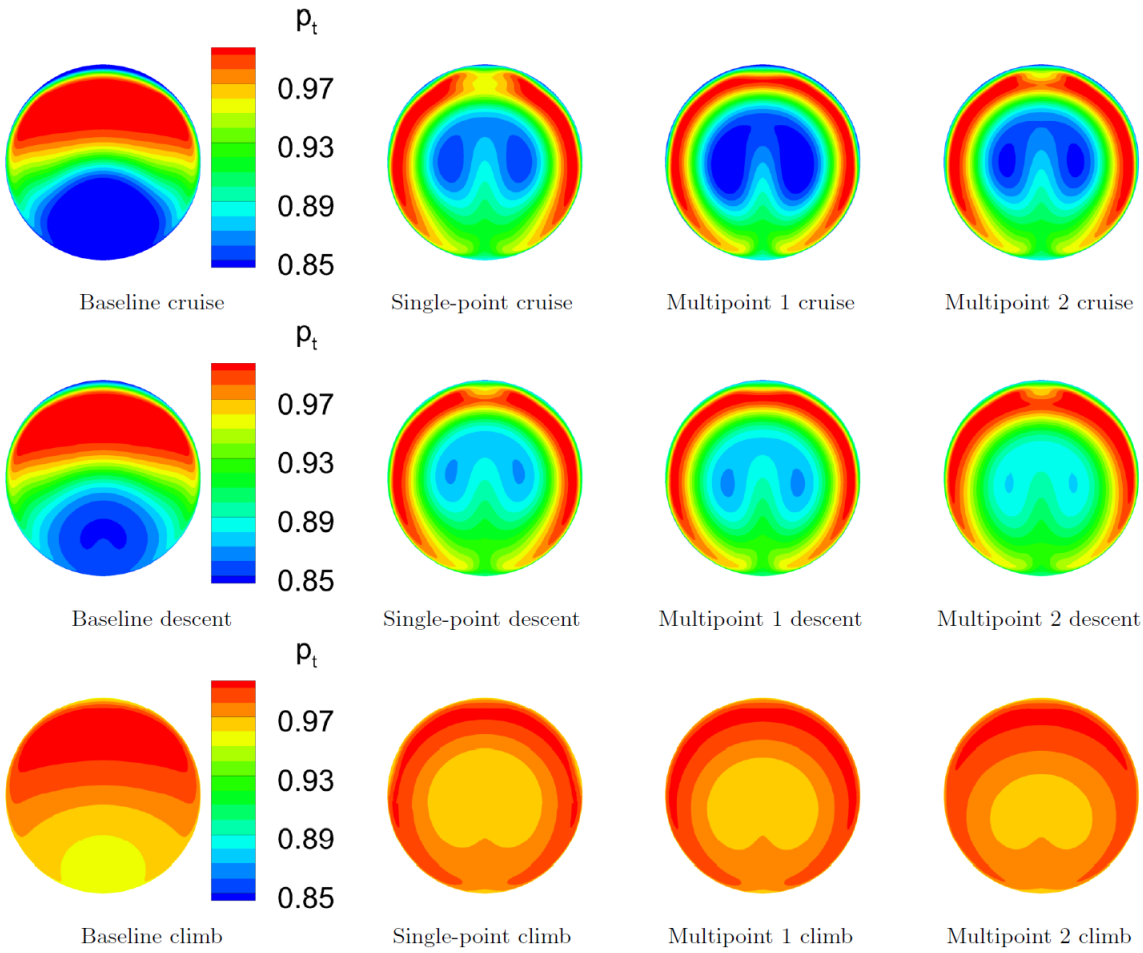

Fig. 13 Outlet total pressure contours optimizing with $\alpha = 1.00$, $\beta = 0.80$ (PR and D).

appear to mitigate distortion fairly well at every operating condition. In reference to climb, the multipoint 1 contour is similar in size and shape to single-point. A small trade-off in climb could be made in multipoint 2 if better total pressure recovery is desired at cruise.

Next, Table 6 is analyzed, which presents the results for $\alpha = 0.95$ and $\beta = 0.90$ to maximize pressure recovery and minimize distortion and swirl. From the table, it is evident that the optimizer was able to simultaneously improve all three objectives at every operating condition for both single-point and multipoint. The single-point objective function values at cruise and descent are superior to both multipoint values, but only slightly. The single-point climb performance for distortion and swirl are reduced even further by approximately a factor of two compared to the cruise-optimized geometry, multipoint 1, and multipoint 2. The discrepancy between the performance at multipoint 1 and multipoint 2 is also much smaller here. This may suggest that the geometry is insensitive to the multipoint weighting between operating conditions once swirl is included in the objective function. It can also be noted that blade load and blade incidence variation were reduced by generally similar percentages as the distortion.

Observing the outlet total pressure contours in Figure 14, the flow features appear to be quite similar. At cruise

Table 6 Performance at multiple operating points optimizing with $\alpha = 0.95$ and $\beta = 0.90$ (PR, D, and S).

Condition	Objective	Baseline	Cruise-Optimized	Single-point Morphing	Multipoint 1 (2:1:1)	Multipoint 2 (6:1:1)
Cruise	PR	0.956	0.961 (+0.5%)		0.960 (+0.4%)	0.961 (+0.5%)
	$D \times 10^2$	5.205	1.360 (−74%)		1.348 (−74%)	1.355 (−74%)
	$S \times 10^2$	9.569	1.455 (−85%)		1.701 (−82%)	1.468 (−85%)
	L_{var}	0.827	0.259 (−69%)		0.249 (−70%)	0.249 (−70%)
	ϕ_{var}	0.537	0.087 (−84%)		0.082 (−85%)	0.082 (−85%)
	$\mathcal{J}_{1,\text{ cruise}}$	1.000	0.882		0.884	0.883
Descent	PR	0.963	0.969 (+0.6%)	0.969 (+0.6%)	0.968 (+0.5%)	0.969 (+0.6%)
	$D \times 10^2$	3.696	1.322 (−64%)	1.288 (−65%)	1.303 (−65%)	1.316 (−64%)
	$S \times 10^2$	10.24	1.257 (−88%)	1.316 (−87%)	1.454 (−86%)	1.237 (−88%)
	L_{var}	0.620	0.259 (−58%)	0.243 (−61%)	0.247 (−60%)	0.248 (−60%)
	ϕ_{var}	0.588	0.144 (−76%)	0.129 (−78%)	0.133 (−77%)	0.135 (−77%)
	$\mathcal{J}_{1,\text{ descent}}$	1.000	0.891	0.890	0.892	0.891
Climb	PR	0.992	0.993 (+0.1%)	0.992 (0%)	0.993 (+0.1%)	0.993 (+0.1%)
	$D \times 10^2$	0.203	0.122 (−40%)	0.071 (−65%)	0.129 (−36%)	0.122 (−40%)
	$S \times 10^2$	7.390	5.267 (−29%)	1.231 (−83%)	4.660 (−37%)	5.183 (−30%)
	L_{var}	0.178	0.073 (−59%)	0.065 (−63%)	0.094 (−47%)	0.092 (−48%)
	ϕ_{var}	0.085	0.027 (−68%)	0.023 (−73%)	0.045 (−47%)	0.043 (−49%)
	$\mathcal{J}_{1,\text{ climb}}$	1.000	0.946	0.897	0.946	0.946

and descent, the high total pressure region in the top half is bisected with a core of lower pressure and the low total pressure region is flattened. Although not clear due to contour scaling, this type of flow redistribution does also appear during climb conditions. Therefore, this particular pair of values for α and β weights also show the lack of benefit for single-point morphing. At every operating condition, the presence of the low total pressure region remains in the lower portion of the outlet near the tip. This could potentially risk exposing the fan to unfavourable conditions that outweigh the concerns of minimizing swirl in the flow. It is difficult to determine which aspect is more detrimental to fan performance and further work is required to incorporate a method to evaluate fan performance.

The third and final case is summarized in Table 7 with $\alpha = 0.90$ and $\beta = 1.00$ for an objective function containing pressure recovery and blade incidence variation. Once again, the optimizer was able to significantly reduce the blade incidence variation while maintaining or improving the pressure recovery of the baseline. The blade load variation was not included in the objective function but was reduced by nearly the same percentage as the blade incidence variation. Distortion and swirl were also not optimized for, but the distortion was reduced as well. On the other hand, the swirl increased around one and a half times that of the baseline. This was also observed in the first set of cases optimizing for pressure recovery and distortion, and raises an important question about the impact of swirl on fan performance. The

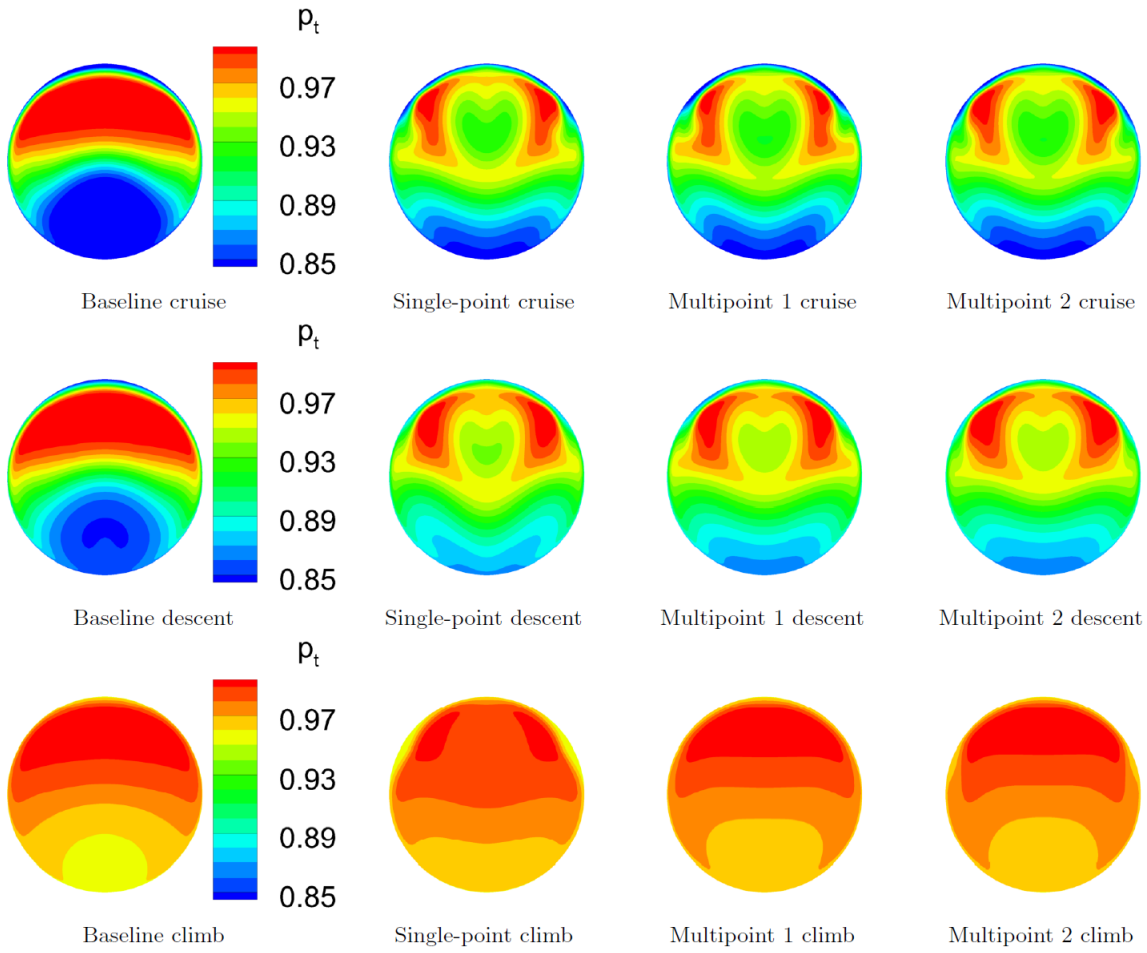

Fig. 14 Outlet total pressure contours optimizing with $\alpha = 0.95$, $\beta = 0.90$ (PR, D, and S).

single-point objective function values are only slightly better than the multipoint values at all three operating conditions. The cruise-optimized geometry at descent and climb have better pressure recoveries with a small trade-off in blade incidence variation compared to single-point and multipoint. Furthermore, there is not a large difference between the performance at multipoint 1 and multipoint 2 indicating that the larger prioritization of cruise in multipoint 2 did not improve cruise performance by a significant margin.

The outlet total pressure contours in Figure 15 display a similar flow field to that of Figure 13, where the high total pressure region is pulled around the outer circumference and the low total pressure region is lifted from the bottom surface and placed near the outlet centre. As such, the benefit of morphing is relatively small. A notable difference however is that the low total pressure region is divided into two distinct lobes on each side of the central symmetry plane. This once again demonstrates the trade-off between different flight conditions, in both on- and off- design points as well as the reliance on the multipoint weights to determine those trade-offs.

Table 7 Performance at multiple operating points optimizing with $\alpha = 0.90$ and $\beta = 1.00$ (PR and ϕ_{var}).

Condition	Objective	Baseline	Cruise-Optimized	Single-point Morphing	Multipoint 1 (2:1:1)	Multipoint 2 (6:1:1)
Cruise	PR	0.956	0.969 (+1.4%)		0.965 (+0.9%)	0.966 (+1.0%)
	L_{var}	0.827	0.089 (−89%)		0.109 (−87%)	0.115 (−86%)
	ϕ_{var}	0.537	0.035 (−93%)		0.041 (−92%)	0.045 (−92%)
	$D \times 10^2$	5.205	0.697 (−87%)		0.999 (−81%)	1.057 (−80%)
	$S \times 10^2$	9.569	16.74 (+75%)		16.31 (+70%)	15.66 (+64%)
	$\mathcal{J}_{2,\text{ cruise}}$	1.000	0.895		0.899	0.898
Descent	PR	0.963	0.977 (+1.5%)	0.974 (+1.1%)	0.974 (+1.1%)	0.975 (+1.2%)
	L_{var}	0.620	0.105 (−83%)	0.063 (−90%)	0.081 (−87%)	0.091 (−85%)
	ϕ_{var}	0.588	0.069 (−88%)	0.034 (−94%)	0.048 (−92%)	0.055 (−91%)
	$D \times 10^2$	3.696	0.845 (−77%)	0.445 (−88%)	0.688 (−81%)	0.765 (−79%)
	$S \times 10^2$	10.24	14.74 (+44%)	16.97 (+66%)	16.04 (+57%)	15.30 (+49%)
	$\mathcal{J}_{2,\text{ descent}}$	1.000	0.900	0.897	0.899	0.899
Climb	PR	0.992	0.995 (+0.3%)	0.992 (0%)	0.994 (+0.2%)	0.994 (+0.2%)
	L_{var}	0.178	0.034 (−81%)	0.011 (−94%)	0.021 (−88%)	0.025 (−86%)
	ϕ_{var}	0.085	0.017 (−80%)	0.006 (−93%)	0.008 (−91%)	0.012 (−86%)
	$D \times 10^2$	0.203	0.055 (−73%)	0.017 (−92%)	0.035 (−83%)	0.043 (−79%)
	$S \times 10^2$	7.390	11.95 (+62%)	16.99 (+130%)	13.82 (+87%)	12.85 (+74%)
	$\mathcal{J}_{2,\text{ climb}}$	1.000	0.918	0.904	0.908	0.912

VI. Conclusions

The current aerodynamic shape optimization framework has demonstrated its capability in handling the multi-objective nature of intake design for a BLI S-duct. Two composite objective functions pertaining to aerodynamic performance were investigated. One combines pressure recovery, distortion, and swirl, the other pressure recovery, blade load variation, and blade incidence variation. Pareto fronts are presented that demonstrate the trade-offs between competing objectives. Both single-point and multipoint optimizations are performed, considering cruise, descent, and climb conditions to consider the potential benefits of a morphing S-duct. From the results, the following conclusions can be drawn:

- Pressure recovery and fan-face distortion are competing objectives. In addition, distortion and swirl at the fan interface plane also compete. Distortion can be significantly reduced at the expense of a high amount of swirl. It is more challenging to reduce both distortion and swirl at the fan face. It is difficult to determine if swirl can be ignored; hence it is important to gain a better understanding of the importance of fan-face swirl on BLI efficiency.
- To improve performance in terms of fan-face distortion, blade load variation, or blade incidence variation, the optimizer increases the swirl in the flow to shift the low total pressure region towards the centre and

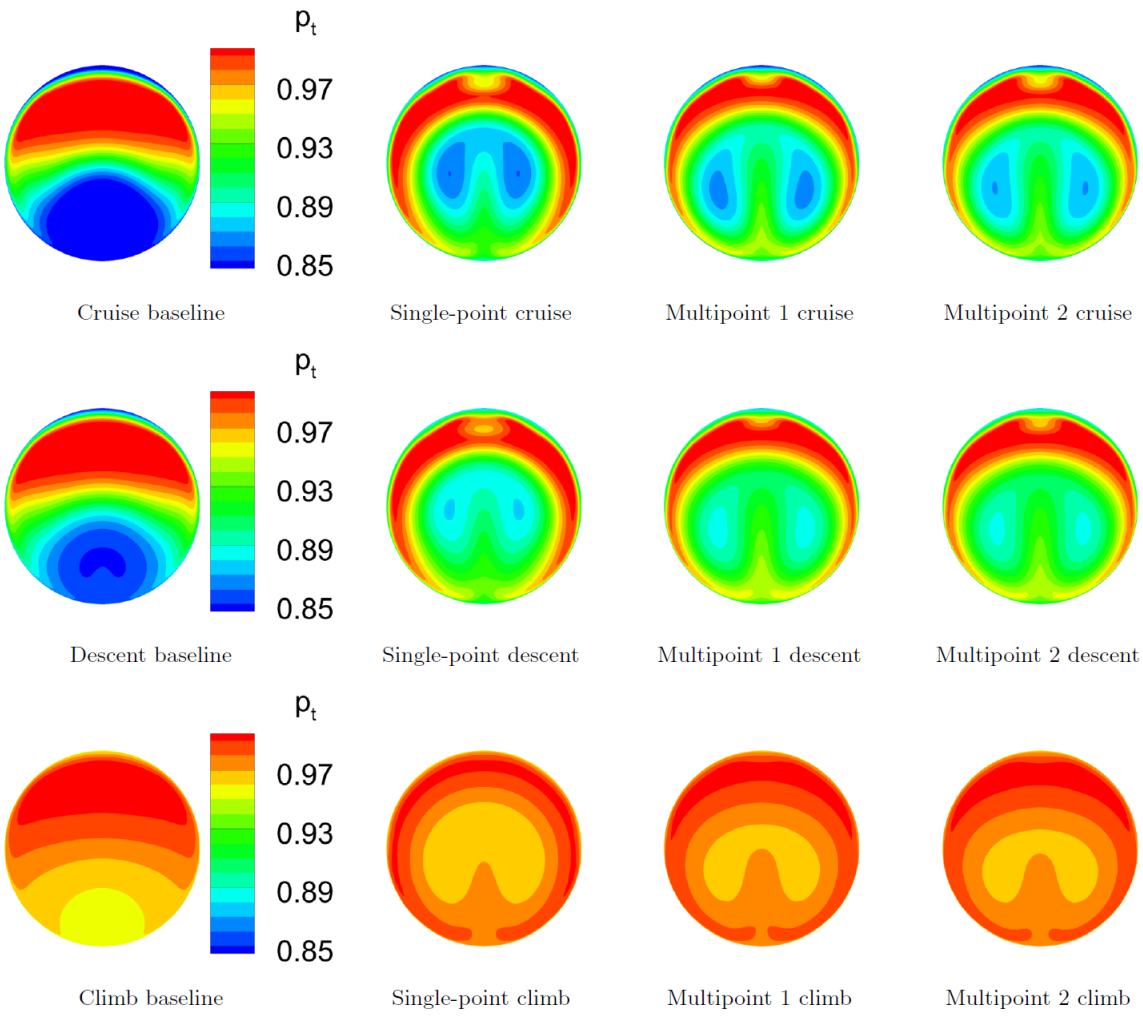

Fig. 15 Outlet total pressure contours optimizing with $\alpha = 0.90$, $\beta = 1.00$ (PR and ϕ_{var}).

circumferentially redistributes the high total pressure region. If swirl performance cannot be ignored, the low total pressure region is reduced in size but remains in the bottom portion of the fan interface plane near the fan tip.

- The results indicate that a single geometry determined from a multipoint optimization can provide good performance at multiple operating points for the range of conditions considered in this paper. This suggests that the benefits of S-duct morphing may be insignificant for this application.
- Multimodality was observed with this aerodynamic shape optimization problem. A gradient-based optimization algorithm was used for its computational efficiency; however, this method may be susceptible to finding multiple local minima. The authors were careful to revisit any anomalous optimization results with alternate initial geometries to locate the highest quality local optimum.

For future work, the consideration of additional geometric freedom in the optimization and experimental validation are topics of short term interest that would provide additional confidence in the results. The problem of multimodality

observed in this aerodynamic shape optimization problem could be mitigated with the application of the gradient-based multistart algorithm developed by Chernukhin and Zingg [48] to provide more extensive coverage of the design space with a range of initial geometries. In parallel, work is underway to investigate a single BLI S-duct installation on a regional-class blended wing body configuration. Additional work towards a predictive capability for fan performance under inlet distortion is also important to understand the relative importance of swirl on fan efficiency.

Acknowledgements

The authors gratefully acknowledge the financial assistance from the Ontario Graduate Scholarship program and the University of Toronto. Computations were performed on the General Purpose Cluster supercomputer at the SciNet HPS 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] Uranga, A., Drela, M., Hall, D. K., and Greitzer, E. M., "Analysis of the Aerodynamic Benefit from Boundary Layer Ingestion for Transport Aircraft," *AIAA Journal*, Vol. 56, No. 11, 2018, pp. 4271–4281. https://doi.org/10.2514/1.J056781.

[2] Betz, A., *Introduction to the Theory of Flow Machines*, 1st ed., Pergamon, New York, 1966, pp. 215, 217. https://doi.org/10.1016/C2013-0-05426-6.

[3] Küchemann, D., and Weber, J., *Aerodynamics of Propulsion*, McGraw-Hill, New York, 1953, pp. 205–209.

[4] Smith Jr., L. H., "Wake Ingestion Propulsion Benefit," *Journal of Propulsion and Power*, Vol. 9, No. 1, 1993, pp. 74–82. https://doi.org/10.2514/3.11487.

[5] Uranga, A., Drela, M., Greitzer, E. M., Hall, D. K., Titchener, N. A., Lieu, M. K., Siu, N. M., Casses, C., Huang, A. C., Gatlin, G. M., and Hannon, J. A., "Boundary Layer Ingestion Benefit of the D8 Transport Aircraft," *AIAA Journal*, Vol. 55, No. 11, 2017, pp. 3693–3708. https://doi.org/10.2514/1.J055755.

[6] Felder, J. L., Kim, H. D., and Brown, G. V., "Turboelectric Distributed Propulsion Engine Cycle Analysis for Hybrid-Wing-Body Aircraft," *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, AIAA 2009-1132, Orlando, Florida, January 2009. https://doi.org/10.2514/6.2009-1132.

[7] Gray, J. S., Mader, C. A., Kenway, G. K. W., and Martins, J. R. R. A., "Coupled Aeropropulsive Optimization of a Three-Dimensional Boundary-Layer Ingestion Propulsor Considering Inlet Distortion," *Journal of Aircraft*, Vol. 57, No. 6, 2020, pp. 1014–1025. https://doi.org/10.2514/1.C035845.

[8] Papadopoulos, F., Valakos, I., and Nikolos, I. K., "Design of an S-duct Intake for UAV Applications," *Aircraft Engineering and Aerospace Technology*, Vol. 84, No. 6, 2012, pp. 439–456. https://doi.org/10.1108/00022661211272990.

[9] Vaccaro, J. C., Elimelech, Y., Chen, Y., Sahni, O., Jansen, K., and Amitay, M., "Experimental and Numerical Investigation of Steady Blowing Flow Control Within a Compact Inlet Duct," *International Journal of Heat and Fluid Flow*, Vol. 54, August 2015, pp. 143–152. https://doi.org/10.1016/j.ijheatfluidflow.2015.05.011.

[10] Hui, Y., Feng, L., Yaoying, S., and Baigang, S., "Numerical Investigation of Electrohydrodynamic (EHD) Flow Control in an S-Shaped Duct," *Plasma Science and Technology*, Vol. 14, No. 10, 2012, pp. 897–904. https://doi.org/10.1088/1009-0630/14/10/08.

[11] Seddon, J., and Goldsmith, E. L., *Intake Aerodynamics*, AIAA Education Series, AIAA, Washington, D.C., 1985, pp. 12–14.

[12] Sóbestor, A., "Tradeoffs in Jet Inlet Design: A Historical Perspective," *Journal of Aircraft*, Vol. 44, No. 3, 2007, pp. 705–717. https://doi.org/10.2514/1.26830.

[13] Ramachandra, S. M., Sudhakar, K., Perumal, P. V. K., and Jayasimha, P., "Air-Inlet Engine Matching Problems Encountered in a Jet Trainer Re-Engining Program," *Journal of Aircraft*, Vol. 19, No. 8, 1982, pp. 609–614. https://doi.org/10.2514/3.57441.

[14] Hercock, R. G., *Effects on Intake Flow Distortion on Engine Stability*, Paper 20 of AGARD-CP-324, 1982.

[15] Ludwig, G., *Tomahawk Engine/Inlet Compatibility Study for F107-WR-400/402 Engines*, Williams International Report CMEP 5003-2025, 1989.

[16] Motycka, D. L., "Ground Vortex–Limit to Engine/Reverser Operation," *Journal of Engineering for Power*, Vol. 98, No. 2, 1976, pp. 258–263. https://doi.org/10.1115/1.3446157.

[17] Lotter, K. W., and Jörg, J., "The Effect of Intake Flow Disturbances on APU Compressor Blade High Cycle Fatigue in the Airbus A300," *Proceedings of the 13th Congress of ICAS*, ICAS-82-4.6.2, Seattle, Washington, August 1982, pp. 1072–1081.

[18] Guo, R. W., and Seddon, J., "An Investigation of the Swirl in an S-duct," *The Aeronautical Quarterly*, Vol. 33, No. 1, 1982, pp. 25–58. https://doi.org/10.1017/S0001925900009288.

[19] SAE International, "A Methodology for Assessing Inlet Swirl Distortion," Soc. of Automotive Engineers, AIR5686, March 2022. https://doi.org/10.4271/AIR5686.

[20] Fidalgo, V. J., Hall, C. A., and Colin, Y., "A Study of Fan-Distortion Interaction Within the NASA Rotor 67 Transonic Stage," *Journal of Turbomachinery*, Vol. 134, No. 5, 2012, Paper 051011. https://doi.org/10.1115/1.4003850.

[21] Cousins, W. T., "History, Philosophy, Physics, and Future Directions of Aircraft Propulsion System/Inlet Integration," *Proceedings of ASME Turbo Expo 2004: Power for Land, Sea, and Air*, GT2004-54210, Vienna, Austria, June 2004, pp. 305–320. https://doi.org/10.1115/GT2004-54210.

[22] McLelland, G., Macmanus, D. G., Zachos, P. K., Gil-Prieto, D., and Migliorini, M., "Influence of Upstream Total Pressure Profiles on S-Duct Intake Flow Distortion," *Journal of Propulsion and Power*, Vol. 36, No. 3, 2020, pp. 346–356. https://doi.org/10.2514/1.B37554.

[23] Owens, L. R., Allan, B. G., and Gorton, S. A., "Boundary-Layer-Ingesting Inlet Flow Control," *Journal of Aircraft*, Vol. 45, No. 4, 2008, pp. 1431–1440. https://doi.org/10.2514/1.36989.

[24] Gil-Prieto, D., Macmanus, D. G., Zachos, P. K., Tanguy, G., Wilson, F., and Chiereghin, N., "Delayed Detached-Eddy Simulation and Particle Image Velocimetry Investigation of S-Duct Flow Distortion," *AIAA Journal*, Vol. 55, No. 6, 2017, pp. 1893–1908. https://doi.org/10.2514/1.J055468.

[25] Rein, M., and Koch, S., "Experimental Study of Boundary-Layer Ingestion into a Diverterless S-Duct Intake," *AIAA Journal*, Vol. 53, No. 11, 2015, pp. 3487–3492. https://doi.org/10.2514/1.J053902.

[26] Tanguy, G., Macmanus, D. G., Zachos, P. K., Gil-Prieto, D., and Garnier, E., "Passive Flow Control Study in an S-Duct Using Stereo Particle Image Velocimetry," *AIAA Journal*, Vol. 55, No. 6, 2017, pp. 1862–1877. https://doi.org/10.2514/1.J055354.

[27] Connolly, B. J., Loth, E., and Smith, C. F., "Unsteady Separated Flows in an S-Duct and a Bifurcating Duct," *Journal of Aircraft*, Vol. 59, No. 1, 2022, pp. 47–57. https://doi.org/10.2514/1.C036383.

[28] Giuliani, J. E., and Chen, J. P., "Fan Response to Boundary-Layer Ingesting Inlet Distortions," *AIAA Journal*, Vol. 54, No. 10, 2016, pp. 3232–3243. https://doi.org/10.2514/1.J054762.

[29] Harrison, N. A., Anderson, J., Fleming, J. L., and Ng, W. F., "Active Flow Control of a Boundary Layer-Ingesting Serpentine Inlet Diffuser," *Journal of Aircraft*, Vol. 50, No. 1, 2013, pp. 262–271. https://doi.org/10.2514/1.C031818.

[30] Wojewodka, M. M., White, C., Shahpar, S., and Kontis, K., "A Review of Flow Control Techniques and Optimisation in S-Shaped Ducts," *International Journal of Heat and Fluid Flow*, Vol. 74, December 2018, pp. 223–235. https://doi.org/10.1016/j.ijheatfluidflow.2018.06.016.

[31] Willmer, A., Brown, T., and Goldsmith, E., *Effects of Intake Geometry on Circular Pitot Intake Performance at Zero and Low Forward Speeds*, AGARD-CP-301, 1981, Chap. 5.

[32] Whitelaw, J., and Yu, S. M., "Turbulent Flow Characteristics in an S-Shaped Diffusing Duct," *Flow Measurement and Instrumentation*, Vol. 4, No. 3, 1993, pp. 171–179. https://doi.org/10.1016/0955-5986(93)90053-L.

[33] Zhang, W., Knight, D., and Smith, D., "Automated Design of a Three-Dimensional Subsonic Diffuser," *Journal of Propulsion and Power*, Vol. 16, No. 6, 2000, pp. 1132–1140. https://doi.org/10.2514/2.5688.

[34] D'Ambros, A., Kipouros, T., Zachos, P., Savill, M., and Benini, E., "Computational Design Optimization for S-Ducts," *Designs*, Vol. 2, No. 4, 2018. https://doi.org/10.3390/designs2040036.

[35] Chiereghin, N., Guglielmi, L., Savill, A. M., Kipouros, T., Manca, E., Rigobello, A., Barison, M., and Benini, E., "Shape Optimization of a Curved Duct with Free Form Deformations," *23rd AIAA Computational Fluid Dynamics Conference*, AIAA 2017-4144, Denver, Colorado, June 2017. https://doi.org/10.2514/6.2017-4114.

[36] Sederberg, T. W., and Parry, S. R., "Free-Form Deformation of Solid Geometric Models," *ACM SIGGRAPH Computer Graphics*, Vol. 20, No. 4, 1986, pp. 151–160. https://doi.org/10.1145/15886.15903.

[37] Lee, B. J., and Kim, C., "Automated Design Methodology of Turbulent Internal Flow Using Discrete Adjoint Formulation," *Aerospace Science and Technology*, Vol. 11, No. 2, 2007, pp. 163–173. https://doi.org/10.1016/j.ast.2006.12.001.

[38] Koo, D., Zingg, D. W., Chishty, W. A., and abo el Ella, H. M., "Optimizing Intakes for Embedded Engines," *ISABE 2019*, ISABE-2019-24290, Canberra, Australia, September 2019.

[39] Asghar, A., Stowe, R. A., Allan, W. D. E., and Alexander, D., "Performance Evaluation of an S-Duct Diffuser of a Flight-Vehicle Inlet in High-Subsonic Flow," *Proceedings of ASME Turbo Expo 2015: Turbine Technical Conference and Exposition*, GT2015-43740, Montréal, Canada, June 2015. https://doi.org/10.1115/GT2015-43740.

[40] 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. https://doi.org/10.2514/1.J053575.

[41] Truong, A. H., Oldfield, C. A., and Zingg, D. W., "Mesh Movement for a Discrete-Adjoint Newton-Krylov Algorithm for Aerodynamic Optimization," *AIAA Journal*, Vol. 46, No. 7, 2008, pp. 1695–1704. https://doi.org/10.2514/1.33836.

[42] 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. 400–413. https://doi.org/10.2514/1.44033.

[43] Osusky, M., and Zingg, D. W., "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. https://doi.org/10.2514/1.J052487.

[44] Spalart, P. R., "Strategies for Turbulence Modeling and Simulations," *International Journal of Heat and Fluid Flow*, Vol. 21, No. 3, 2000, pp. 252–263. https://doi.org/10.1016/S0142-727X(00)00007-2.

[45] Del Rey Fernandez, D. C., Hicken, J. E., and Zingg, D. W., "Review of Summation-By-Parts Operators with Simultaneous Approximation Terms for the Numerical Solution of Partial Differential Equations," *Computers and Fluids*, Vol. 95, May 2014, pp. 171–196. https://doi.org/10.1016/j.compfluid.2014.02.016.

[46] Yaras, M. I., and Grosvenor, A. D., "Evaluation of One- and Two-Equation Low-Re Turbulence Models. Part II - Vortex-Generator Jet and Diffusing S-Duct Flows," *International Journal for Numerical Methods in Fluids*, Vol. 42, No. 12, 2003, pp. 1321–1343. https://doi.org/10.1002/fld.587.

[47] 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. https://doi.org/10.3166/remn.17.103-126.

[48] Chernukhin, O., and Zingg, D. W., "Multimodality and Global Optimization in Aerodynamic Design," *AIAA Journal*, Vol. 51, No. 6, 2013, pp. 1342–1354. https://doi.org/10.2514/1.J051835.

[49] Streuber, G. M., and Zingg, D. W., "Evaluating the Risk of Local Optima in Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 59, No. 1, 2021, pp. 75–87. https://doi.org/10.2514/1.J059826.

[50] Jameson, A., Martinelli, L., and Pierce, N. A., "Optimum Aerodynamic Design Using the Navier-Stokes Equations," *Theoretical and Computational Fluid Dynamics*, Vol. 10, January 1998, pp. 213–237. https://doi.org/10.1007/s001620050060.

[51] Rashad, R., and Zingg, D. W., "Aerodynamic Shape Optimization for Natural Laminar Flow Using a Discrete-Adjoint Approach," *AIAA Journal*, Vol. 54, No. 11, 2016, pp. 3321–3337. https://doi.org/10.2514/1.J054940.

[52] Osusky, L., Buckley, H. P., Reist, T. A., 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. https://doi.org/10.2514/1.J053457.

[53] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. https://doi.org/10.1137/S0036144504446096.

[54] Toot, L., "Vibration and Mission Simulation Testing on Engine 828 XF107-WR-400 Cruise Missile Engine," Tech. rep., Williams Research Corporation, Report No. 79-106-39, June 1981.